

Yes, we GAN: Applying Adversarial Techniques for Autonomous Driving

Michal Uříčář¹, Pavel Křížek¹, David Hurych¹, Ibrahim Sobh², Senthil Yogaman³ and Patrick Denny³

¹ Valeo DVS Prague, Czech Republic

² Valeo CDV AI Cairo, Egypt

³ Valeo Vision Systems, Ireland

ABSTRACT

Generative Adversarial Networks (GAN) have gained a lot of popularity from their introduction in 2014 till present. Research on GAN is rapidly growing and there are many variants of the original GAN focusing on various aspects of deep learning. GAN are perceived as the most impactful direction of machine learning in the last decade. This paper focuses on the application of GAN in autonomous driving including topics such as advanced data augmentation, loss function learning, semi-supervised learning, etc. We formalize and review key applications of adversarial techniques and discuss challenges and open problems to be addressed.

I INTRODUCTION

Autonomous driving is becoming a common feature in modern vehicles. Fully autonomous driving still remains a challenging task and there is a lot of active research to solve technical problems. Figure 2 illustrates the standard modules in an autonomous driving system. The first stage is perception using a suite of sensors. Common sensors which are already deployed are Cameras, Radar, Ultrasonics and Lidar. The first stage in the processing is perception where semantic objects like lanes/vehicles or geometric objects like freespace and generic obstacles are detected. They are then fused into a generic abstract representation typically a 2D or 3D map of objects with respect to the ego vehicle. A driving policy algorithm uses this map to decide a trajectory for the car to maneuver. The traditional approach is to have the modules independently designed but there are also attempts to do end to end learning.

Progress in Deep learning has accelerated the maturity of autonomous driving systems [1]. Deep learning models have become a standard in perception and gradually becoming competitive for other modules like fusion and driving policy. As shown in Figure 1, the commonly used deep learning models are Convolutional Neural Networks (CNN) and Recurrent Neural Networks (RNN). These are discriminative models which are trained for classification or regression problems. Discriminative models extract features which are sufficient to solve the classification problem and do not typically capture the complete information in the data. Generative models on the other hand try to capture the data distribution and hence form a more powerful representation. Generative Adversarial Networks (GAN) belong to this category and has become an effective generative model in various domains and tasks.

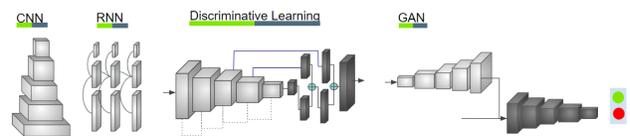


Figure 1: Commonly used deep learning models - CNN, RNN and GAN.

GAN have been progressing very rapidly and are seen as one of the most impactful models in the field of machine learning. The realism of generation of images using GAN has been impressive and recent results can capture subtle expressions. In spite of its popularity in the field of machine learning, it is relatively less explored for applications of autonomous driving. The main application where GAN are used in AD are image-to-image translation for style transfer from synthetic to realistic or transfer across different conditions of lighting, weather, etc. Autonomous driving systems have to be extremely robust and this requires training the model with all possible scenarios which can happen in real life. Collecting such a dataset is infeasible in practice and synthetic data simulators are commonly used to ameliorate this issue. Generative models, like GAN, provide a promising path to generate realistic datasets in this case.

The motivation of this paper is to survey adversarial techniques, distill the key ideas for applications of automated driving and provide a list of challenges and open problems from our industrial experience. The rest of the paper is organized into four sections. The first section describes the vanilla GAN as proposed in the original paper [2], prominent derivatives and recent advance. The second section discusses applications of GAN for autonomous driving. The third section discusses the main applications of GAN in autonomous driving. The fourth section summarizes results of our experiments on using GAN for soiling and adverse weather classification. The fifth section provides the discussion of the main challenges arising with GAN. Finally, the last section summarizes the paper and provides concluding remarks.

II OVERVIEW OF GAN

Generative Adversarial Networks were introduced in 2014 and were immediately recognized as a perspective direction of upcoming deep learning research, especially in domains such as unsupervised and semi-supervised learning, or advanced data augmentation. The idea behind GAN is very simple and intuitive, which might be one of the reasons for such a big popularity in the research community. Moreover, it seems that GAN are attracting in-

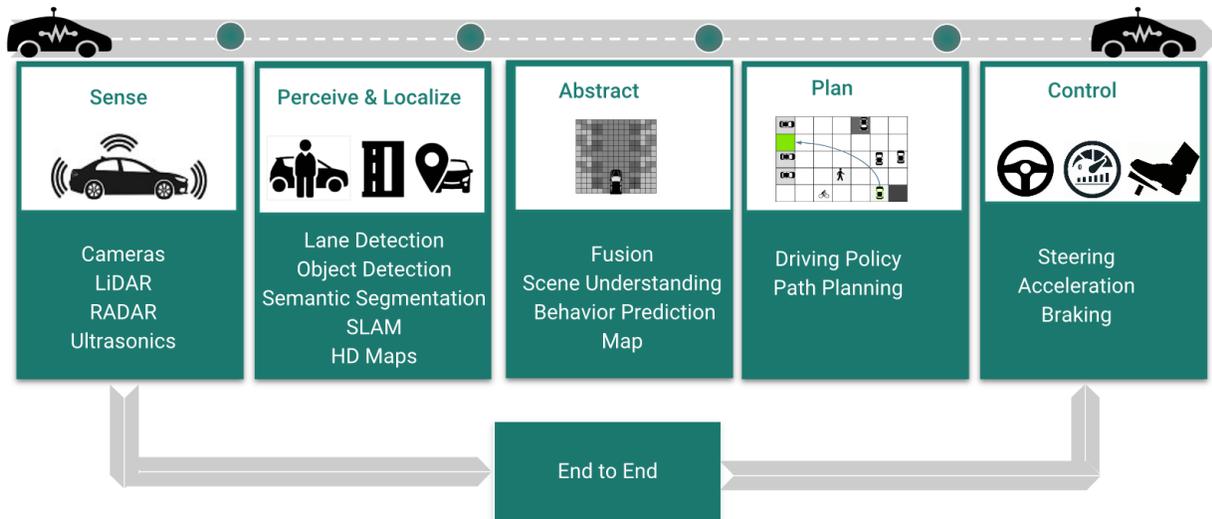


Figure 2: Modules in an autonomous driving pipeline.

terest of not only research pioneers, but as well the proof-seeking theoreticians, since both branches are getting more or less equal attention, and there appear both papers about new applications as well as theoretical improvements of the existing drawbacks of GAN. Therefore, we can see interesting applications such as unsupervised image-to-image translation, where for example images with horses are successfully transformed to realistically looking images with zebras, or a shot of Yosemite National Park taken in summer, transformed to the same scenery how it would probably look like in winter [3, 4]. The theory-deepening papers on the other hand deal with the stabilization of the complicated GAN learning, which often tends to get stuck in the mode collapse, i.e. the state when the discriminator is fooled to believe in unrealistic samples or degeneration of the generator, providing only a very limited set of samples [5, 6, 7]. The mode collapse problem arises when the generated data do not reflect the diversity of the underlying data distribution.

Vanilla GAN

Vanilla GAN [2] were introduced as a two-player minimax game, where each player is represented by a neural network. One network is a generator and the other one is a discriminator. The generator's task is to generate samples, which are as similar to the real data samples as possible, while the discriminator's task is to distinguish the real samples from the generated ones. See Figure 3 for the concept idea. The optimization task should end up at an equilibrium point, where the generator is able to generate samples, which the discriminator cannot distinguish from the real ones. In other words, the discriminator should output a probability equal to 0.5 for either of the real or generated input.

By using the neural networks, one of the biggest problems connected to generative modeling is mitigated—generative modeling in computer vision problems usually requires very complicated sampling functions, or complicated structures, and often only approximate inference computation is possible. However, using the neural networks is very simple, gradients are computed by a simple, yet effective, back-propagation algorithm, and the intuition behind GAN is that the complicated sampling function

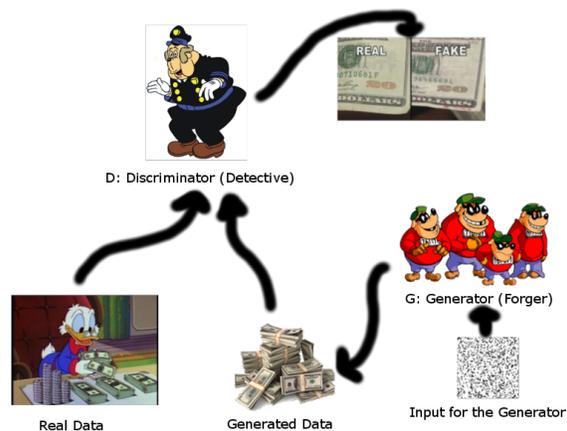


Figure 3: The conceptual idea of Generative Adversarial Networks. Generator (Forger) takes random noise on input and generates fake data. Discriminator (Detective) takes samples from both real and generated data and is trying to distinguish both categories, by reporting real or fake output for each sample.

can be constructed by implicit learning.

Let us describe GAN more formally. We assume for a simplicity that both models (generator and discriminator) are multi-layer perceptrons. The task is to learn the generator's distribution p_g , using the data samples x . Then, let us define a prior on the input noise variables $p_z(z)$. Using these, we can define a mapping to the data space as a differentiable function G , represented by a multi-layer perceptron with parameters θ_g : $G(z; \theta_g)$. Function G represents our generator. Next, we define a second function D , also as a multi-layer perceptron, with parameters θ_d : $D(x; \theta_d)$, the goal of which is to represent a probability that x comes from the data, and not from p_g . The function D represents our discriminator.

D is trained to maximize the probability of assigning the correct label to both the training samples and samples generated by G . Simultaneously, G is trained to minimize the $\log(1 - D(G(z)))$. This refers to what we have already mentioned in the beginning— D and G are playing a two-player minimax game defined as fol-

lows:

$$\min_G \max_D V(D, G) = \mathbb{E}_{x \sim p_{\text{data}}(x)} [\log D(x)] + \mathbb{E}_{z \sim p_z(z)} [\log (1 - D(G(z)))], \quad (1)$$

where $V(G, D)$ is a value function.

Because of the computational complexity of the discriminator maximization, the game has to be solved by an iterative numerical approach, which is summarized in Algorithm 1.

Algorithm 1 Mini-batch Stochastic Gradient Descent (SGD) training of GAN. The number of steps to apply to the discriminator, k , is a hyper-parameter. In the original paper [2], the authors used $k = 1$.

- 1: **for** number of training iterations **do**
- 2: **for** k steps **do**
- 3: Sample mini-batch of m noise samples $\{z^{(1)}, \dots, z^{(m)}\}$ from noise prior $p_g(z)$.
- 4: Sample mini-batch of m samples $\{x^{(1)}, \dots, x^{(m)}\}$ from data generating distribution $p_{\text{data}}(x)$.
- 5: Update the discriminator by ascending its stochastic gradient:

$$\nabla_{\theta_d} \frac{1}{m} \sum_{i=1}^m [\log D(x^{(i)}) + \log (1 - D(G(z^{(i)})))] .$$

- 6: **end for**
- 7: Sample mini-batch of m noise samples $\{z^{(1)}, \dots, z^{(m)}\}$ from noise prior $p_g(z)$.
- 8: Update the generator by descending its stochastic gradient:

$$\nabla_{\theta_g} \frac{1}{m} \sum_{i=1}^m \log (1 - D(G(z^{(i)}))) .$$

- 9: **end for**
-

We will here omit the theoretical results, such as the global optimality of $p_g = p_{\text{data}}$, or the proof of convergence of the Algorithm 1, those are described in the original paper [2] and are not particularly interesting for the potential applications of the GAN in various tasks. Instead, we will just summarize the advantages and disadvantages of GAN in the next paragraphs.

The main disadvantage of GAN framework is the absence of explicit representation of $p_g(x)$, and the requirement of simultaneous optimization of the discriminator D with the generator G . The authors say explicitly that “ G must not be trained too much without updating D ”.

There are several advantages. One was already mentioned—no inference computation is needed, only back-propagation is needed to compute the gradients. Moreover, a wide variety of functions can be incorporated in the model. The comparison to different approaches of generative modeling is summarized in [2, Table 2]. Another advantage is that GAN, in contrast to methods based on Markov Chains, can represent very sharp (even degenerate) distributions.

More information about the Vanilla GAN can be found in [8], which summarizes the NIPS 2016 tutorial on GAN.

Prominent Derivatives of GAN

In this section, we outline several prominent derivatives of GAN. While the list of all derivatives is nowadays containing tens of different instances, we will aim only at a few of them. For a reader’s convenience, we also list the references in the Table 1.

Conditional Generative Adversarial Nets (CGAN) were introduced in [9]. They are basically a straightforward extension of the Vanilla GAN [2], whereby conditioning the model on additional information one can influence the data generation process. This is crucial for example in the unsupervised learning scenario since we can exploit the conditioning to generate samples of a required label. In other words, we have a possibility to omit the costly annotation process.

The objective function of CGAN is updated by the conditional variable y (e.g. a class label) as follows:

$$\min_G \max_D V(D, G) = \mathbb{E}_{x \sim p_{\text{data}}(x)} [\log D(x|y)] + \mathbb{E}_{z \sim p_z(z)} [\log (1 - D(G(z|y)|y))] . \quad (2)$$

CGAN were shown to successfully generate MNIST [10] digits conditioned on class labels. However, the concept of CGAN was successfully used on more advanced tasks as well, e.g. the fast-converging conditional GAN (FC-GAN) [11] presents results on CIFAR10 dataset [12] as well.

Wasserstein GAN (WGAN) is another important derivative of GAN, they were first described in [6]. While most of the existing derivatives of GAN are trying to come up with novel or more interesting problems to solve, WGAN is pursuing a different goal—they focus solely on the learning of GAN.

In order to learn a probability distribution, one usually use the maximum likelihood estimation (MLE) over the training data. Let us denote the real data distribution (density) as P_r , and the parametrized density as P_θ . The MLE amounts to minimization of the Kullback-Leibler divergence $\text{KL}(P_r || P_\theta)$. This requires the model density P_θ to exist, this requirement is often broken. When we are dealing with distributions supported by low dimensional manifolds, the KL divergence is either not defined or is simply infinite.

The authors of [6] provide a theoretical analysis how different distances (such as total variation, KL divergence, Jensen-Shannon divergence, and earth-mover distance, a.k.a. Wasserstein distance) behave in the context of learning distributions. They show on a simple example (which, however, holds for the low dimensional manifolds as well) that except for the Wasserstein distance all other distances are not continuous, and therefore, do not provide clearly defined gradient everywhere. Moreover, Wasserstein distance, besides the guarantees of continuity and differentiability, is also the weakest¹ of the distances. There is also a proof that a small Wasserstein distance corresponds to a small distance in distributions.

¹We say that a distance d is weaker than distance d' if every sequence that converges under d' converges also under d .

The goal of WGAN is, therefore, to minimize the Wasserstein distance $W(P_r, P_\theta)$:

$$W(P_r, P_\theta) = \inf_{\gamma \in \Pi(P_r, P_\theta)} \mathbb{E}_{(x,y) \sim \gamma} [\|x - y\|]. \quad (3)$$

However, the infimum in (3) is highly intractable, so the Kantorovich-Rubinstein duality is used, which transforms (3) into

$$W(P_r, P_\theta) = \sup_{\|f\|_L \leq 1} \mathbb{E}_{x \sim P_r} [f(x)] - \mathbb{E}_{x \sim P_\theta} [f(x)], \quad (4)$$

where the supremum is taken over all 1-Lipschitz functions. The K -Lipschitz function is defined as follows: let d_X and d_Y be distance functions on spaces X and Y , respectively. Then a function $f: X \rightarrow Y$ is a K -Lipschitz if for all $x_1, x_2 \in X$,

$$d_Y(f(x_1), f(x_2)) \leq K d_X(x_1, x_2). \quad (5)$$

The gradient of (4) is defined as

$$\nabla_\theta W(P_r, P_\theta) = -\mathbb{E}_{z \sim p(z)} [\nabla_\theta f(g_\theta(z))]. \quad (6)$$

Finding f maximizing (4) is still intractable. We can, however, more easily approximate the problem—we can use a similar trick as in GAN—we can train a neural network parametrized with weights w lying in a compact space \mathcal{W} and then backpropagate through $\mathbb{E}_{z \sim p(z)} [\nabla_\theta f(g_\theta(z))]$. In order to have parameters w lying in a compact space, we can clamp the weights to a fixed box (e.g. $\mathcal{W} = [-0.01, 0.01]^l$) after each gradient update. The WGAN learning procedure is described in Algorithm 2.

Algorithm 2 WGAN algorithm. Authors of [6] propose to use the following default values: $\alpha = 0.00005, c = 0.01, m = 64, n_{\text{critic}} = 5$.

Require: α , the learning rate; c , the clipping parameter; m , the batch size; n_{critic} , the number of iterations of the critic per generator iteration.

Require: w_0 , initial critic parameters; θ_0 , initial generator's parameters.

```

1: while  $\theta$  has not converged do
2:   for  $t = 0, \dots, n_{\text{critic}}$  do
3:     Sample  $\{x^{(i)}\}_{i=1}^m \sim P_r$  a batch from the real data.
4:     Sample  $\{z^{(i)}\}_{i=1}^m \sim p(z)$  a batch of prior samples.
5:      $g_w \leftarrow \nabla_w \left[ \frac{1}{m} \sum_{i=1}^m f_w(x^{(i)}) - \frac{1}{m} \sum_{i=1}^m f_w(g_\theta(z^{(i)})) \right]$ 
6:      $w \leftarrow w + \alpha \cdot \text{RMSProp}(w, g_w)$ 
7:      $w \leftarrow \text{clip}(w, -c, c)$ 
8:   end for
9:   Sample  $\{z^{(i)}\}_{i=1}^m \sim p(z)$  a batch of prior samples.
10:   $g_\theta \leftarrow -\nabla_\theta \frac{1}{m} \sum_{i=1}^m f_w(g_\theta(z^{(i)}))$ 
11:   $\theta \leftarrow \theta - \alpha \cdot \text{RMSProp}(\theta, g_\theta)$ 
12: end while

```

The biggest advantage over vanilla GAN is that we can train the critic² till the optimality. This helps to prevent collapsing modes, which is a behavior frequently reported in vanilla GAN.

²Critic is a novel name for the discriminator from GAN. The reason for this renaming is, that unlike in GAN, in WGAN the output of this network is not a probability.

Improved WGAN introduced in [7] noticed that the quick and dirty solution of the 1-Lipschitzness by weight clipping can still lead to flawed learning, where only poor quality samples are generated, or where still convergence failure happens. In [7] a better solution is provided—the authors propose to penalize the norm of the gradient of the critic with respect to its input and prove that optimal 1-Lipschitz function's gradient for the optimization criterion in WGAN should have unit norm almost everywhere under P_r , and P_θ .

We list the improved WGAN in Algorithm 3, so the reader can compare both variants comfortably.

Algorithm 3 Improved WGAN algorithm. Authors of [7] propose to use the following default values: $\lambda = 10, \alpha = 0.0001, \beta_1 = 0, \beta_2 = 0.9, n_{\text{critic}} = 5$.

Require: The gradient penalty coefficient λ , the number of critic iterations per generator iteration n_{critic} , the batch size m , Adam [13] hyperparameters α, β_1, β_2 .

Require: Initial critic parameters w_0 , initial generator's parameters θ_0 .

```

1: while  $\theta$  has not converged do
2:   for  $t = 0, \dots, n_{\text{critic}}$  do
3:     for  $i = 0, \dots, m$  do
4:       Sample  $\{x^{(i)}\}_{i=1}^m \sim P_r$  a batch from the real data;
          $\{z^{(i)}\}_{i=1}^m \sim p(z)$  a batch of latent variables;  $\varepsilon \sim U[0, 1]$ .
5:        $\tilde{x} \leftarrow G_\theta(z)$ 
6:        $\hat{x} \leftarrow \varepsilon x + (1 - \varepsilon)\tilde{x}$ 
7:        $L^{(i)} \leftarrow D_w(\tilde{x}) - D_w(x) + \lambda (\|\nabla_{\hat{x}} D_w(\hat{x})\|_2 - 1)^2$ 
8:     end for
9:      $w \leftarrow \text{Adam}(\nabla_x \frac{1}{m} \sum_{i=1}^m L^{(i)}, w, \alpha, \beta_1, \beta_2)$ 
10:  end for
11:  Sample  $\{z^{(i)}\}_{i=1}^m \sim p(z)$  a batch of latent variables.
12:   $\theta \leftarrow \text{Adam}(\nabla_\theta \frac{1}{m} \sum_{i=1}^m -D_w(G_\theta(z)), \theta, \alpha, \beta_1, \beta_2)$ 
13: end while

```

Boundary-Seeking Generative Adversarial Networks (BGAN) were introduced in [14]. Similarly as WGAN, BGAN are focusing on GAN learning, they train a generator to match the target distribution that converges to the data distribution at the limit of a perfect discriminator. The interpretation of this is training the generator in order to produce the samples lying on the decision boundary of the current discriminator (which explains the name “Boundary-Seeking” GAN).

We omit the details of BGAN here and refer the reader to the original BGAN paper [14]. The authors claim that it produces similar results as WGAN, which is proved to be more stable against the mode collapse. We list BGAN here because of its main strength, which is a definition of a unified learning framework for both discrete and continuous variables, and interesting intuition about the samples lying on the decision boundary of the discriminator.

GAN: Recent Advances

In this section, some of the related and recent GAN advances are briefly described. BigGAN [15], a class-conditional image

Table 1: Summary of the prominent GAN derivatives.

GAN Derivative	References
General functionality	Vanilla GAN [2], Conditional GAN (CGAN) [9]
Stabilizing functionality	Wasserstein GAN (WGAN) [6], Improved WGAN (iWGAN) [7], Boundary-Seeking GAN (BGAN) [14]

synthesis approach, achieved a new level of performance for the large scale of ImageNet GAN models while allowing fine control over the trade-off between sample fidelity and variety by truncating the latent space. Moreover the authors presented an analysis of the training stability. A new training methodology for GAN is proposed in [16] where the generator and discriminator are made larger starting from a low resolution and adding new layers progressively, leading to generating a high resolution, high quality and fine details outputs. The authors described different implementation details to discourage the competition between the generator and discriminator. Furthermore, the proposed progressive training is showed to speed up and stabilize the training process. A novel generator architecture for GAN based on style transfer approaches is proposed in [17], leading to unsupervised separation of attributes and stochastic effects, that enabled fine control based on different levels of the generated images. Moreover, the experiments showed better interpolation properties. Self-Attention Generative Adversarial Network (SAGAN) [18] integrates a self-attention mechanism into GAN framework that allows attention-driven modeling for image generation where details are generated using cues from all feature locations. Moreover, it was shown that spectral normalization applied to the generator stabilizes the GAN training process.

III GAN APPLICATIONS FOR AUTONOMOUS DRIVING

In this section, we discuss the potential domains of application of GAN (in the, basically, arbitrary variant) with a focus on autonomous driving. Thanks to the increasing popularity of GAN, a lot of applications have been already identified. Most of them are, based on the GAN nature, related to the image-to-image translation or possibilities of the semi/unsupervised learning. We further discuss the following main categories: 1) advanced data augmentation (which is split into several subcategories, such as 2D & 3D synthesis, video synthesis, super resolution, or inpainting); 2) semi-supervised/unsupervised learning; 3) loss function learning; 4) adversarial training & testing. A short description of each of these categories follows. For the reader's convenience, we also summarized all key references in a Table 2.

Advanced Data Augmentation

Data Augmentation is a natural application of GAN. There are numerous papers dedicated to Image-to-Image translation on the top-level conferences from last few years ([3, 4, 19], to name a few). GAN create realistic looking images, automatic conversion from a black and white image to a colored one, areal image to map, edges to a photo-realistic images of the sketched objects, or even some advanced stuff like day to night or summer to winter, or context-aware object placement [20] (which is a very interesting

application for autonomous driving, giving us the possibility to enhance existing dataset by realistic looking images which we are lacking— e.g. small number of images with pedestrians, etc.). A natural extension of the Image-to-Image translation is a Video-to-Video translation, where the idea remains the same. However, the task is much more difficult thanks to the temporal information, which also has to remain consistent.

In [3], the authors formulate the problem of Image-to-Image translation as an instance of CGAN, where the objective loss, inspired by [21], is mixed with a traditional L1 loss. Note, that within such formulation the discriminator's task remains unchanged. However, the generator's task, besides fooling the discriminator, is to produce samples near the ground-truth in the L1 sense. The authors argue, that using L1 leads to less blurring, compared to the L2. The \mathcal{L}_{L1} term is formulated as follows:

$$\mathcal{L}_{L1} = \mathbb{E}_{x,y \sim p_{\text{data}}(x,y), z \sim p_c(z)} [\|y - G(x,y)\|_1] . \quad (7)$$

If we denote the equation (2) as $\mathcal{L}_{\text{CGAN}}(G,D)$, the final objective of [3] can be expressed as:

$$G^* = \arg \min_G \max_D \mathcal{L}_{\text{CGAN}}(G,D) + \lambda \mathcal{L}_{L1}(G) , \quad (8)$$

where λ tells how much we want the \mathcal{L}_{L1} loss influence the objective. Even though the results look impressive at the first sight, we should point out here, that there are certain artifacts appearing basically in all of the examples.

The authors of [3] continued working on the Image-to-Image translation topic, and in a follow-up paper [4], they move further, by introducing an algorithm for an unpaired Image-to-Image translation. Briefly, the authors are learning a mapping $G: X \rightarrow Y$, such that the distribution of images from $G(X)$ is indistinguishable from the distribution Y . Because such mapping is highly under-constrained, they couple it with an inverse mapping $F: Y \rightarrow X$ and introduce a cycle consistent loss enforcing $F(G(X)) \approx X$, and vice versa. The proposed algorithm is called CycleGAN.

Because of the inverse mapping, the CycleGAN uses a slightly different definition of \mathcal{L}_{GAN} than what we introduced here:

$$\begin{aligned} \mathcal{L}_{\text{GAN}}(G,D_Y,X,Y) &= \mathbb{E}_{y \sim p_{\text{data}}(y)} [\log D_Y(y)] \\ &+ \mathbb{E}_{x \sim p_{\text{data}}(x)} [\log(1 - D_Y(G(x)))] , \end{aligned} \quad (9)$$

The cycle consistency loss is defined as follows:

$$\begin{aligned} \mathcal{L}_{\text{cyc}}(G,F) &= \mathbb{E}_{x \sim p_{\text{data}}(x)} [\|F(G(x)) - x\|_1] \\ &+ \mathbb{E}_{y \sim p_{\text{data}}(y)} [\|G(F(y)) - y\|_1] . \end{aligned} \quad (10)$$

The full objective of CycleGAN is defined as

$$\begin{aligned} \mathcal{L}(G,F,D_X,D_Y) &= \mathcal{L}_{\text{GAN}}(G,D_Y,X,Y) \\ &+ \mathcal{L}_{\text{GAN}}(F,D_X,Y,X) \\ &+ \lambda \mathcal{L}_{\text{cyc}}(G,F) , \end{aligned} \quad (11)$$

where again λ controls the relative importance of the objective components. The full optimization task is simply

$$G^*, F^* = \arg \min_{G,F} \max_{D_X,D_Y} \mathcal{L}(G,F,D_X,D_Y) . \quad (12)$$

As the main benefit compared to the previous work [3], the authors consider the ability of CycleGAN to operate without a specific supervision. However, the main drawback remains unsolved—the generated images, after a careful inspection, show the same artifacts as the previous work. On the other hand, CycleGAN showed even more advanced examples of possible data augmentation (changing of an object class, like apples \leftrightarrow oranges, or zebras \leftrightarrow horses).

Since both [3], and [4] are based on GAN, they also require some compromises in the network architecture for both D , and G . It might be interesting to try to adopt the ideas and using WGAN instead.

Slightly different approach than CycleGAN is described in [19]. They also deal with an UNsupervised Image-to-image Translation (UNIT). However, instead of explicitly modeling the cycle consistency, they use an assumption of shared-latent space, which, as the authors show, implies the cycle consistency constraint. The proposed UNIT framework is based on GAN and variational autoencoders (VAE). We omit the technical details about UNIT here and refer the reader to the original paper [19]. The results are similar to those presented by CycleGAN, including the artifacts appearing in the generated images.

In [22], an algorithm for conditional image synthesis with auxiliary classifier GAN (AC-GAN) is introduced. The authors claim that even though structurally AC-GAN are not much different from the existing models, they seem to stabilize the training. Despite the AC-GAN formulation, the paper focuses on measuring the extent to which a model make use of its given output resolution, and on measuring the perceptual variability of samples generated by the model. In these areas, this work can be proclaimed as a pioneer.

Synthesis: Data Synthesis using GAN can be exploited for different autonomous driving applications. For example, data augmentation to enable better generalization over different weather and lighting conditions conducting visual perception by generating abstracted views such as semantic segmentation of input camera frames. In addition to sensor correcting such as fixing noisy inputs and sensor modeling. Synthesis can be achieved spatially in two and three dimensional spaces, in addition to spatio-temporal spaces such as videos.

2D Synthesis: GAN have been proposed as a generative framework that maps random noise to synthetic, realistically looking images following the training data distribution. In image-to-image translation, conditional GAN are adopted to enable the generative framework to condition both the generator and the discriminator on prior knowledge. In this case, the model is trained to map from images in a source domain to images in a target domain. Image-to-Image translation can be approached based on two main directions: 1) paired or unpaired; 2) unimodal or multimodal. In the unimodal paired image translation, for example Pix2Pix [3], SRGAN [23], the model learns to map images where the training data is organized in pairs of input and output samples. In many cases, the paired training data could not be available. In the unimodal unpaired approach, for example CycleGAN [4], DiscoGAN [24], StarGAN [25], UNIT [26], the image translation is conducted on unpaired data from two domains, where it learns a mapping between the two

Table 2: Summary of applications of GAN for autonomous driving.

GAN Applications	References
2D Synthesis	Pix2Pix [3], SRGAN [23], CycleGAN [4], DiscoGAN [24], StarGAN [25], UNIT [26], Pix2PixHD [27], BicycleGAN [28], MUNIT [29], Augmented GAN [30]
3D Synthesis	3D-GAN[31], PrGAN [32], PC-GAN [33]
Video Synthesis	video-to-video [34], TGAN [35], [36]
Domain Adaptation	Pixel level [37], GraspGAN [38], [39], [40]
Object Detection	SeGAN [41], [42], Perceptual-GAN [43]
Super Resolution	SRGAN [23]
Inpainting	[44], [45], [46]
Advanced Data Augmentation	Context-aware Synthesis and Placement of Object Instances [20]

domains without supervision. In multimodal image translation, it is possible to generate several images of different styles based on a single source image. Multimodal translation can be paired Pix2PixHD [27], BicycleGAN [28], or unpaired MUNIT [29], Augmented GAN [30].

3D Synthesis: Recently, one of the essential sensors used for automated driving is LiDAR, mainly because of its physical ability to perceive accurate depth and to produce 3D point clouds, regardless of lighting conditions. Most of GAN approaches are not applicable to 3D point clouds, however Point Cloud GAN (PC-GAN) [33] proposed a two fold modification to GAN algorithm for learning to generate point clouds. Moreover, studies were provided for transforming images into point clouds. 3D-GAN framework is introduced in [31] to map from a low-dimensional probabilistic space to the space of 3D objects. PrGAN [32], investigated the task of generating a distribution over 3D structures given 2D views of multiple objects taken from unknown viewpoints. This approach produces 3D shapes of comparable quality to GAN trained on 3D data, allows generating new views from an input image in an unsupervised manner.

Video Synthesis: The goal of video-to-video synthesis is to learn a mapping from an input video to a realistic output video. Simply applying image to image synthesis results in temporally incoherent videos. Recently, a video synthesis approach based on GAN framework is proposed in [34] where a spatio-temporal adversarial objective is used to synthesis 2k resolution videos of street scenes up to 30 seconds long. This allows developers and artists to create new interactive 3D virtual worlds for different domains including automotive domain. Temporal GAN (TGAN) [35], on the other hand, learns a semantic representation of unlabeled videos and generates videos, using a temporal generator and an image generator. Earlier, a GAN network for video with a spatio-temporal convolutional architecture is proposed in [36] where the foreground of the scene is separated from the background, generating small one second videos.

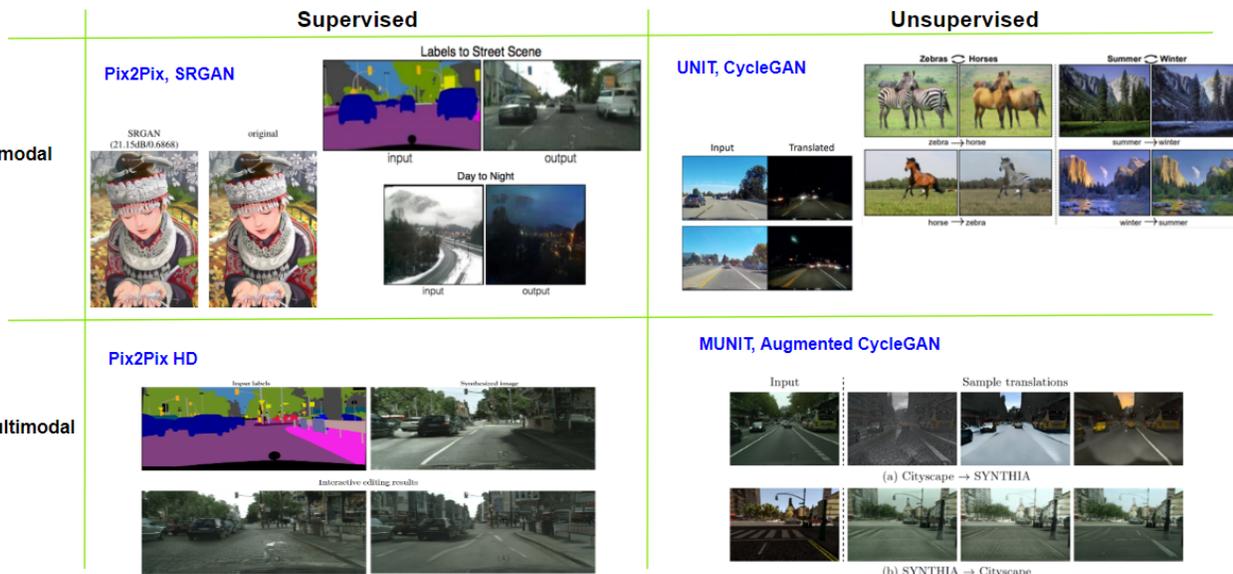


Figure 4: Successful application of GAN for autonomous driving - Image to Image Translation

Domain adaptation from simulation to real: Autonomous driving systems usually require collecting and annotating a lot of training data. On the other hand, using simulated environments enables much easier collection, but models trained on simulated environments often fail to generalize on real environments. Domain adaptation allows a machine learning model trained on samples from a source domain to generalize to a target domain. GAN based Pixel-Level domain adaptation method is proposed in [37] where the adaptation process showed to produce plausible samples and to generalize well to object classes unseen during the training. GraspGAN [38] extended the pixel-level domain adaptation to reduce the number of real world samples needed by up to 50 times for vision-based grasping system. Reinforcement Learning for Autonomous Driving model trained in virtual environment is shown to perform well in real environment [39]. Two image-to-image translation networks are used. The first network translates virtual images to their segmentation, the second network translates segmented images into their realistic counterpart. Accordingly, the driving policy can be easily adapted to real environment. A method for transferring a vision-based lane following driving policy from a simulated to a real environment is presented in [40] where a model for end-to-end driving is constructed by learning to translate between simulated and real images, jointly learning a control policy from this common latent space using labels from an expert driver in the simulated environment. It was shown that the proposed system is capable of leveraging simulation to learn a driving policy to directly transfer to real world scenarios.

Object Detection: In real life situations, especially for autonomous driving, objects often occlude each other and inferring the occluded objects is essential for scene understanding and taking decisions. SeGAN [41] is an approach for both segmentation and generation of the occluded parts of objects, where the proposed network has three parts: segmentor, generator, and discriminator. On the other hand, some occlusions are very rare in the training data. Learning a model invariant to such

occurrences is proposed in [42] where an adversarial network learns to generate challenging occlusions and deformations examples. Detection of small objects is often challenging task. Perceptual-GAN [43] model narrows representation difference of small and large objects, where the generator learns to transfer the small objects representations large ones to fool the competing discriminator.

Super Resolution: In autonomous driving domain, low resolution sensors may be used. Generating the corresponding high resolution representations could enable and enhance the systems that were trained on high resolution inputs. However, estimating a high resolution representation from a low resolution counterpart is a highly challenging task. SRGAN [23] is a generative adversarial network for image super-resolution framework able to infer photo-realistic natural images for 4x upscaling factors. A perceptual loss function is proposed consisting of both an adversarial loss for natural output, and a content loss for perceptual similarity.

Inpainting: In real life autonomous driving situations, sensors may read noisy data or may suffer from failures causing incomplete readings. Inpainting can provide a solution. Globally and locally consistent image completion approach is proposed in [44] where global and local context discriminators are trained to distinguish real images from completed ones. The global discriminator is used for the entire coherent of the generated image and the local discriminator ensures the local consistency. The approach showed to naturally complete the missing fragments. Inpainting is a challenging task especially for large missing parts. [45] proposed a method for semantic image inpainting by conditioning on the available data making the inference possible irrespective of how the missing parts are structured. [46] introduced context encoders that predict missing parts of a scene from their surroundings. Using an adversarial loss is found to produce sharper results. Moreover, the network learns a representation that captures both the appearance and the semantics of the scene.

Semi-supervised/Unsupervised Learning

Unsupervised learning is the holy grail of machine learning. It is the hardest scenario of machine learning, as it requires just a bunch of unannotated data, and the output is a learned meaningful model, which pursues a certain task (which can be inferred from the data).

Recently, one of the most popular approaches to unsupervised learning are Variational Autoencoders (VAE). They are built on top of the standard function approximators (neural networks) and can be trained by stochastic gradient descent. Our description of VAE framework follows [47].

In principle, VAE is aiming to maximize the probability of each x in the training set under the generative process, according to

$$p(x) = \int p(x|z; \theta)p(z)dz \quad (13)$$

where $z \in \mathcal{Z}$ is a *latent* variable specifying the class of the generated object. The intuition behind this framework is nothing else than maximum likelihood principle—if the model is likely to produce training set samples, then it is also likely to produce similar samples, and unlikely to produce dissimilar ones. VAE approximately maximizes the probability function. The name “autoencoders” comes from the training objective, which, derived from this setup, have an encoder and a decoder, and resembles a traditional autoencoder. The benefit of VAE is that we can sample from $p(x)$ (without performing Markov Chain Monte Carlo).

To solve the maximization, VAE have to deal with the problem of how to define latent variables z (i.e. decide what information they represent), as well as how to deal with the integral over z .

VAE assume that there is no simple interpretation of dimensions of z , rather they claim that samples of z might be drawn from a simple distribution, i.e. $\mathcal{N}(0; \mathcal{I})$, where \mathcal{I} is the identity matrix (notice, that any distribution in d dimensions can be generated by taking a set of d variables that are normally distributed, and map them through a sufficiently complicated function). Conceptually, the approximation of $p(x)$ is straightforward one just needs to sample a large number of samples $\{z_1, \dots, z_n\}$, and then compute $p(x) \approx \frac{1}{n} \sum_{i=1}^n p(x|z_i)$. A problem arises in high dimensional spaces because then n might be extremely large, to get an accurate estimate of $p(x)$.

The key idea behind VAE is to attempt to sample values of z that are likely to produce X and compute $P(X)$ just from those. For this purpose, let us define a new function $Q(z|X)$, which takes a value of X and give us a distribution over z values, that are likely to produce X . Hopefully, the space of z values that are likely under Q will be much smaller than the space of all z 's, that are likely under the prior $P(z)$. This lets us compute $\mathbb{E}_{z \sim Q} P(X|z)$. However, if z is sampled from an arbitrary distribution with a probability density function $Q(z)$ (i.e. which is not necessarily $\mathcal{N}(0, I)$), we need to relate $\mathbb{E}_{z \sim Q} P(X|z)$ with $P(X)$. By defining a Kullback-Leibler divergence between $P(z|X)$ and $Q(z)$, applying a Bayes rule to $P(z|X)$, and some basic algebra, we finally arrive at a core equation of VAE:

$$\log P(X) - \text{KL}[Q(z|X)||P(z|X)] = \mathbb{E}_{z \sim Q} [\log P(X|z)] - \text{KL}[Q(z|X)||P(z)]. \quad (14)$$

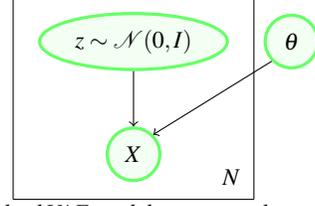


Figure 5: Standard VAE model represented as a graphical model. The rectangle is a “plate notation”, meaning that we can sample from z and X N times, while the model parameters θ remain fixed.

The left hand side is the quantity we want to maximize ($\log P(X)$ plus an error term making Q produce z 's that can reproduce a given X ; this term should be small if Q is high-capacity). The right hand side is optimizable via stochastic gradient descent, given the right choice of Q . Note, that the right hand side already looks like an autoencoder, since Q is “encoding” X into z and P is “decoding” it to reconstruct X .

In (14), we are maximizing $\log P(X)$, while simultaneously minimizing $\text{KL}[Q(z|X)||P(z|X)]$. The second term on the left-hand side is pulling $Q(z|X)$ to match $P(z|X)$. With an assumption of an arbitrarily high-capacity model for $Q(z|X)$ it will actually match $P(z|X)$, in which case the KL divergence term will be zero, and we will be optimizing $\log P(X)$ directly.

In order to make VAE work, it is essential to drive Q to produce codes for X , that P can reliably decode. See Figure 6 left, for a different view on the problem. The forward pass of this network works fine. However, the backpropagation of the error through a layer sampling z from $Q(z|X)$ is not possible, since this is a non-continuous operation and therefore has no gradient. The solution is a so-called “reparametrization trick” [13], which moves the sampling to an input layer. Figure 6 right depicts the reparametrization scheme. We should point out that this trick works only if sampling from $Q(z|X)$ is possible through evaluation of a continuous function $h(\eta, X)$ in X , where η is a noise from a distribution that is not learned. This basically means that $Q(z|X)$ (and therefore also $P(z)$) cannot be a discrete distribution. At the test time, when we want to generate new samples, we simply input values of $z \sim \mathcal{N}(0, I)$ into the decoder. This is schematically shown in Figure 7.

We will skip the definition of Conditional VAE here since the extension is quite straightforward, and refer the reader to the VAE tutorial paper [47] for details.

In [48], the authors show how to combine GAN and VAE in such a way, that the best of both worlds is used, and the limitations of both methods are mitigated. They propose to use a hybrid loss function which combines VAE and GAN:

$$\mathcal{L}(\theta, \eta) = \mathbb{E}_{q_{\eta}(z|x)} \left[-\lambda \|x - \mathcal{G}_{\theta}(z)\|_1 + \log \frac{\mathcal{D}_{\theta}(\mathcal{G}_{\theta}(z))}{1 - \mathcal{G}_{\theta}(z)} + \log \frac{\mathcal{C}_{\omega}(z)}{1 - \mathcal{C}_{\omega}(z)} \right]. \quad (15)$$

We need to build four networks here. The classifier $\mathcal{D}_{\theta}(x)$, which is trained to discriminate between reconstructions from an autoencoder, and real data points. A second classifier, which is trained to discriminate between latent samples produced by the encoder and samples from a standard Gaussian. The deep generative model $\mathcal{G}_{\theta}(z)$. And also the encoder network $q_{\eta}(z|x)$, which can

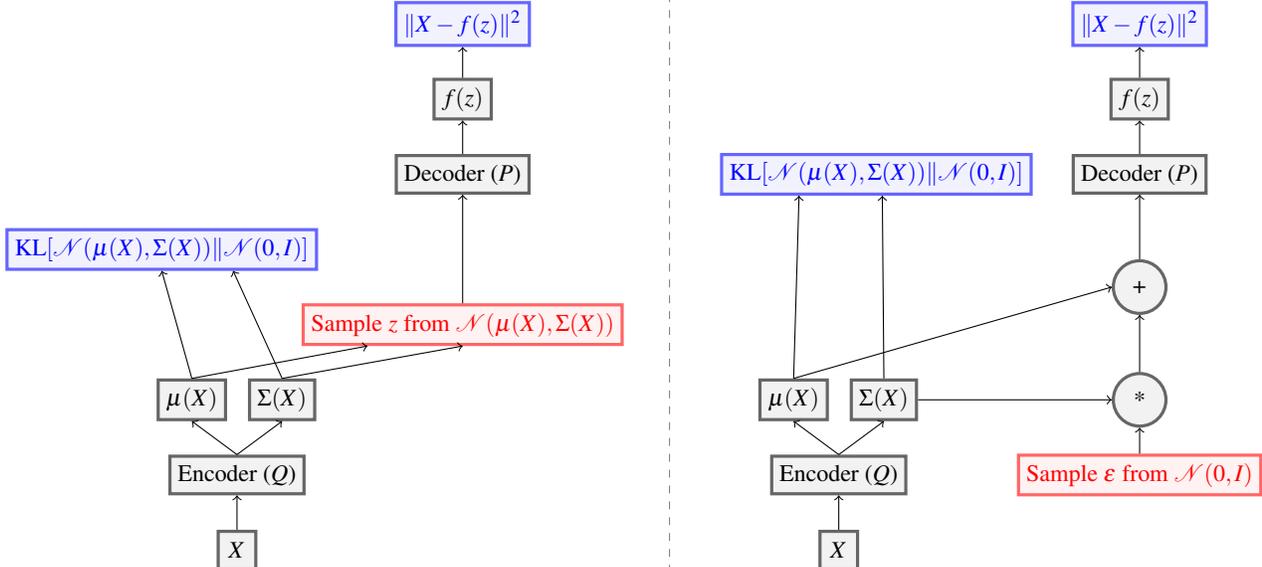


Figure 6: A training-time variational autoencoder implemented as a feed-forward neural network, where $P(X|z)$ is Gaussian. Red shows non-differentiable sampling operations. Blue shows loss layers. The left diagram is without using a reparameterization trick, while the right is with it. The feed-forward behavior of both networks is identical. However, the backpropagation algorithm can be applied only to the one on the right. Taken from [47].

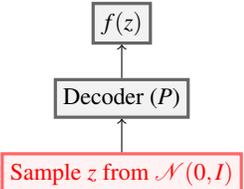


Figure 7: The testing-time variational autoencoder allowing us to generate new samples. The “encoder” pathway is simply discarded.

be implemented using arbitrary deep network. The authors refer to (15) as α -GAN. The algorithm alternates between updates of parameters of the generator θ , encoder η , synthetic likelihood discriminator ϕ , and the latent code discriminator ω . For the detailed description of the algorithm, we refer the reader to the original paper [48].

The problem of semi-supervised learning is covered in [49], which shows, that deep generative models and approximate Bayesian inference exploiting recent advances in variational methods can be used to provide significant improvements, making generative approaches highly competitive for semi-supervised learning. The authors describe a new framework for semi-supervised learning with generative models, employing rich parametric density estimators formed by the fusion of the probabilistic modeling and deep neural networks.

Interesting view of unsupervised learning of visual representations is shown in [50], where the problem is posed as solving a jigsaw puzzle. The authors introduce the context-free network (CFN), which takes image tiles as input and explicitly limits the receptive field. The experimental evaluation shows that the learned features capture semantically relevant content.

Authors of [51] argue that unsupervised pre-training is beneficial for deep learning in general. Their results suggest, that

unsupervised pre-training guides the learning towards basins of attraction of minima that support a better generalization from the training dataset.

Learned Loss Functions

Last but not least domain of GAN applications is its implicit ability of learning non-trivial loss functions. In certain setup, the GAN optimization criterion expresses the loss function indirectly, and therefore by optimizing the complex GAN criterion, we are also optimizing a loss function which we do not have to formulate explicitly. Unlike the previous domains, this is so far the least explored topic, so there are not many papers available in this area.

In [52], the authors propose discriminative adversarial networks (DAN) for semi-supervised learning and loss function learning³. Unlike the vanilla GAN, DAN uses two discriminators, instead of a generator and a discriminator. It can be seen as a framework to learn a loss function for predictors, that also implements semi-supervised learning.

DAN [52] are adversarial networks framework, that uses only discriminators. The authors propose a DAN formulation suitable for semi-supervised learning. However, we believe that other formulations suitable for the fully supervised learning are also possible. The original DAN use two discriminators: the predictor P , and the judge J . While P receives a data point x on input and outputs a prediction $p(x)$, J receives a data point x together with a label y , and produces a single scalar $J(x, y)$ representing the probability that x, y came from the labeled training data, rather than being predicted by P . See Figure 8 for an illustration of the DAN framework. Note the similarity between DAN and CGAN—while in CGAN the idea is to generate x conditioned on y , in DAN

³We should point out that in [52] the DAN is used on two different tasks of Natural Language Processing field. However, there do not seem to be any burdens prohibiting usage in the computer vision field.

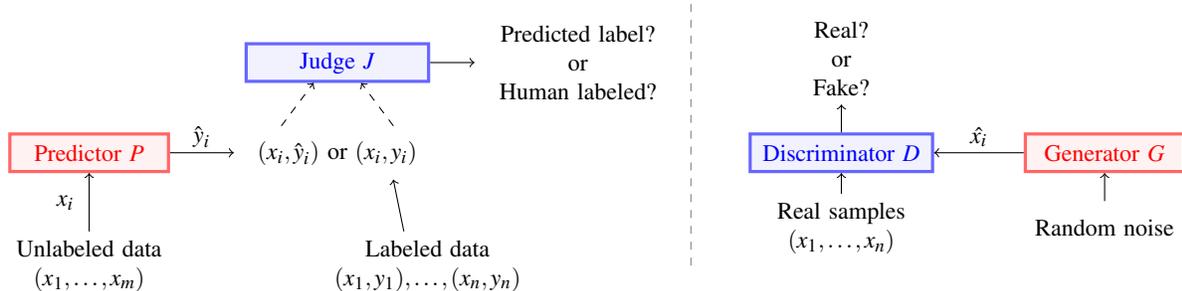


Figure 8: Left: DAN framework; Right: GAN framework.

the goal is to predict y conditioned on x . The cost function for DAN looks as follows:

$$\min_P \max_J V(J, P) = \mathbb{E}_{x, y \sim p_{\text{data}}(x, y)} [\log J(x, y)] + \mathbb{E}_{x \sim p_{\text{data}}(x)} [\log(1 - J(x, P(x)))] . \quad (16)$$

The important characteristic of this DAN formulation is that P does not make use of labels, so the semi-supervised learning is pretty straightforward within this framework. Even more importantly, there is no need for specification of a loss function for the predictor, it is learned implicitly by the judge.

Another mention of an implicit loss function learning is in [3], where the model learned to do image-to-image mapping without explicit specification of the loss function. A connection of GAN-based loss function learning for the generative model, and cost function learning in reinforcement learning (aka inverse reinforcement learning) is presented in [53]. In contrast, DAN concentrates on learning loss functions for discriminative models.

Adversarial training/testing

Adversarial examples were first introduced as attacks to weaken the performance of CNN by addition of noise [54]. This led to adversarial training [55] where adversarial examples were added to training to make the model more robust. Follow-up work on this eventually led to GAN model. The topic of adversarial training can also be interpreted as loss function learning, i.e. we can use adversarial loss for improving the final classifier robustness. A nice example of adversarial training is EL-GAN [56], where a GAN framework is used for loss embedding, by which the problem of ill-posed formulation of some tasks is mitigated. Since there are very stringent requirements on safety in AD, the adversarial examples generation might be used as a tool for testing corner cases and robustness. An automated testing mechanism for autonomous driving using deep learning were provided in [57], [58] but they do not leverage GAN.

IV OUR RESULTS

In this section, we would like to present some of our results from autonomous driving GAN application. To be more specific, we present the results on the soiling and adverse weather classification/enhancement.

The problem of soiling and adverse weather classification consist of the recognition of image deterioration and its possible enhancement. The image deterioration by soiling and adverse

weather is caused either by presence of some “soiling categories” (e.g. splashes of mud, rain drops, freeze, dust, etc.) on the camera lens, or by adverse weather conditions (e.g. heavy rain, snow, blizzards, etc.). The reasons for dealing with this problem are mainly two-fold: 1) by recognition of soiling type on the camera lens, we can leverage this information to initiate the cleaning system for the lens; 2) we can leverage the information to enhance the image quality (we call this “desoiling” in general, it consist of effects such as image “dehazing”, “deraining”, etc.). While the reasons for doing 1) are obvious, we can use 2) for example to improve the quality of the recognition pipeline.

We decided to use GAN for this problem generally because, as the reader can imagine, obtaining the relevant data is both very problematic and expensive (just imagine that someone has to annotate manually each rain drop on the camera lens during the heavy rain conditions). In Figure 9, we show how the mud splashed on camera can impact the image quality as well as an example of how the image taken during heavy rain looks like. Another reason is the semi-supervised learning potential, which can be easily achieved by GAN.

We started with a proof of concept experiment— we sorted our images to two categories, namely “clean” category, which consists of images that are not affected by soiling, and “soiled” category, consisting of images deteriorated by the presence of soiling. In Figure 10, we show the representatives of both classes. This allowed us to try using the CycleGAN [4]. We were happy to see, that just after a few epochs the generator started to correctly recognize which parts of the image are soiled. In the end of the training, we got a generator which is capable of “desoiling” the image as well as a generator which can introduce some soiling to the image, see Figure 11 for some examples. Note, that due to the relatively small dataset used for this experiment, the “desoiling” generator learned to introduce shadow of the car body to the image. This is because the vast majority of images in the “clean” category contained shadow of the car body. On the other hand, the “soiling” generator learned that the weather was usually cloudy on our images from the “soiled” category.

The CycleGAN experiment encouraged us in the presumption that GAN could be a nice solution for our problem. It also started several hypotheses, such as it might be eventually possible to get semantic segmentation of the soiling without having such explicit annotations, which would be very tedious and expensive to obtain. Another one is that by learning a high quality generators which introduce soiling in the image, we might be able to use them for the advanced data augmentation and create much bigger annotation using all data from various project, which are intentionally “clean”.



Figure 9: From left to right: a) soiled camera lens mounted to the car body; b) the image quality of the soiled camera from the previous image; c) an example of image soiled by a heavy rain.



Figure 10: The examples of the “clean” and “soiled” categories for our CycleGAN experiment.

In the latter direction, the experiment contains an important flaw— it is not able to produce soiling with variable output. Because of that, we tried another experiment with the MUNIT [29] approach. The main motivation for switching to MUNIT was its ability to split content from the style, which would help our intention to possess the control over generated images and therefore ease the further classifiers training. Our proof of concept for MUNIT failed due to the lack of high quality data. The resulting images from the generator contained a lot of artifacts and the control over the generated classes was not to our satisfaction. However, we still think that it is a promising branch to explore further. We depict our results from MUNIT experiment in Figure 12.

We believe that thanks to GAN, the problem of soiling and adverse weather classification/enhancement is possible to be solved even with a minimalist annotation efforts. We continue our research in the soiling and adverse weather classification/enhancement problematic with a lots of further research direction specified.

V DISCUSSION

In this section, we would like to discuss the main challenges of GAN. Namely, we discuss the problem of quantitative evaluation, adversarial attacks with respect to safety, or optimization stability.

Quantitative Evaluation

One of the critical challenges for GAN is their quantitative evaluation. The classic approach for generative models evaluation is based on the model likelihood. However, this approach is usu-

ally intractable [59]. The Inception Score (IS) [5], on the other hand, provides a way to quantitatively evaluate the quality and diversity of the generated samples, where conditional label distribution of samples containing meaningful objects should have low entropy and the variability of the samples should be high. Using the generic Inception Net [60] trained on ImageNet [61], it was found that IS is well correlated with scores from human annotators. However, IS is found to be insensitive to the prior distribution over labels. Fréchet Inception Distance (FID) [62] provides another approach for quantifying the quality of generated samples, where the samples are embedded into a feature space given by a specific layer of the Inception Net. Then these are modeled as a continuous multivariate Gaussian distribution. Finally, to quantify the quality of samples, the mean and covariance is estimated for the generated and the real data and the Fréchet distance is evaluated. FID score showed to be consistent with human judgment. Unlike IS, FID can detect intra-class mode dropping. IS mainly captures precision where FID captures both precision and recall.

Adversarial examples and Safety

Adversarial examples are inputs to machine learning models that have been intentionally modified in a way to fool the model. These modifications are, usually, not even noticed by a human observer, yet the classifier still makes wrong classifications. Moreover, adversarial examples can be used to perform attacks on machine learning systems even in a physical world. It was shown in [63] that machine learning systems are vulnerable to adversarial examples in physical world scenarios. Most of adversarial example attacks require knowledge of either the

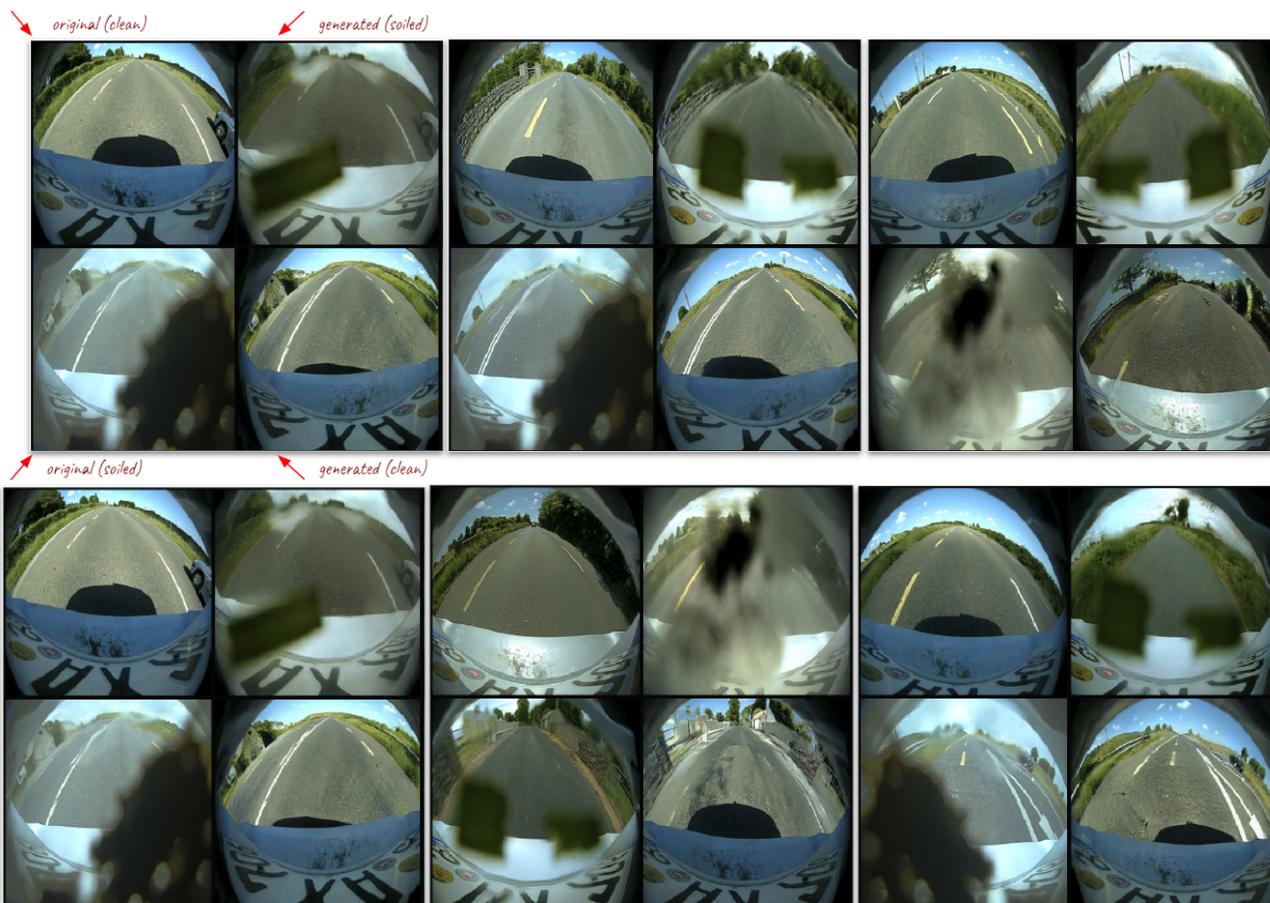


Figure 11: The results obtained from the CycleGAN proof of concept for the problem of soiling and adverse weather classification. Note the legend on the first image, which is self-explanatory.

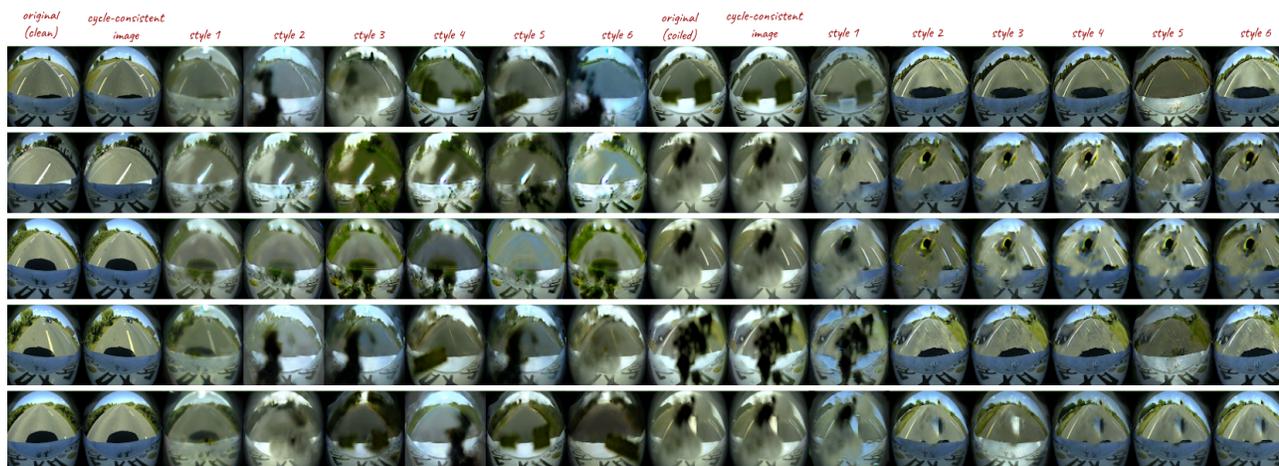


Figure 12: The results from the MUNIT experiment. Note that there are frequent artifacts present. However, it seems that sometimes the generated clean image did not contain the shadow of the body car, and the soiling generator is capable to introduce high variety of soiling categories.

model or its training data. However, [64] introduced a practical demonstration of an attacker with no such knowledge known as black box attacks. For the autonomous driving, attackers can target autonomous vehicles by using stickers to create an adversarial sign that the vehicle would interpret as another sign leading to performing unwanted or even dangerous behaviour. In the Universal Adversarial Training [65], instead of adding tailored perturbations to an image, an update can be added to any image in a broad class of images, while changing the predicted classes. One example for defensive approaches is the Defensive Distillation mechanism [66] that trains a model whose surface is smoothed in the directions an attacker will typically try to exploit, making it difficult to discover adversarial examples. Furthermore, Reinforcement Learning (RL) agents can also be manipulated by adversarial examples which causes a degraded or even a dangerous policy. Adversarial examples represent a concrete problem in AD safety and designing methods for preventing adversarial examples is an active area of research.

Optimization Stability

GAN training requires to find a Nash equilibrium [67] of non-convex non-cooperative game with high-dimensional parameters. The training typically consist of some form of gradient descent. Moreover, the optimization is, due to convergence issues, designed to minimize some loss function, rather than to find the Nash equilibrium. In [5], the authors use the heuristic understanding of the non-convergence problem to introduce several techniques for its improving, such as feature matching trying to prevent overtraining of the discriminator, minibatch discrimination, which, by allowing the discriminator to see several data samples in combination, prevents generator from collapsing to always generate the same sample, and many other “hacks” trying to overcome the identified problems.

The authors of [68] tried to approach the known convergence and stability problems of GAN training differently. They described the source cause of these problems by identifying the Kullback-Leibler (KL) divergence minimization task hidden in the probability distribution learning. The problem arises when dealing with distributions supported by low-dimensional manifolds, because for these the KL is either not defined or simply infinite. As a solution to overcome this problem, the authors propose to use a different distance function, which consists of the earth-mover, or Wasserstein, distance. In experimental evaluation, they show that the desirable stabilization effect is achieved. The algorithm proposed in [68] have one flaw, which is the clipping of weights to enforce 1-Lipshitzness required in the optimization theory of WGAN. This flaw was corrected by [7], where the gradient penalty was introduced.

VI CONCLUSIONS

We compiled a detailed overview of GAN models and provide a taxonomic survey of various applications of GAN in autonomous driving. GAN have a potential for high impact for autonomous driving applications but there is slow progress in this area. There are plenty of other applications beyond the standard image translation application. We also discussed the main challenges and

open problems which have to be resolved in order for it to be more practically used. We hope that this paper encourages further research in applying GAN for autonomous driving applications.

ACKNOWLEDGMENTS

The authors would like to thank their employer for the opportunity to perform fundamental research. We would also like to thank the company Borealis AI for the GAN T-shirts which inspired the title of this paper⁴. Last but not least, we would like to express our very special thanks to our colleague GANesh Sistu, who gave us permission to use his nice figures.

REFERENCES

- [1] M. Siam, S. Elkerdawy, M. Jagersand, and S. Yogamani, “Deep Semantic Segmentation for Automated Driving: Taxonomy, Roadmap and Challenges,” in *Intelligent Transportation Systems (ITSC), 2017 IEEE 20th International Conference on*, pp. 1–8, IEEE, 2017.
- [2] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative Adversarial Nets,” in *Advances in Neural Information Processing Systems 27* (Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, eds.), pp. 2672–2680, Curran Associates, Inc., 2014.
- [3] P. Isola, J. Zhu, T. Zhou, and A. A. Efros, “Image-to-Image Translation with Conditional Adversarial Networks,” in *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pp. 5967–5976, 2017.
- [4] J. Zhu, T. Park, P. Isola, and A. A. Efros, “Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks,” in *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*, pp. 2242–2251, 2017.
- [5] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, “Improved Techniques for Training GANs,” in *Advances in Neural Information Processing Systems*, pp. 2234–2242, 2016.
- [6] M. Arjovsky, S. Chintala, and L. Bottou, “Wasserstein GAN,” *CoRR*, vol. abs/1701.07875, 2017.
- [7] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville, “Improved Training of Wasserstein GANs,” in *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pp. 5769–5779, 2017.
- [8] I. J. Goodfellow, “NIPS 2016 tutorial: Generative adversarial networks,” *CoRR*, vol. abs/1701.00160, 2017.
- [9] M. Mirza and S. Osindero, “Conditional Generative Adversarial Nets,” *CoRR*, vol. abs/1411.1784, 2014.
- [10] Y. LeCun and C. Cortes, “MNIST handwritten digit database,” 2010.
- [11] C. Li, Z. Wang, and H. Qi, “Fast-Converging Conditional Generative Adversarial Networks for Image Synthesis,” *CoRR*, vol. abs/1805.01972, 2018.
- [12] A. Krizhevsky, V. Nair, and G. Hinton, “CIFAR-10 (Canadian Institute for Advanced Research),”
- [13] D. P. Kingma and M. Welling, “Auto-Encoding Variational Bayes,” in *ICLR 2014*, vol. abs/1312.6114, 2013.
- [14] D. R. Hjelm, A. P. Jacob, T. Che, K. Cho, and Y. Ben-

⁴<https://www.borealisai.com/en/blog/yes-we-gan/>

- gio, "Boundary-Seeking Generative Adversarial Networks," *CoRR*, vol. abs/1702.08431v2, 2017.
- [15] A. Brock, J. Donahue, and K. Simonyan, "Large Scale GAN Training for High Fidelity Natural Image Synthesis," *arXiv preprint arXiv:1809.11096*, 2018.
- [16] T. Karras, T. Aila, S. Laine, and J. Lehtinen, "Progressive Growing of GANs for Improved Quality, Stability, and Variation," *arXiv preprint arXiv:1710.10196*, 2017.
- [17] T. Karras, S. Laine, and T. Aila, "A Style-Based Generator Architecture for Generative Adversarial Networks," *arXiv preprint arXiv:1812.04948*, 2018.
- [18] H. Zhang, I. Goodfellow, D. Metaxas, and A. Odena, "Self-Attention Generative Adversarial Networks," *arXiv preprint arXiv:1805.08318*, 2018.
- [19] M.-Y. Liu, T. Breuel, and J. Kautz, "Unsupervised Image-to-Image Translation Networks," in *Advances in Neural Information Processing Systems 30* (I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, eds.), pp. 700–708, Curran Associates, Inc., 2017.
- [20] D. Lee, S. Liu, J. Gu, M. Liu, M. Yang, and J. Kautz, "Context-Aware Synthesis and Placement of Object Instances," *CoRR*, vol. abs/1812.02350, 2018.
- [21] D. Pathak, P. Krähenbühl, J. Donahue, T. Darrell, and A. A. Efros, "Context Encoders: Feature Learning by Inpainting," in *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pp. 2536–2544, 2016.
- [22] A. Odena, C. Olah, and J. Shlens, "Conditional Image Synthesis with Auxiliary Classifier GANs," in *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, pp. 2642–2651, 2017.
- [23] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, *et al.*, "Photo-realistic Single Image Super-Resolution Using a Generative Adversarial Network," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 105–114, IEEE, 2017.
- [24] T. Kim, M. Cha, H. Kim, J. K. Lee, and J. Kim, "Learning to Discover Cross-Domain Relations with Generative Adversarial Networks," *arXiv preprint arXiv:1703.05192*, 2017.
- [25] Y. Choi, M. Choi, M. Kim, J.-W. Ha, S. Kim, and J. Choo, "StarGAN: Unified Generative Adversarial Networks for Multi-Domain Image-to-Image Translation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8789–8797, 2018.
- [26] M.-Y. Liu, T. Breuel, and J. Kautz, "Unsupervised Image-to-Image Translation Networks," in *Advances in Neural Information Processing Systems*, pp. 700–708, 2017.
- [27] T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, A. Tao, J. Kautz, and B. Catanzaro, "High-resolution Image Synthesis and Semantic Manipulation With Conditional GANs," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8798–8807, 2018.
- [28] J.-Y. Zhu, R. Zhang, D. Pathak, T. Darrell, A. A. Efros, O. Wang, and E. Shechtman, "Toward Multimodal Image-to-Image Translation," in *Advances in Neural Information Processing Systems*, pp. 465–476, 2017.
- [29] X. Huang, M.-Y. Liu, S. Belongie, and J. Kautz, "Multimodal Unsupervised Image-to-Image Translation," *arXiv preprint arXiv:1804.04732*, 2018.
- [30] A. Almahairi, S. Rajeswar, A. Sordoni, P. Bachman, and A. Courville, "Augmented CycleGAN: Learning Many-to-Many Mappings from Unpaired Data," *arXiv preprint arXiv:1802.10151*, 2018.
- [31] J. Wu, C. Zhang, T. Xue, B. Freeman, and J. Tenenbaum, "Learning a Probabilistic Latent Space of Object Shapes via 3D Generative-Adversarial Modeling," in *Advances in Neural Information Processing Systems*, pp. 82–90, 2016.
- [32] M. Gadelha, S. Maji, and R. Wang, "3D Shape Induction from 2D Views of Multiple Objects," in *3D Vision (3DV), 2017 International Conference on*, pp. 402–411, IEEE, 2017.
- [33] C.-L. Li, M. Zaheer, Y. Zhang, B. Poczoz, and R. Salakhutdinov, "Point Cloud GAN," *arXiv preprint arXiv:1810.05795*, 2018.
- [34] T.-C. Wang, M.-Y. Liu, J.-Y. Zhu, G. Liu, A. Tao, J. Kautz, and B. Catanzaro, "Video-to-Video Synthesis," *arXiv preprint arXiv:1808.06601*, 2018.
- [35] M. Saito, E. Matsumoto, and S. Saito, "Temporal Generative Adversarial Nets with Singular Value Clipping," in *IEEE International Conference on Computer Vision (ICCV)*, vol. 2, p. 5, 2017.
- [36] C. Vondrick, H. Pirsaviash, and A. Torralba, "Generating Videos with Scene Dynamics," in *Advances In Neural Information Processing Systems*, pp. 613–621, 2016.
- [37] K. Bousmalis, N. Silberman, D. Dohan, D. Erhan, and D. Krishnan, "Unsupervised Pixel-Level Domain Adaptation with Generative Adversarial Networks," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 1, p. 7, 2017.
- [38] K. Bousmalis, A. Irpan, P. Wohlhart, Y. Bai, M. Kelcey, M. Kalakrishnan, L. Downs, J. Ibarz, P. Pastor, K. Konolige, *et al.*, "Using Simulation and Domain Adaptation to Improve Efficiency of Deep Robotic Grasping," *arXiv preprint arXiv:1709.07857*, 2017.
- [39] X. Pan, Y. You, Z. Wang, and C. Lu, "Virtual to Real Reinforcement Learning for Autonomous Driving," *arXiv preprint arXiv:1704.03952*, 2017.
- [40] A. Bewley, J. Rigley, Y. Liu, J. Hawke, R. Shen, V.-D. Lam, and A. Kendall, "Learning to Drive from Simulation without Real World Labels," *arXiv preprint arXiv:1812.03823*, 2018.
- [41] K. Ehsani, R. Mottaghi, and A. Farhadi, "SEGAN: Segmenting and Generating the Invisible," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6144–6153, 2018.
- [42] X. Wang, A. Shrivastava, and A. Gupta, "A-Fast-RCNN: Hard Positive Generation via Adversary for Object detection," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [43] J. Li, X. Liang, Y. Wei, T. Xu, J. Feng, and S. Yan, "Perceptual Generative Adversarial Networks for Small Object detection," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1951–1959, IEEE, 2017.
- [44] S. Iizuka, E. Simo-Serra, and H. Ishikawa, "Globally and Locally Consistent Image Completion," *ACM Transactions on Graphics (TOG)*, vol. 36, no. 4, p. 107, 2017.
- [45] R. A. Yeh, C. Chen, T. Yian Lim, A. G. Schwing, M. Hasegawa-Johnson, and M. N. Do, "Semantic Image Inpainting with Deep Generative Models," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5485–5493, 2017.
- [46] D. Pathak, P. Krahenbuhl, J. Donahue, T. Darrell, and A. A. Efros, "Context Encoders: Feature Learning by Inpainting," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2536–2544, 2016.
- [47] C. Doersch, "Tutorial on Variational Autoencoders," *CoRR*, vol. abs/1606.05908, 2016.
- [48] M. Rosca, B. Lakshminarayanan, D. Warde-Farley, and S. Mohamed, "Variational Approaches for Auto-Encoding Generative Ad-

- versarial Networks,” *CoRR*, vol. abs/1706.04987, 2017.
- [49] D. P. Kingma, S. Mohamed, D. J. Rezende, and M. Welling, “Semi-supervised Learning with Deep Generative Models,” in *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pp. 3581–3589, 2014.
- [50] M. Noroozi and P. Favaro, “Unsupervised Learning of Visual Representations by Solving Jigsaw Puzzles,” in *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part VI*, pp. 69–84, 2016.
- [51] D. Erhan, Y. Bengio, A. C. Courville, P. Manzagol, P. Vincent, and S. Bengio, “Why Does Unsupervised Pre-training Help Deep Learning?,” *Journal of Machine Learning Research*, vol. 11, pp. 625–660, 2010.
- [52] C. N. dos Santos, K. Wadhawan, and B. Zhou, “Learning Loss Functions for Semi-supervised Learning via Discriminative Adversarial Networks,” *CoRR*, vol. abs/1707.02198, 2017.
- [53] C. Finn, P. F. Christiano, P. Abbeel, and S. Levine, “A Connection between Generative Adversarial Networks, Inverse Reinforcement Learning, and Energy-Based Models,” in *NIPS Workshop on Adversarial Training (WAT 2016)*, vol. abs/1611.03852, 2016.
- [54] A. Nguyen, J. Yosinski, and J. Clune, “Deep neural Networks are Easily Fooled: High Confidence Predictions for Unrecognizable Images,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 427–436, 2015.
- [55] I. Goodfellow, J. Shlens, and C. Szegedy, “Explaining and Harnessing Adversarial Examples,” in *International Conference on Learning Representations*, 2015.
- [56] M. Ghafoorian, C. Nugteren, N. Baka, O. Booi, and M. Hofmann, “EL-GAN: Embedding Loss Driven Generative Adversarial Networks for Lane Detection,” in *Computer Vision - ECCV 2018 Workshops - Munich, Germany, September 8-14, 2018, Proceedings, Part I*, pp. 256–272, 2018.
- [57] Y. Tian, K. Pei, S. Jana, and B. Ray, “DeepTest: Automated Testing of Deep-Neural-Network-Driven Autonomous Cars,” in *Proceedings of the 40th international conference on software engineering*, pp. 303–314, ACM, 2018.
- [58] K. Pei, Y. Cao, J. Yang, and S. Jana, “Deepxplore: Automated Whitebox Testing of Deep Learning Systems,” in *Proceedings of the 26th Symposium on Operating Systems Principles*, pp. 1–18, ACM, 2017.
- [59] F. Huszár, “How (not) to Train Your Generative Model: Scheduled Sampling, Likelihood, Adversary?,” *arXiv preprint arXiv:1511.05101*, 2015.
- [60] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. A. Alemi, “Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning,” in *AAAI*, vol. 4, p. 12, 2017.
- [61] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “ImageNet: A Large-Scale Hierarchical Image Database,” in *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pp. 248–255, Ieee, 2009.
- [62] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, “GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium,” in *Advances in Neural Information Processing Systems*, pp. 6626–6637, 2017.
- [63] A. Kurakin, I. Goodfellow, and S. Bengio, “Adversarial examples in the physical world,” *arXiv preprint arXiv:1607.02533*, 2016.
- [64] N. Papernot, P. McDaniel, I. Goodfellow, S. Jha, Z. B. Celik, and A. Swami, “Practical black-box attacks against machine learning,” in *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security*, pp. 506–519, ACM, 2017.
- [65] A. Shafahi, M. Najibi, Z. Xu, J. Dickerson, L. S. Davis, and T. Goldstein, “Universal Adversarial Training,” *arXiv preprint arXiv:1811.11304*, 2018.
- [66] N. Papernot, P. McDaniel, X. Wu, S. Jha, and A. Swami, “Distillation as a defense to adversarial perturbations against deep neural networks,” in *2016 IEEE Symposium on Security and Privacy (SP)*, pp. 582–597, IEEE, 2016.
- [67] J. Nash, “Non-Cooperative Games,” *Annals of Mathematics*, vol. 54, no. 2, pp. 286–295, 1951.
- [68] M. Arjovsky, S. Chintala, and L. Bottou, “Wasserstein Generative Adversarial Networks,” in *Proceedings of the 34th International Conference on Machine Learning (D. Precup and Y. W. Teh, eds.)*, vol. 70 of *Proceedings of Machine Learning Research*, (International Convention Centre, Sydney, Australia), pp. 214–223, PMLR, 06–11 Aug 2017.

AUTHORS BIOGRAPHY

Michal Uříčář received MSc. degree (major in computer graphics and minor in computer vision) from Czech Technical University in Prague in 2011, and Ph.D. from the same university (in 2018) while working in the machine learning group of the Center for Machine Perception under the supervision of Dr. Vojtěch Franc. Since 2016, he is working as a researcher in Valeo R&D, Prague. His main research focus is machine learning, with applications in computer vision. He received the best paper prize at VISAPP in 2012 and the 3rd place at Looking at People Challenge: Age estimation track, organized at CVPR 2016.

Pavel Křížek received M.Sc. in Control Engineering from the Czech Technical University in Prague in 2003, and Ph.D. in Bio-Cybernetics and Artificial Intelligence from the same university in 2008 at the Center for Machine Perception. During Ph.D., he spent two years at the Center for Vision, Speech and Signal Processing at University of Surrey, UK. Over five years he was working as a Postdoc on several projects focusing on super-resolution imaging in fluorescence microscopy at the Charles University in Prague. Since 2014, he is a researcher in Valeo R&D, Prague. He has over 16 years of experience in image processing, computer vision, and machine learning applications including 4 years of experience in industrial automotive systems. He is an author of 30 peer reviewed journal and conference papers and 4 patents.

David Hurych received his Ph.D. in Bio-Cybernetics and Artificial Intelligence from Czech Technical University in Prague in 2014. He received his M.Sc. from Computer Science from VSB - Technical University of Ostrava in 2007. Since 2014 he is a researcher and team leader in Valeo R&D, Prague and was awarded machine learning expert price. His main research focus is machine learning with applications in computer vision. He is an author of 10 conference and journal papers (including one TPAMI) and 4 awarded patents. He received the best student paper award at VISAPP 2011. He has 11 years of machine learning and computer vision experience including 5 years of experience in automotive industry and autonomous driving.

Ibrahim Sobh received Ph.D in Deep Reinforcement Learning for fast learning agents acting in 3D environments directly from high-dimensional sensory inputs, from Cairo University Faculty of Engineering, in 2017. He received his B.Sc., in 1997, and M.Sc., in 2009, from the same University. His M.Sc. thesis is in the field of Machine Learning and applied on automatic documents summarization. Ibrahim has participated in several related national and international mega projects, conferences and summits. Ibrahim's publications including international journals and conference papers are mainly in the machine and deep learning fields for Natural language processing, Speech processing, Computer vision and autonomous driving.

Senthil Yogamani is a computer vision architect and technical leader at Valeo Vision systems. He is currently focused on research and design of the overall computer vision algorithm architecture for surround-view camera visual perception in autonomous driving systems. He has over 13 years of experience in computer vision and machine learning including 10 years of experience in industrial automotive systems. He is an author of 50 peer reviewed publications and 33 patents. He serves in the editorial board of various leading IEEE automotive conferences including ITSC and ICVES and advisory board of various industry consortia including Khronos, Cognitive Vehicles and IS Auto. He is a recipient of best associate editor award at ITSC 2015 and best paper award at ITST 2012.

Patrick Denny is a Senior Research Engineer and Senior Expert at Valeo Vision Systems, where he has been a Technical Lead and internal consultant on RF, camera and imaging system developments for the last 17 years for the leading automotive marques. He is Adjunct Professor of Automotive Electronics in NUI Galway. He received his B.Sc. in Experimental Physics and Mathematics from NUI Maynooth in 1993, an M.Sc. in Pure Mathematics from NUI Galway in 1994 and a Ph.D in Physics in 2000 from NUI Galway and a Professional Diploma in Data Analytics from University College Dublin in 2016. He has 119 patent applications in 60 patent families and over 30 publications in refereed journals, conference proceedings and books. He co-founded and chairs and is on the steering committee of the AutoSens and Electronic Imaging's Autonomous Vehicles and Machines conferences and co-founded and chaired IS Auto. He is also a co-founder and board member of the IEEE P2020 Automotive Imaging Standards body.

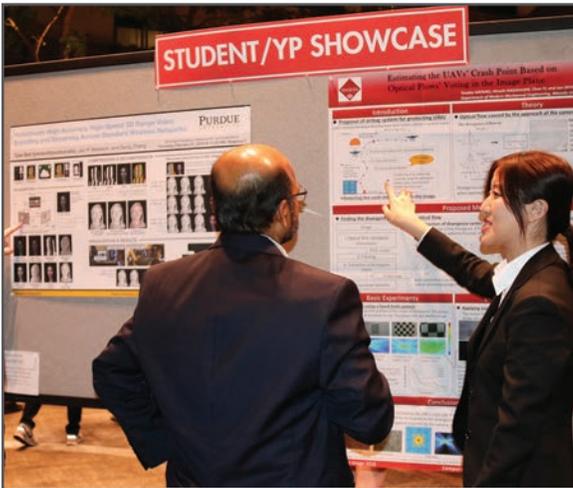
JOIN US AT THE NEXT EI!

IS&T International Symposium on

Electronic Imaging

SCIENCE AND TECHNOLOGY

Imaging across applications . . . Where industry and academia meet!



- **SHORT COURSES • EXHIBITS • DEMONSTRATION SESSION • PLENARY TALKS •**
- **INTERACTIVE PAPER SESSION • SPECIAL EVENTS • TECHNICAL SESSIONS •**

www.electronicimaging.org

