King Abdullah University of Science and Technology

KAUST

Studying Noise Sensitivity of Deep Neural Networks

Prof. Bernard Ghanem, Assoc. Prof. of EE & CS

VISUAL COMPUTING CENTER

Main Research Themes @ IVUL

Activity Understanding

- Activity Detection
- Efficient Search
- Object Tracking

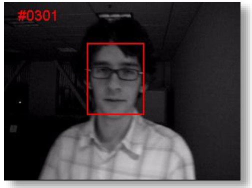


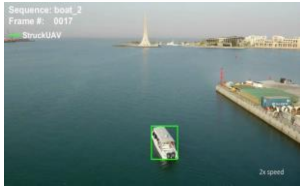

ACTIVITYNET

Vision for Automated Navigation

- Sim4CV
- Transfer Learning
- Applications

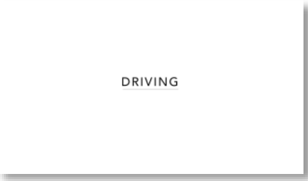


Fundamentals

- Optimization for CV&ML (sparse, low-rank, integer)
- Deep DNN Understanding

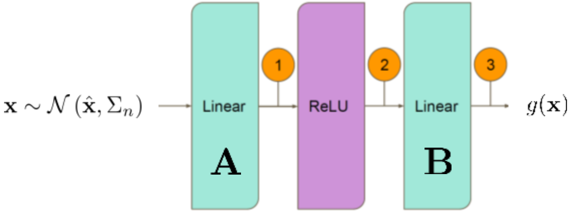

$$\min_{\mathbf{x}} \ f(\mathbf{x}) \quad \text{s.t. } \mathbf{x} \in \{1, -1\}^n; \mathbf{x} \in \Omega$$

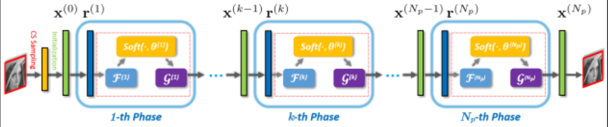

THEME: ACTIVITY UNDERSTANDING

جامعة الملك عبدالله للعلوم والتقنية
King Abdullah University of Science and Technology
VCC VISUAL COMPUTING CENTER

Bernard Ghanem

Fun Facts

By 2017, online video will account for 74% of all online traffic[3]

55% of people watch videos online every day[1]

45% of people watch more than an hour of Facebook or YouTube videos a week[2]

Almost 50% of internet users look for videos related to a product or service before visiting a store[4]

85% of Facebook video is watched without sound[5]

Source: 1) MWP Statistics, 2015; 2) HubSpot, 2016 3) KPCB, 2016 4) Google, 2016; 5) DIGIDAY, 2016

Applications of Activity Understanding

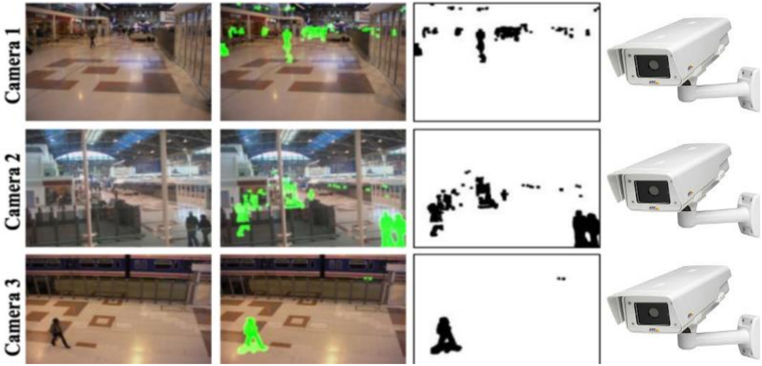

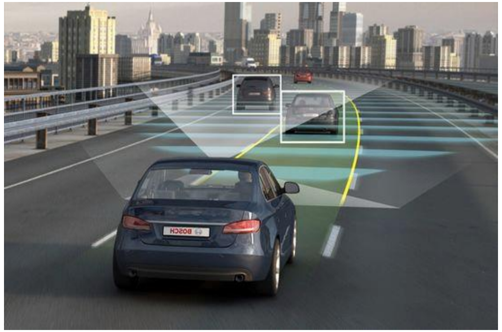

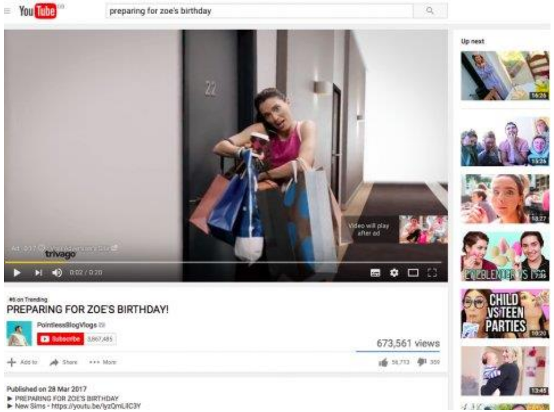

Activity Detection @ IVUL

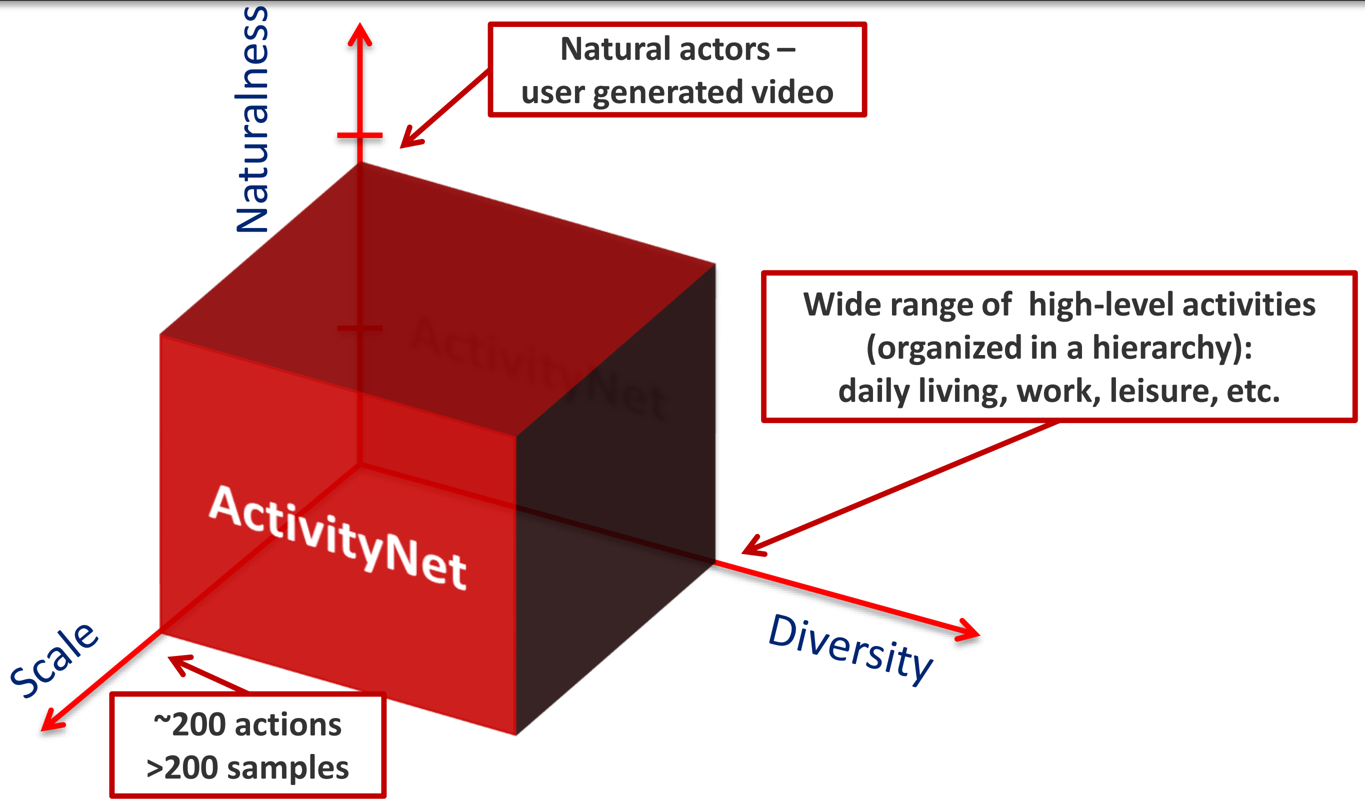

Activity Detection @ IVUL

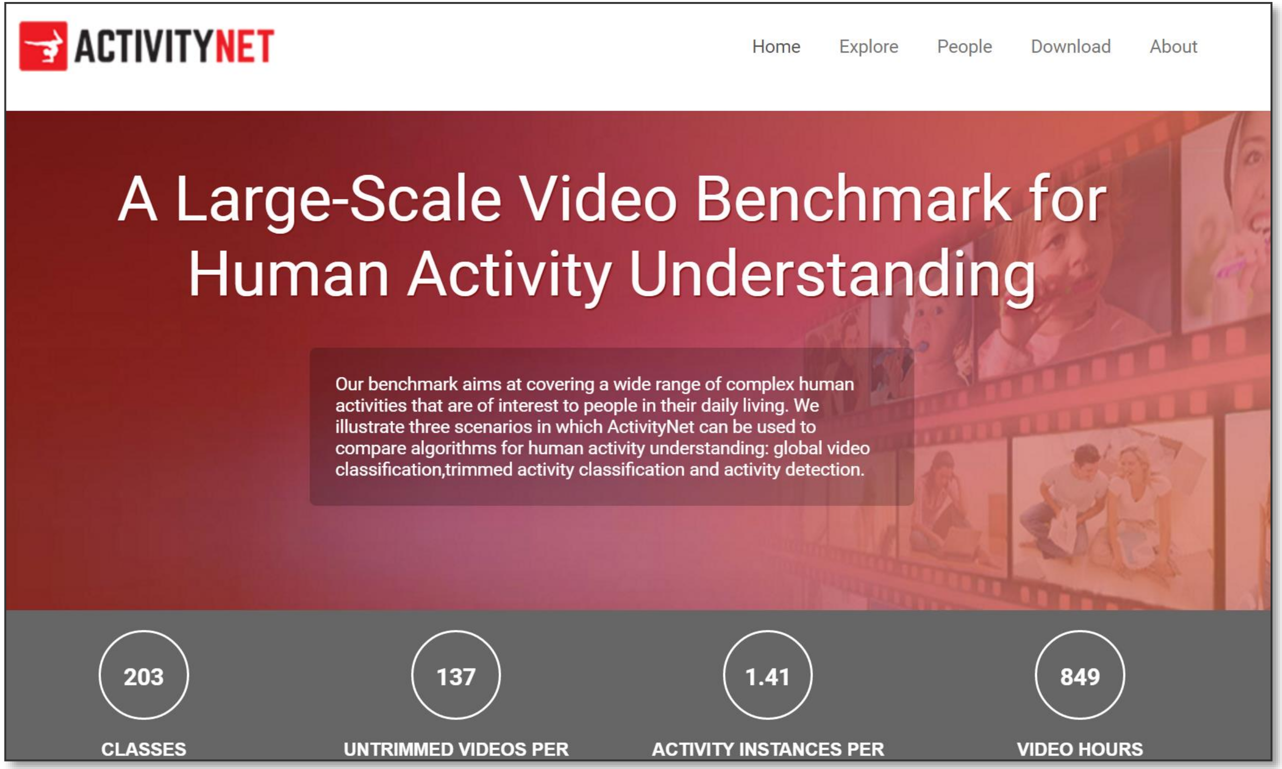

Google

Google Faculty Research Award in 2015; 1st in MENA for Machine Perception; 1st in Saudi Arabia

1st Version (R1.1):
- ~200 classes
- ~850 hours
- class hierarchy

ActivityNet: A Large-Scale Video Benchmark for Human Activity Understanding [CVPR'15]

Activity Detection @ IVUL

ACTIVITYNET
Large Scale Activity Recognition Challenge

HOME PEOPLE CHALLENGE PROGRAM DATES EVALUATION CONTACT

Challenge Description

Challenge Introduction

We are proud to announce that this year the challenge will host six diverse tasks which aim to push the limits of semantic visual understanding of videos as well as bridging visual content with human captions. Three out of the seven tasks are based on the ActivityNet dataset, which was introduced in CVPR 2015 and organized hierarchically in a semantic taxonomy. These tasks focus on trace evidence of activities in time in the form of proposals, class labels, and captions.

In this installment of the challenge, we will host three guest tasks which enrich the understanding of visual information in videos. These tasks focus on complementary aspects of the activity recognition problem at large scale and involve challenging and recently compiled activity/action datasets, including Kinetics (Google DeepMind), AVA (Berkeley and Google), and Moments in Time (MIT and IBM Research).

ActivityNet Tasks

TASK 1
Temporal Action Proposals (ActivityNet)
This task is intended to evaluate the ability of algorithms to generate high quality action proposals. The goal is to produce a set of candidate temporal segments that are likely to contain a human action.
DETAILS

TASK 2
Temporal Action Localization (ActivityNet)
This task is intended to evaluate the ability of algorithms to temporally localize activities in untrimmed video sequences. Here, videos can contain more than one activity instance, and mutiple activity categories can appear in the video.
DETAILS

TASK 3
Dense-Captioning Events in Videos (ActivityNet Captions)
This task involves both detecting and describing events in a video. For this task, participants will use the ActivityNet Captions dataset, a new large-scale benchmark for dense-captioning events.
DETAILS

HOME PEOPLE CHALLENGE PROGRAM DATES EVALUATION CONTACT

Challenge Description

Challenge Introduction

We are proud to announce that this year the challenge will host six diverse tasks which aim to push the limits of semantic visual understanding of videos as well as bridging visual content with human captions. Three out of the seven tasks are based on the ActivityNet dataset, which was introduced in CVPR 2015 and organized hierarchically in a semantic taxonomy. These tasks focus on trace evidence of activities in time in the form of proposals, class labels, and captions.

In this installment of the challenge, we will host three guest tasks which enrich the understanding of visual information in videos. These tasks focus on complementary aspects of the activity recognition problem at large scale and involve challenging and recently compiled activity/action datasets, including Kinetics (Google DeepMind), AVA (Berkeley and Google), and Moments in Time (MIT and IBM Research).

Guest Tasks

TASK A
Trimmed Activity Recognition (Kinetics)
This task is intended to evaluate the ability of algorithms to recognize activities in trimmed video sequences. Here, videos contain a single activity, and all the clips have a standard duration of ten seconds. For this task, participants will use the Kinetics dataset, a large-scale benchmark for trimmed action classification.
DETAILS

TASK B
Spatio-temporal Action Localization (AVA)
This task is intended to evaluate the ability of algorithms to localize human actions in space and time. Each labeled video segment can contain multiple subjects, each performing potentially multiple actions. The goal is to identify these subjects and actions over continuous 15-minute video clips extracted from movies. For this task, participants will use the new AVA atomic visual actions dataset.
DETAILS

TASK C
Trimmed Event Recognition (Moments in Time)
This task is intended to evaluate the ability of algorithms to classify events in trimmed 3-second videos. Here, videos contain a single activity, and all clips have a standard duration of 3 seconds. There will be two tracks. The first track will use the Moments in Time dataset, a new large-scale dataset for video understanding, which has 800K videos in the training set. For the second track, participants will use the Moments in Time Mini dataset, a subset of Moments in Time with 100k videos provided in the training set.
DETAILS

At CVPR 2018 (June 22 – All Day)
http://activity-net.org/challenges/2018

Sponsored by:

DeepMind Google KAUST facebook

ActivityNet: A Large-Scale Video Benchmark for Human Activity Understanding [CVPR'15]

Activity Detection @ IVUL

1. **Fast Temporal Activity Proposals for Efficient Detection of Human Actions in Untrimmed Videos** [CVPR’16]
   proposals are represented as sparse combinations of STIPs (10FPS on single CPU core)
2. **DAPs: Deep Action Proposals for Action Understanding** [ECCV’16]
   multi-scale (sparse) proposals are output by an LSTM in one pass (130FPS on single GPU)
3. **SST: Single-Stream Temporal Action Proposals** [CVPR’17]
   multi-scale (dense) proposals are scored by a GRU in one pass + streaming (300FPS on single GPU)

action proposals

4. **SCC: Semantic Context Cascade for Efficient Action Detection** [CVPR’17]
   incorporating objects and scenes in more efficient and accurate activity detection
5. **End-to-End, Single-Stream Temporal Action Detection in Untrimmed Videos** [BMVC’17]
   multi-scale (dense) detector for streaming video (700FPS on single GPU)
6. **Action Search: Spotting Actions in Videos and Its Application to Temporal Action Localization** [ECCV’18]
   learning to detect activities using human search sequences in video
7. **What do I Annotate Next? An Empirical Study of Active Learning for Action Localization** [ECCV’18]
   learnable active learner for efficient annotation and activity detector training
8. **Diagnosing Error in Temporal Action Detectors** [ECCV’18]
   tools to help diagnose detector errors (e.g. localization, double detection, classification, etc.)

action detectors

Activity Detection Examples

Key
Detection
Ground-truth
Time

(Actions are played at 1x speed, Background video is sped up)

Object Tracking @ IVUL

1. **TrackingNet: A Large-Scale Dataset and Benchmark for Object Tracking in the Wild** [ECCV'18]
   large-scale dataset for single object tracking with withheld testing sequences
2. **A Benchmark and Simulator for UAV Tracking** [ECCV'16]
   simulation based tracking benchmark and large dataset for aerial tracking
3. **Context-Aware Correlation Filter Tracking** [CVPR'17] [oral]
   add-on to any correlation filter tracker to discriminate object from context
4. **Target Response Adaptation for Correlation Filter Tracking** [ECCV'16] [spotlight]
   add-on to any correlation filter tracker to dynamically adapt the target per frame
5. **Persistent Aerial Tracking System for UAVs** [IROS'16]
   STRUCK-based tracker for aerial tracking refined in simulation and transferred to real UAVs
6. **In Defense of Sparse Tracking: Circulant Sparse Tracker** [CVPR'16] [spotlight]
   Revisiting LASSO based tracking with efficient FFT solution in dual domain
7. **3D Part-Based Sparse Tracker with Automatic Synchronization and Registration** [CVPR'16]
   sparsity based tracker in 3D exploiting automatic registration from frame-to-frame

tracking datasets

tracking frameworks

sample trackers

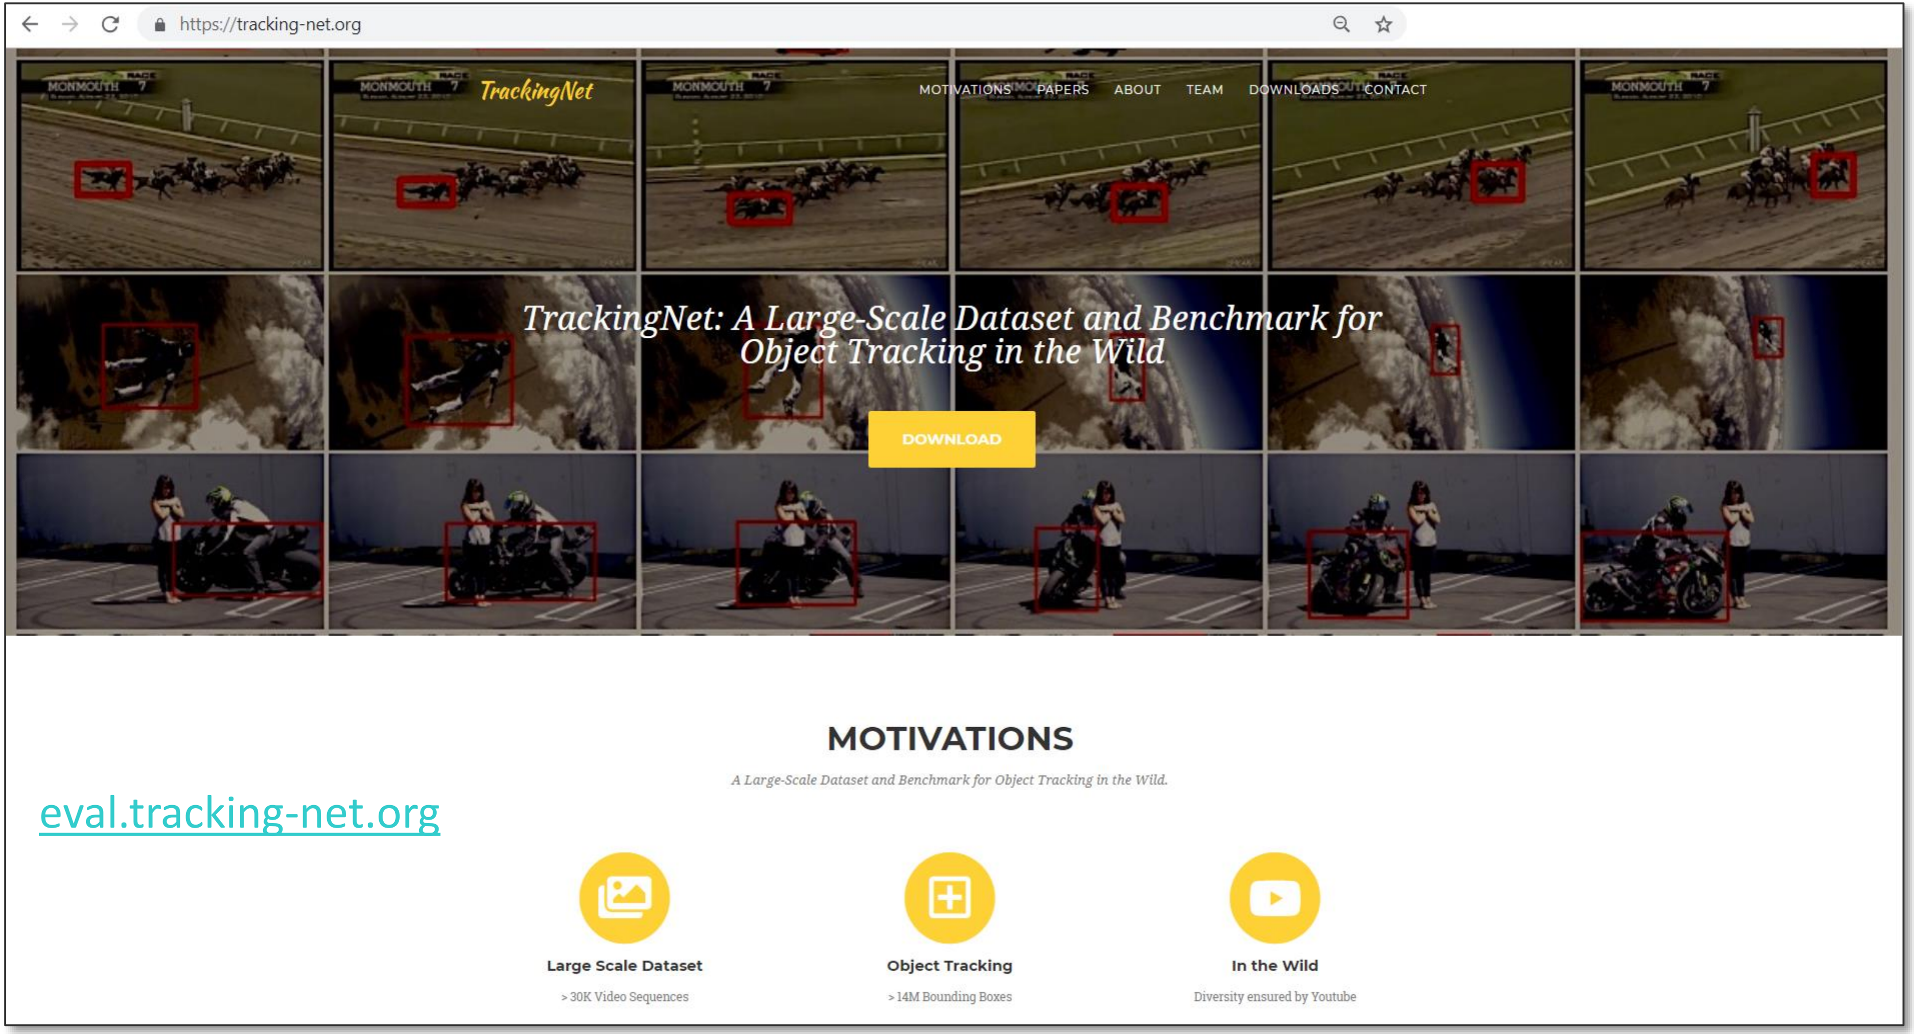

eval.tracking-net.org

Sim4CV: A Photo-Realistic Simulator for Computer Vision Applications [IJCV'18] (www.sim4cv.org)

Self-Driving Car: Scene Generator

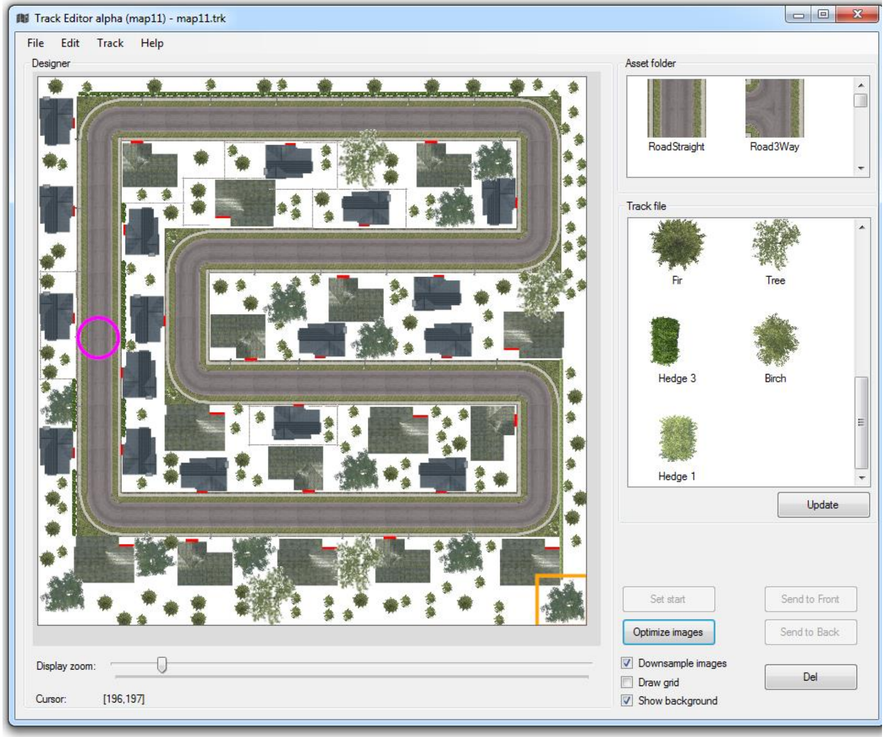

Sim4CV: A Photo-Realistic Simulator for Computer Vision Applications [IJCV'18](www.sim4cv.org)

Single RGB Camera Self-Driving Car Result

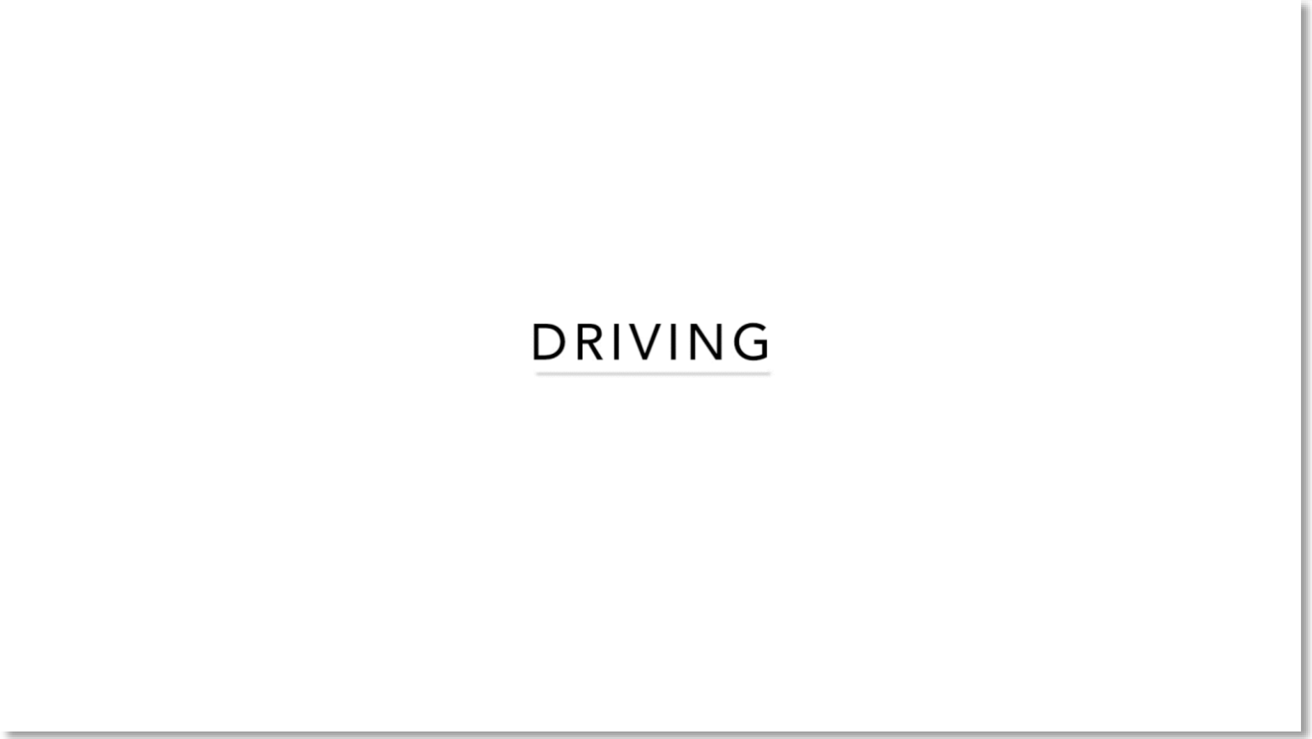

Sim4CV: A Photo-Realistic Simulator for Computer Vision Applications [IJCV'18](www.sim4cv.org)

Self-Driving Car: Real-World Transfer

Driving Policy Transfer via Modularity and Abstraction [CoRL'18][In Collaboration with Intel Labs]

Self-Driving Car: Real-World Transfer

Driving Policy Transfer via Modularity and Abstraction [CoRL'18][In Collaboration with Intel Labs]

Single RGB Camera based Self-Racing UAV

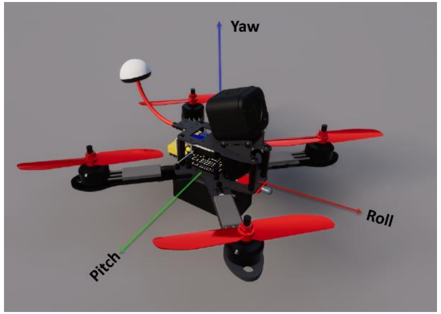

Teaching UAVs to Race: End-to-End Regression of Agile Controls in Simulation [ECCVW’18][Best Paper Award]

THEME: FUNDAMENTALS

Bernard Ghanem

What I'm NOT going to talk about.

$$\min_{\mathbf{c}} \; \|\mathbf{Ac} - \mathbf{b}\|_2^2 + \lambda\|\mathbf{c}\|_1$$

FFTLasso: Large-Scale LASSO in the Fourier Domain [CVPR'17][oral]

$$\min_{\mathbf{x}} \; f(\mathbf{x}) \quad \text{s.t. } \mathbf{x} \in \{1, -1\}^n; \mathbf{x} \in \Omega$$

An Exact Penalty Method for Binary Optimization Based on MPEC Formulation [AAAI'17]

Lp-Box ADMM: A Versatile Framework for Integer Programming [TPAMI'18]

$$\min_{\mathcal{D}, \vec{\mathcal{X}}} \; \frac{1}{2}\sum_{n=1}^{N} \|\vec{\mathcal{Y}}_n - \mathcal{D}\vec{\mathcal{X}}_n\|_F^2 + \lambda\|\vec{\mathcal{X}}_n\|_{1,1,1}$$

High Order Tensor Formulation for Convolutional Sparse Coding [ICCV'17]

$$\min_{\mathbf{x}} \; \left(\frac{1}{2}\mathbf{x}^T\mathbf{A}\mathbf{x} + \mathbf{x}^T\mathbf{b}\right) + h(\mathbf{x})$$

A Matrix Splitting Method for Composite Function Minimization [CVPR'17]

What I AM going to talk about

Analytic Expressions for Probabilistic Moments of PL-DNN with Gaussian Input

Adel Bibi,* Modar Alfadly,* and Bernard Ghanem
King Abdullah University of Science and Technology (KAUST), Saudi Arabia
{adel.bibi,modar.alfadly,bernard.ghanem}@kaust.edu.sa

Oral @CVPR'18

Abstract

The outstanding performance of deep neural networks (DNNs), for the visual recognition task in particular, has been demonstrated on several large-scale benchmarks. This performance has immensely strengthened the line of research that aims to understand and analyze the driving reasons behind the effectiveness of these networks. One important aspect of this analysis has recently gained much attention, namely the reaction of a DNN to noisy input. This has spawned research on developing adversarial input attacks

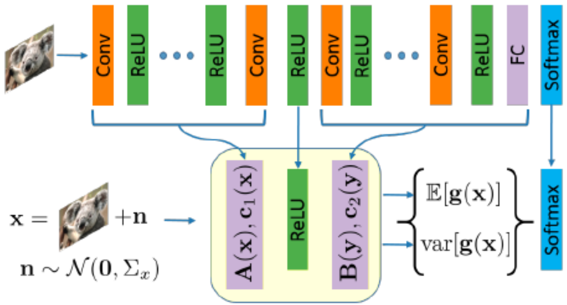

https://github.com/ModarTensai/network_moments

Noise Sensitivity

“panda”

+

=

“gibbon”

“whale”

+

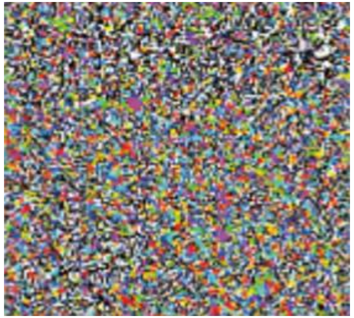

=

“turtle”

Noise Sensitivity

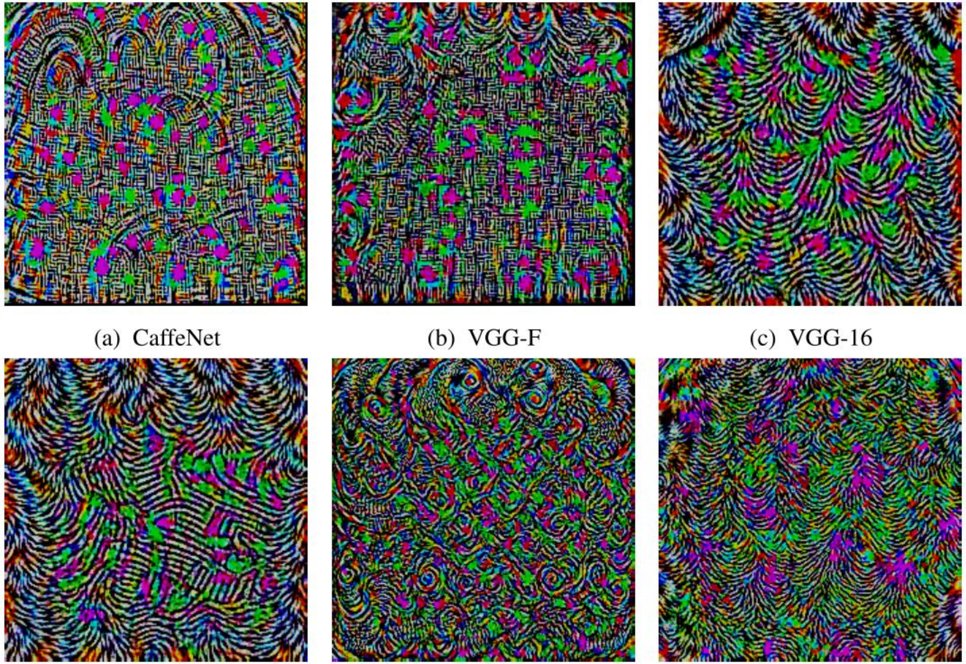

(a) CaffeNet (b) VGG-F (c) VGG-16
(d) VGG-19 (e) GoogLeNet (f) ResNet-152

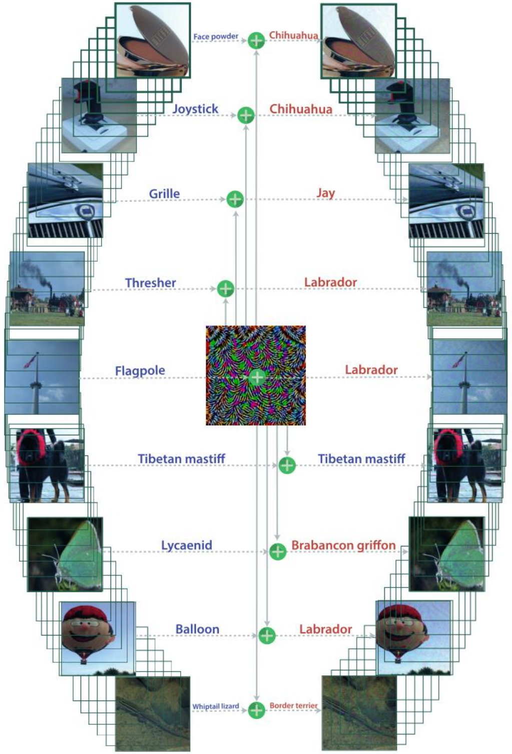

Natural Questions

- Can we derive a closed form expression for the output probability density function? What about the moments?

- Ideally, we want these expressions for any network under any distribution.

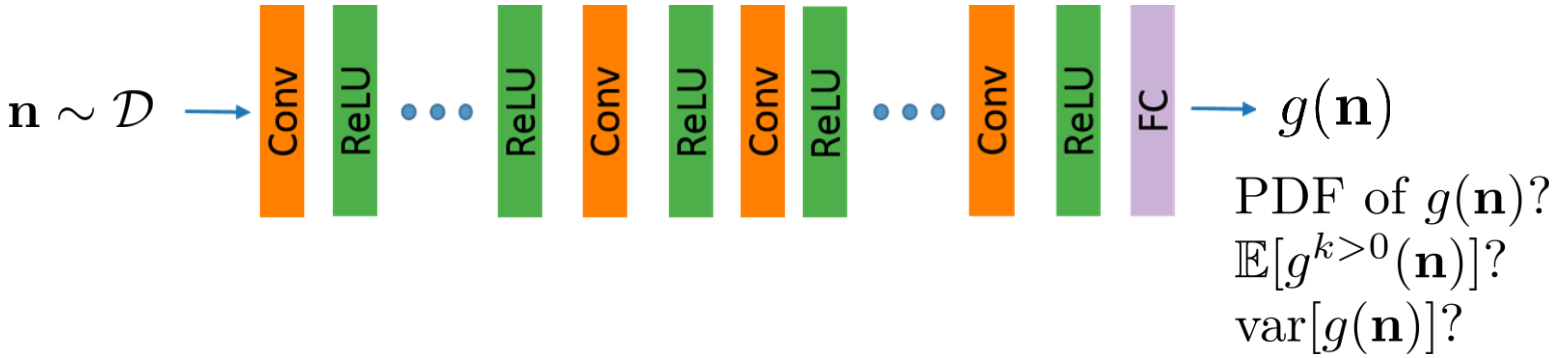

Natural Questions

- Can we derive a closed form expression for the output probability density function? What about the moments?
- Ideally, we want these expressions for any network under any distribution.

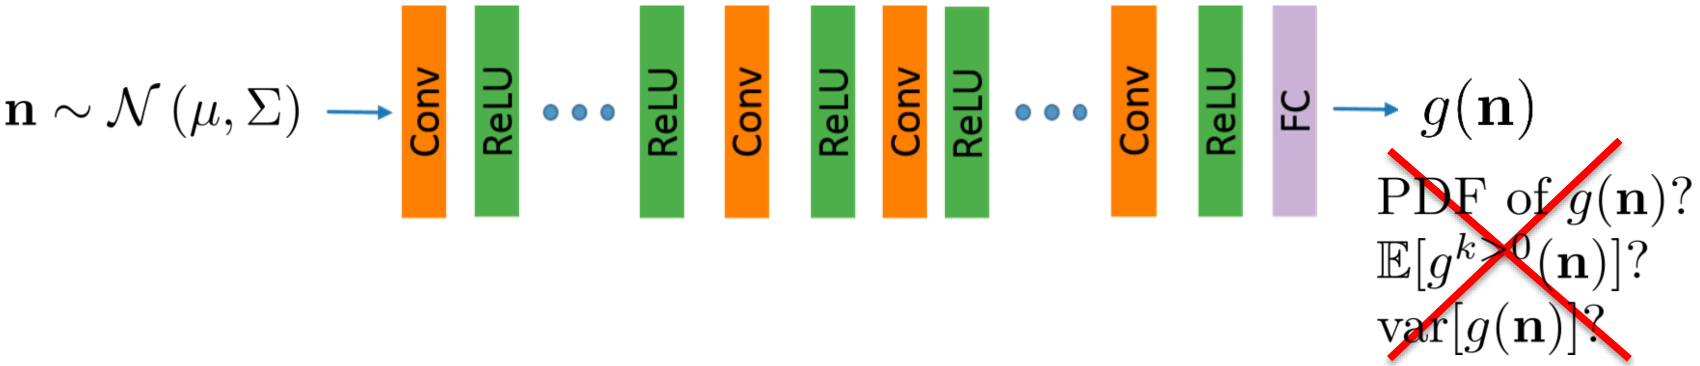

Natural Questions

Maybe that is just too difficult?

Network Moments

Given a Gaussian input, we want to derive analytical expressions for the first and second moments of this shallow piecewise linear NN.

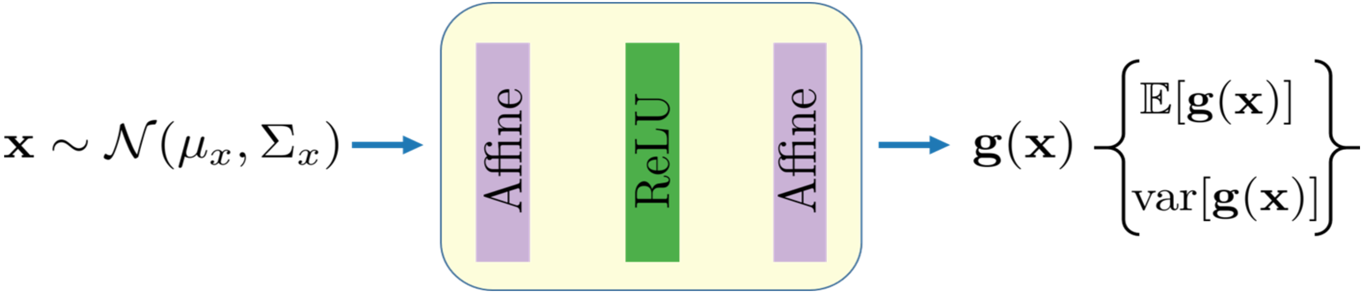

$$\mathbf{g}(\mathbf{x}) = \mathbf{B} \max\left(\mathbf{A}\mathbf{x} + \mathbf{c}_1, \mathbf{0}_p\right) + \mathbf{c}_2$$

where $\mathbf{A} \in \mathbb{R}^{p \times n}$, $\mathbf{B} \in \mathbb{R}^{d \times p}$, $\mathbf{c}_1 \in \mathbb{R}^p$, and $\mathbf{c}_2 \in \mathbb{R}^d$

First Network Moment

$$\mathbf{g}_i(\mathbf{x}) = \mathbf{B}(i,:)\max\left(\mathbf{A}\mathbf{x} + \mathbf{c}_1, \mathbf{0}_p\right) + \mathbf{c}_2(i), \;\; \mathbf{x} \sim \mathcal{N}\left(\mu_x, \Sigma_x\right)$$

Theorem 1. *For any function in the form of* $\mathbf{g}(\mathbf{x})$ *where* $\mathbf{x} \sim \mathcal{N}\left(\mu_x, \Sigma_x\right)$, *we have:*

$$\mathbb{E}[\mathbf{g}_i(\mathbf{x})] = \sum_{v=1}^{p} \mathbf{B}(i,v)\left(\frac{1}{2}\bar{\mu}_v - \frac{1}{2}\bar{\mu}_v \, erf\left(\frac{-\bar{\mu}_v}{\sqrt{2}\bar{\sigma}_v}\right) + \frac{1}{\sqrt{2\pi}}\bar{\sigma}_v \, exp\left(-\frac{\bar{\mu}_v^2}{2\bar{\sigma}_v^2}\right)\right) + \mathbf{c}_2(i)$$

where $\bar{\mu}_v = \left(\mathbf{A}\mu_x + \mathbf{c}_1\right)(v)$, $\bar{\Sigma} = \mathbf{A}\Sigma_x\mathbf{A}^\top$, $\bar{\sigma}_v^2 = \bar{\Sigma}(v,v)$ *and* $erf(x) = \frac{2}{\sqrt{\pi}}\int_0^x e^{-t^2}dt$ *is the error function.*

Second Network Moment (Technical Lemmas)

Lemma 1. *Let* $\mathbf{x} \in \mathbb{R}^n \sim \mathcal{N}(\mu_x, \Sigma_x)$, *for any even* p, *where* $\sigma_{ij} = \Sigma_x(i,j)\ \forall i \neq j$, *and under mild assumptions on the nonlinear map* $\Psi : \mathbb{R}^n \rightarrow \mathbb{R}$, *we have* $\frac{\partial^{\frac{p}{2}} \mathbb{E}[\Psi(\mathbf{x})]}{\prod_{\forall\, odd\, i} \partial \sigma_{i i+1}} = \mathbb{E}[\frac{\partial^p \Psi(\mathbf{x})}{\partial x_1 \ldots \partial x_p}]$.

Second Network Moment (Technical Lemmas)

Lemma 1. *Le*
j, and under

$$\frac{\partial^{\frac{p}{2}} \mathbb{E}[\Psi(\mathbf{x})]}{\prod_{\forall \, odd\, i} \partial \sigma_{i i+1}} =$$

$\Sigma_x(i,j)\ \forall i \neq$
$\to \mathbb{R}$, *we have*

A Useful Theorem for Nonlinear Devices Having Gaussian Inputs*

ROBERT PRICE†

Summary—If and only if the inputs to a set of nonlinear, zero-memory devices are variates drawn from a Gaussian random process, a useful general relationship may be found between certain input and output statistics of the set. This relationship equates partial derivatives of the (high-order) output correlation coefficient taken with respect to the input correlation coefficients, to the output correlation coefficient of a new set of nonlinear devices bearing a simple derivative relation to the original set. Application is made to the interesting special cases of conventional cross-correlation and autocorrelation functions, and Bussgang's theorem is easily proved. As examples, the output autocorrelation functions are simply obtained for a hard limiter, linear detector, clipper, and smooth limiter.

IN THE COURSE of investigating the asymptotic frequency behavior of power spectra resulting from the passage of noise through zero-memory nonlinear devices, an interesting, unique property of Gaussian processes has been encountered, which does not appear to have been previously reported.

Statement of the Theorem

$$\frac{\partial^k R}{\prod_{m=1}^{N} (\partial \rho_{r_m s_m})^{k_m}} = \left(\frac{1}{2}\right)^{\sum_{m=1}^{N} k_m \delta_{r_m s_m}} \overline{\left[\prod_{i=1}^{n} f_i^{\left(\sum_{m=1}^{N} \epsilon_{im} k_m\right)}(x_i)\right]} \quad (3)$$

where r_m and s_m, $m = 1, 2, \cdots, N$, are integers lying between 1 and n, inclusive, and are not necessarily distinct. The k_m are positive integers, with $k = \sum_{m=1}^{N} k_m$. ϵ_{im} is the number of times i appears in (r_m, s_m). $\delta_{r_m s_m}$ is the Kronecker δ function, $\delta_{r_m s_m} = 1$ for $r_m = s_m$, 0 for $r_m \neq s_m$. The symbol $f_i^{(q)}(x_i)$ denotes the qth derivative of $f_i(x)$, taken at x_i.

Furthermore, not only is the above theorem true for inputs having an nth-order joint Gaussian distribution, but it holds true *only* for such inputs if the $f_i(x)$ are allowed to be of general form.

Proof

Second Network Moment

$$\mathbf{g}_i(\mathbf{x}) \approx \mathbf{B}(i,:)\max\left(\mathbf{A}\mathbf{x}, \mathbf{0}_p\right) + \mathbf{c}_2(i), \;\; \mathbf{x} \sim \mathcal{N}\left(\mathbf{0}_n, \Sigma_x\right)$$

Theorem 2. *For any function in the form of* $\mathbf{g}(\mathbf{x})$ *where* $\mathbf{x} \sim \mathcal{N}\left(\mathbf{0}_n, \Sigma_x\right)$ *and that* $\mathbf{c}_1 = \mathbf{0}_p$ *then:*

$$\mathbb{E}[\mathbf{g}_i^2(\mathbf{x})] = 2\sum_{v_1}^{p}\sum_{v_2}^{v_1-1}\mathbf{B}(i,v_1)\mathbf{B}(i,v_2)\left(\frac{\bar{\sigma}_{v_1v_2}}{2\pi}\sin^{-1}\left(\frac{\bar{\sigma}_{v_1v_2}}{\bar{\sigma}_{v_1}\bar{\sigma}_{v_2}}\right) + \frac{\bar{\sigma}_{v_1}\bar{\sigma}_{v_2}}{2\pi}\sqrt{1-\frac{\bar{\sigma}^2_{v_1v_2}}{\bar{\sigma}^2_{v_1}\bar{\sigma}^2_{v_2}}} + \frac{\bar{\sigma}_{v_1v_2}}{4}\right) + \frac{1}{2}\sum_{r}^{p}\mathbf{B}(i,r)^2\bar{\sigma}_r^2 + \mathbf{c}_2(i)$$

Network Moments

$$\mathbf{x} \sim \mathcal{N}(\mu_x, \Sigma_x) \rightarrow$$

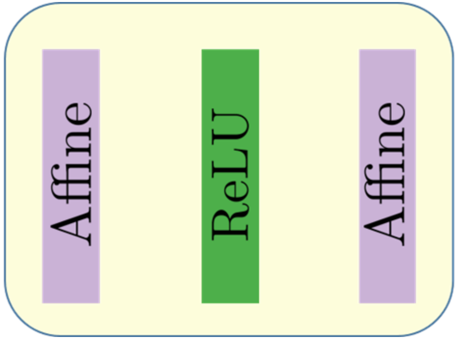

$$\rightarrow \mathbf{g}(\mathbf{x}) \left\{ \begin{matrix} \mathbb{E}[\mathbf{g}(\mathbf{x})] \\ \mathrm{var}[\mathbf{g}(\mathbf{x})] \end{matrix} \right\}$$

$$\mathbf{g}(\mathbf{x}) = \mathbf{B} \max\left(\mathbf{A}\mathbf{x} + \mathbf{c}_1, \mathbf{0}_p\right) + \mathbf{c}_2$$

The mean: $\mathbb{E}[\mathbf{g}_i(\mathbf{x})]$

The variance: $\mathrm{var}[\mathbf{g}_i(\mathbf{x})] \approx \mathbb{E}[\mathbf{g}_i^2(\mathbf{x})] - \mathbb{E}[\mathbf{g}_i(\mathbf{x})]^2|_{\mu_x = \mathbf{0}_n}$

Extending to Deeper Networks

Extending to Deeper Networks

Given any piecewise linear deep neural network (PL-DNN)

Extending to Deeper Networks

We approximate the logits function $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^d$ around a certain input

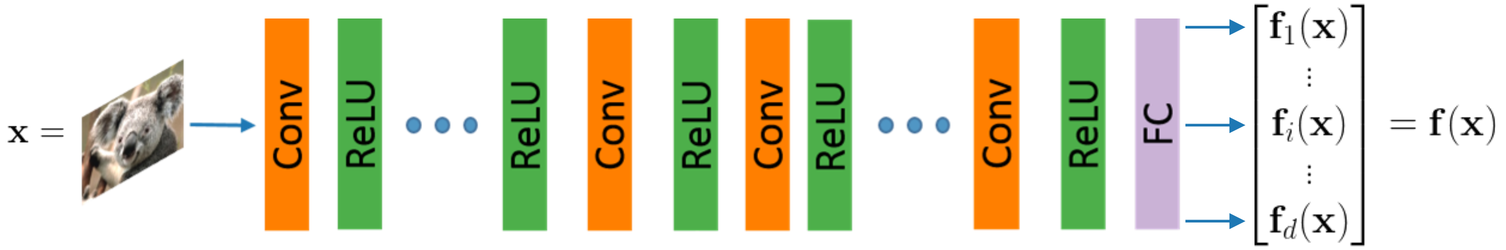

Extending to Deeper Networks

We propose a two-stage linearization strategy at a randomly chosen ReLU

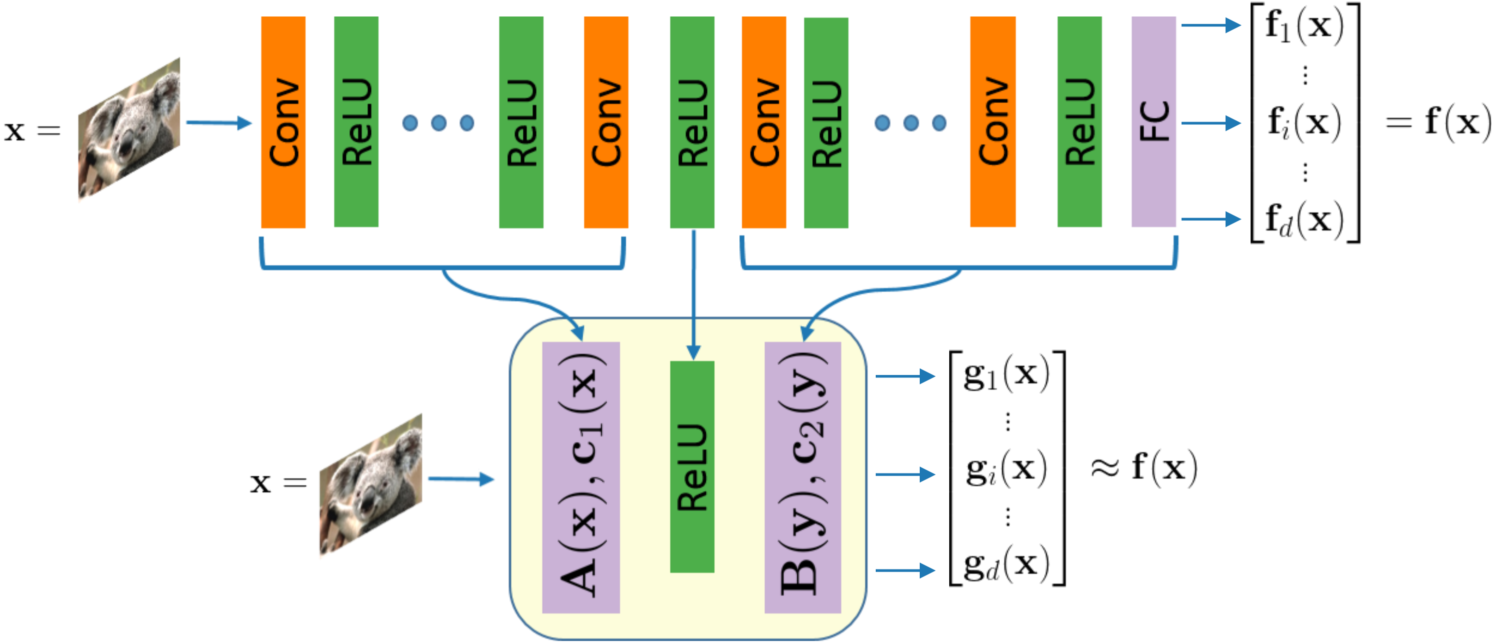

Tightness Verification

Tightness on Synthetic Networks and Data

Verifying tightness by comparing the expressions to Monte Carlo Simulations under various network architectures and various noise regimes.

Fully Connected Networks

Convolutional Networks

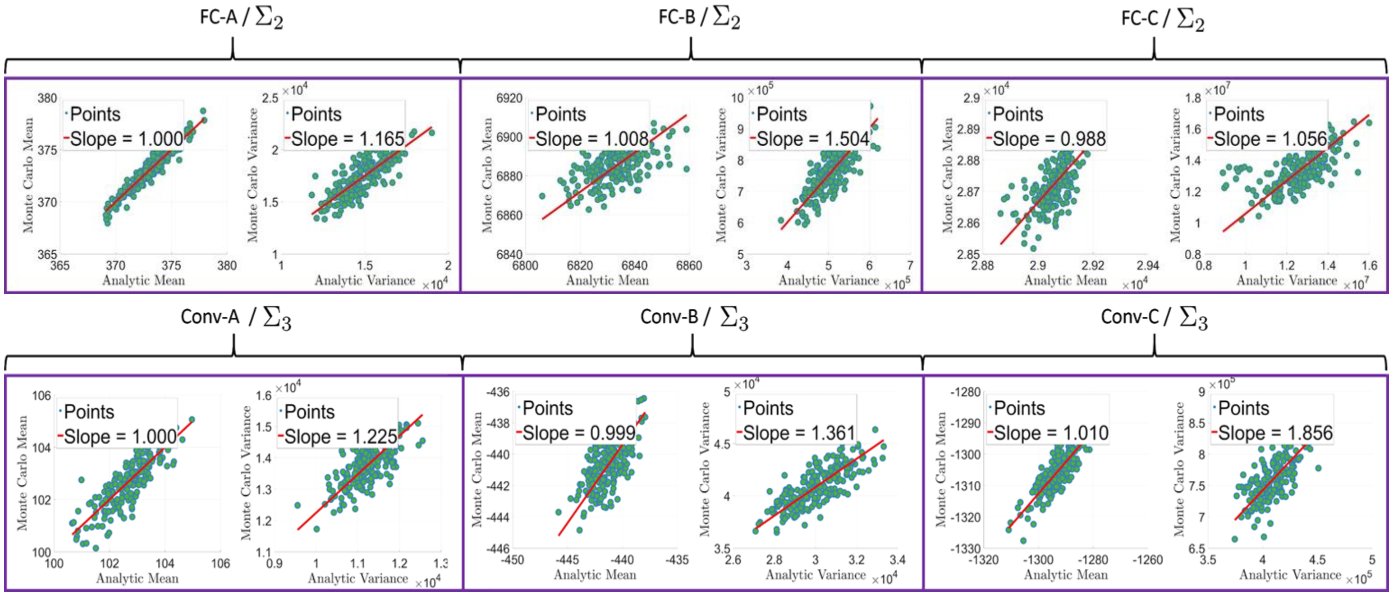

Tightness on AlexNet with ImageNet

Are the expressions tight enough to predict AlexNet top-k score ordering?

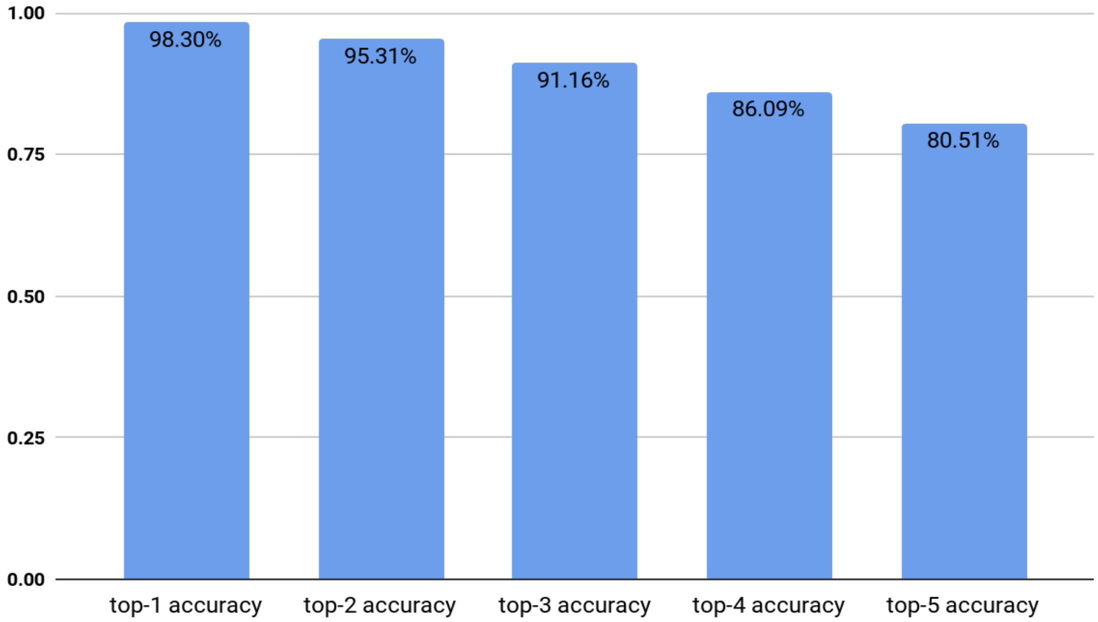

More Experiments

- Tightness on LeNet & explaining other (deterministic) adversarial attacks
- Choice of ReLU for two-stage linearization
- Linearization around cluster centers
- Analyzing the behavior of varying logit score under varying variance

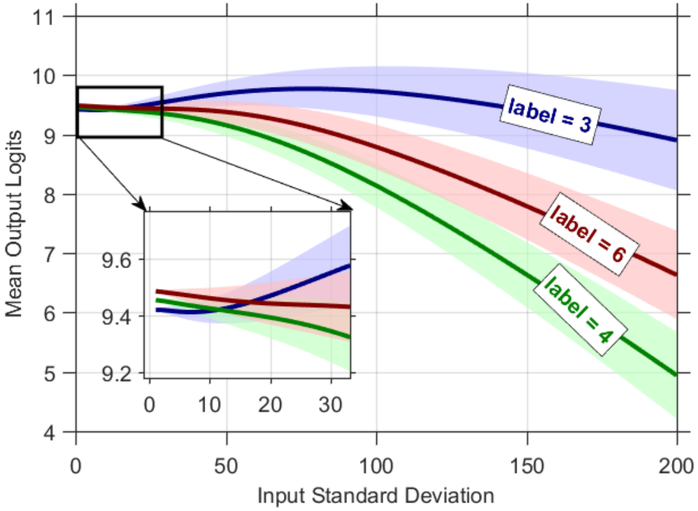

Experiments

Localized Spatial Noise

- Using our expressions, it is now possible to study the effects of adding Gaussian noise around each pixel of the input
- We can visualize a set of heat maps that shows the average fooling rate of LeNet per class label in MNIST validation dataset

Red and blue indicate high and low fooling rates respectively

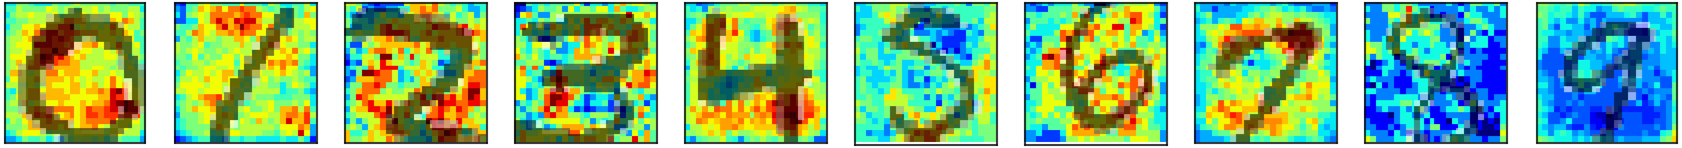

Targeted Adversarial Attack

With $\mathcal{E}_i^{\mathbf{M}}(\mu_x, \sigma^2) = \mathbb{E}[\mathbf{g}_i(\mathbf{M} + \mathbf{x}_{(\mu_x, \sigma^2 \mathbf{I}_n)})]$, we define the targeted attack for image $\mathbf{M}$ to target j as the following optimization:

$$\arg\max_{\mu_x, \sigma} \left(\mathcal{E}_j^{\mathbf{M}}(\mu_x, \sigma^2) - \max_{i \neq j} \left(\mathcal{E}_i^{\mathbf{M}}(\mu_x, \sigma^2) \right) \right) \quad \text{s.t.} \quad 0 < \sigma^2 \leq 2, \; -\beta \mathbf{1}_n \leq \mu_x \leq \beta \mathbf{1}_n$$

Image:

Classified as:	9	2	3	4	7	8

Non-Targeted Adversarial Attack (α- Support)

With μ_x^α having a random support of size $\alpha\%$, we define the non-targeted attack for image $\mathbf{M}$ as the following optimization:

$$\underset{\mu_x^\alpha,\sigma}{\arg\min}\left(\mathcal{E}_i^{\mathbf{M}}(\mu_x^\alpha,\sigma^2) - \max_{j\neq i}\left(\mathcal{E}_j^{\mathbf{M}}(\mu_x^\alpha,\sigma^2)\right)\right) \text{ s.t. } \; 0<\sigma^2\leq 2, \quad -\beta\mathbf{1}_{\alpha n}\leq\mu_x^\alpha\leq\beta\mathbf{1}_{\alpha n}$$

Misclassified MNIST images by LeNet:

4%, ±550

Misclassified ImageNet images by AlexNet:

2%, ±75

Future Work with Network Moments

- More attacks:
  - sparse support optimization (e.g. add L1 regularizer)
  - spatially contiguous attacks (e.g. add TV regularizer)
  - different input noise distributions
  - applications: detection, segmentation, and emotion

- Use in network training
  - No need for noisy data augmentation/sampling

Sneak Peek: Targeted Attacks with Spatially Contiguous Noise

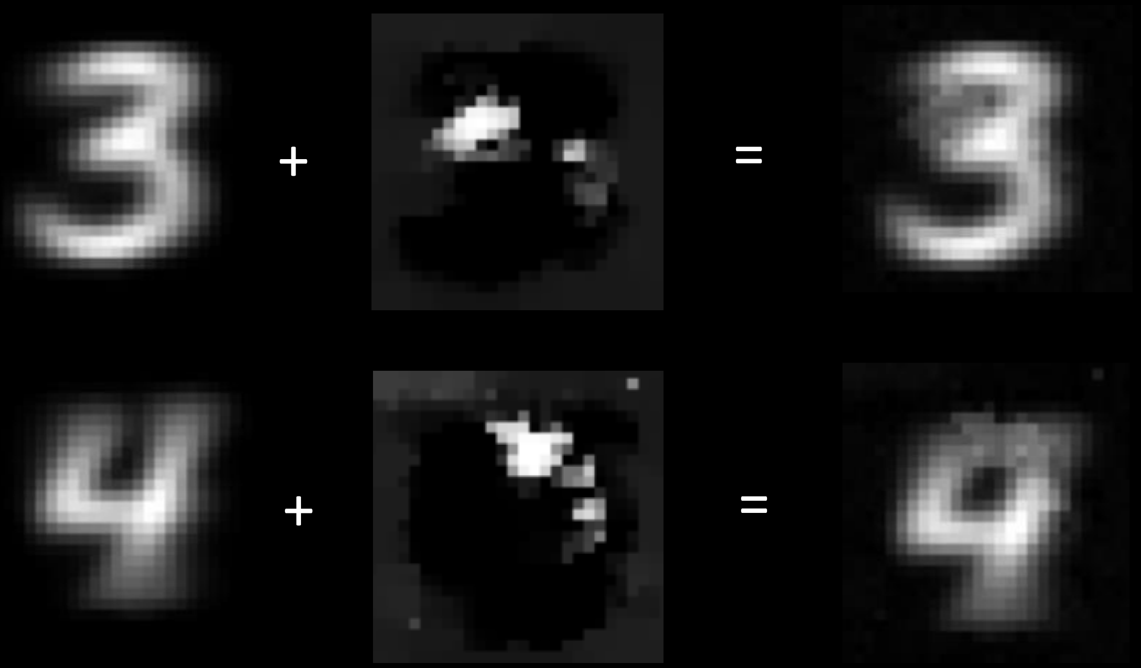

original noise misclassified as 9

IVUL TEAM [ivul.kaust.edu.sa]

PhD


Adel Bibi
Masters to PhD Student
adel.bibi@kaust.edu.sa


Fabian Caba Heilbron
PhD Student
fabian.caba@kaust.edu.sa


Abdullah Hamdi
MS Student
abdullah.hamdi@kaust.edu.sa


Jean Lahoud
PhD Student
jean.lahoud@kaust.edu.sa


Lama Affara
PhD Student
lama.affara@kaust.edu.sa


Humam Alwassel
MS Student
humam.alwassel@kaust.edu.sa


Matthias Mueller
PhD Student
matthias.mueller.2@kaust.edu.sa


Victor Escorcia
PhD Student
victor.escorcia@kaust.edu.sa


Modar Alfadly
MS Student
modar.alfadly@kaust.edu.sa

MS


Noor Bafageeh
MS to PhD Student
noor.bafageeh@kaust.edu.sa


Frost Xu
MS Student
mengmeng.xu@kaust.edu.sa


Salman Alsubaihi
MS Student
salman.subaihi@kaust.edu.sa


Jesus Zarzar
MS Student
jesusalejandro.zarzartorano@kaust.edu.sa

Postdoc


Jian Zhang
Postdoctoral Fellow
jian.zhang@kaust.edu.sa


Silvio Giancola
Postdoctoral Fellow
silvio.giancola@kaust.edu.sa


Chen Zhao
Postdoctoral Fellow
chen.zhao@kaust.edu.sa


Yancheng Bai
Postdoctoral Fellow
yancheng.bai@kaust.edu.sa

Funding Sources

intel LABS
Adobe Research
BOEING
DeepMind
Google AI
NVIDIA
KAUST
Panasonic
QUALCOMM
facebook

Prof. Bernard Ghanem

bernard.ghanem@kaust.edu.sa

ivul.kaust.edu.sa

Image: 

Classified as: 9 2 3 4 7 8