RUHR-UNIVERSITÄT BOCHUM

INI INSTITUT FÜR NEUROINFORMATIK

RUB

Generation of Natural Traffic Sign Images Using Domain Translation with Cycle-Consistent Generative Adversarial Networks

Dominic Spata, Daniela Horn and Sebastian Houben
Institute for Neural Computation, Ruhr University Bochum, Germany

Top: Generated traffic sign images. *Bottom:* Each image's nearest neighbor in terms of Euclidean distance from the training set of real traffic sign photographs.[1]

Motivation

- Machine learning methods for traffic sign recognition require huge amounts of image data
- Challenges for data acquisition:
  - high number of possible classes
  - unbalanced distribution of samples
  - variance in diverting background
  - numerous recording artefacts

Idea:

- Automatic generation of life-like traffic sign images
- Use of artificial samples for image classifier training

Cycle-Consistent Generative Adversarial Networks

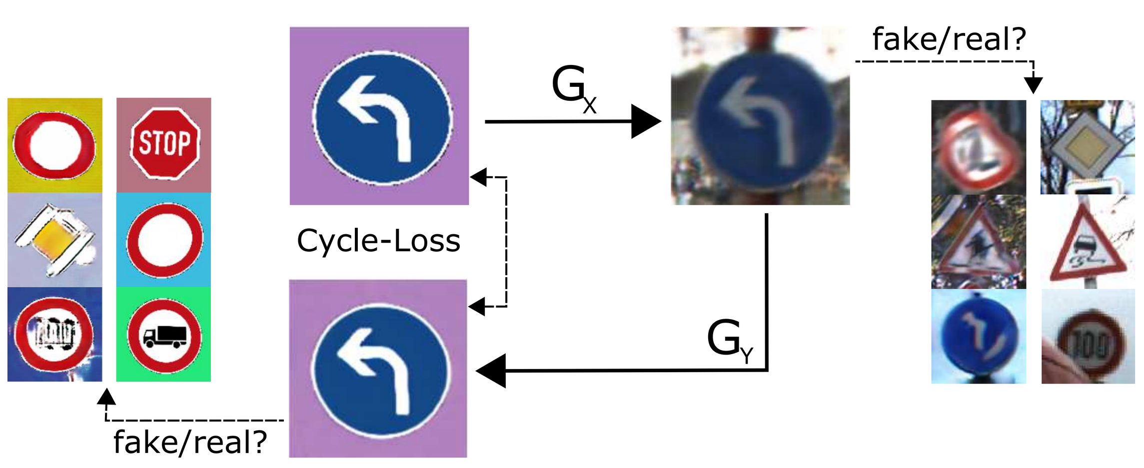

- Image-to-image translation between two domains X and Y
- Cycle-consistency: mappings are the inverse of one another
- Can be trained using unpaired datasets
- Mappings are actual translators rather than being arbitrary
- Generated background style is influenced by simple prototype background color

Learned association from background color in prototype to background and illumination style in generated image

Experimental Setup

- Comparison of multi-class SVMs trained on HOG features
- Baseline classifier (SVM_{Base}) trained solely on real-world data
- Per experiment two further classifiers with varying training inputs

Training on Fully Generated Data

- SVM_{Gen} trained entirely on generated images
- SVM_{Proto} trained entirely on prototype images
- Results grouped into five traffic sign categories

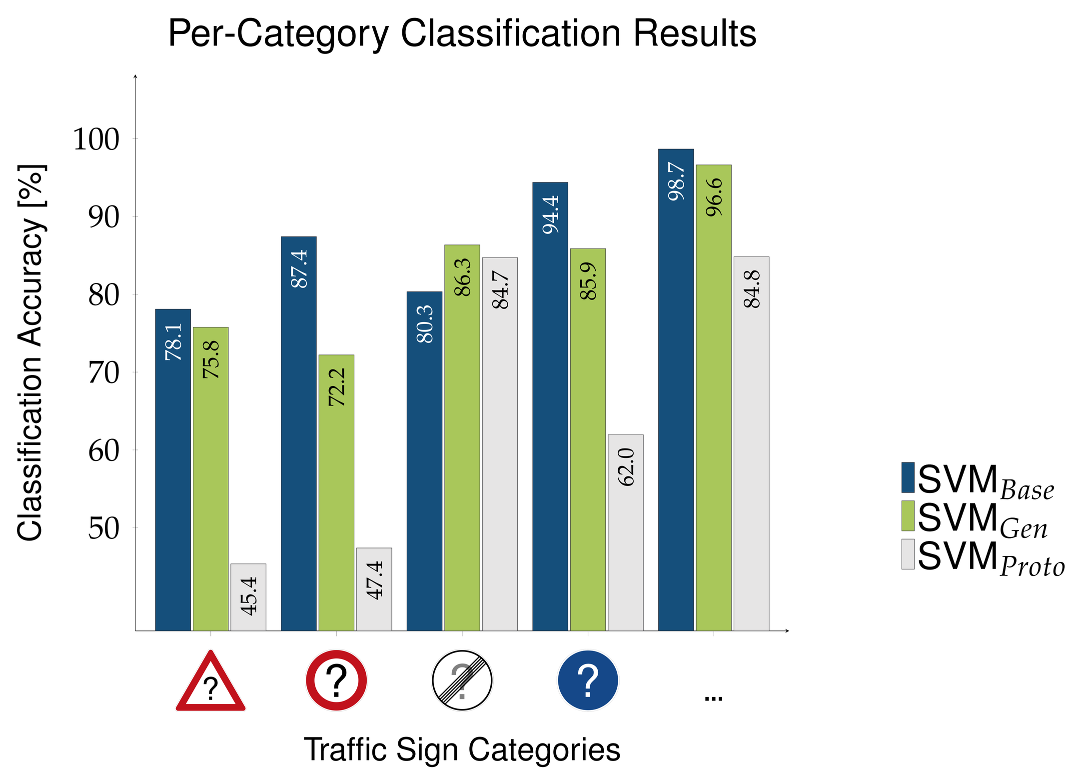

Training on Partially Generated Data

- Replace one sign class from original dataset with synthetic data
- Compare against baseline classifier
- $SVM_{GenClass}$ swaps real-world samples for generated images
- $SVM_{ProtoClass}$ swaps real-world samples for prototype images

Contributions

- Generation method for traffic sign images
- High control over pose and background
- Facilitates data acquisition
- Real-world image data remains essential

$\Longrightarrow$ Use for data augmentation and rare classes

Get the source!

Class	SVM_{Base}	$SVM_{GenClass}$	$SVM_{ProtoClass}$
	Replacing Class "No Entry (Trucks)" — Classification Accuracy [%]		
No entry (trucks)	97.18	88.73 (−8.45)	0.00
Speed limit 100	74.44	73.99 (−0.45)	74.44
Roundabout	70.46	72.73 (+2.27)	70.46
Total	87.97	87.87 (−0.10)	86.86

Class	SVM_{Base}	$SVM_{GenClass}$	$SVM_{ProtoClass}$
	Replacing Class "Slippery Road" — Classification Accuracy [%]		
Slippery Road	67.11	60.53 (−6.58)	0.00
(Non-substituted classes show no performance change for $SVM_{GenClass}$)			
Total	87.97	87.89 (−0.08)	87.17

Class	SVM_{Base}	$SVM_{GenClass}$	$SVM_{ProtoClass}$
	Replacing Class "Pass Right" — Classification Accuracy [%]		
Pass right	95.87	78.47 (−17.40)	14.16
Stop	92.36	93.06 (+0.70)	93.06
Forward or right	96.92	98.46 (+1.54)	98.46
Total	87.97	87.06 (−0.91)	83.61

Results for training on partially generated data. SVMs were trained on real-world data, with the exception of one class, which was replaced by either generated images ($SVM_{GenClass}$) or prototype images ($SVM_{ProtoClass}$). Numbers in parenthesis display difference to the baseline classifier.

[1] Stallkamp, J., Schlipsing, M., Salmen, J., & Igel, C. (2011). The German Traffic Sign Recognition Benchmark: A multi-class classification competition. In *IEEE International Joint Conference on Neural Networks* (pp. 1453–1460).