FlowPic: Encrypted Internet Traffic Classification is as Easy as Image Recognition

TEL AVIV UNIVERSITY אוניברסיטת תל אביב

School of Electrical Engineering

TAL SHAPIRA & YUVAL SHAVITT

Why Classification?

- Traffic Classification
  - Traffic engineering
  - Law enforcement
  - Decryption for MITM attacks
- Most of todays traffic is encrypted
  - Looking at the payload does not help
- Port numbers are no longer indicative

Previous Approaches

- Port based methods
- Handcrafted features extraction
  - Identify many 'important' features
  - Feed a feature vector to some tool
  - Usually apply 'classical' supervised learning techniques
- Cross correlation between payload size distributions (PSDs)
- Payload based traffic classification methods (DPI)
  - Apply neural networks over packet payload content

Motivation

- Our goal is to propose **a generic approach** for Internet traffic classification
  - We use the exact same architecture for all the experiments
  - Dealing with all kinds of Internet traffic classification problems
- **Using all time and size related information available** in a network flow, instead of just using information from manually extracted features
- Dealing with **only a time window of a unidirectional flow** instead of the entire bidirectional session
- **Do not rely on the packet payload content**
  - Do not breach privacy
  - Minimalizing storage requirement
  - Classifying VPN and Tor traffic
  - Capturing an intrinsic characteristic of a category behavior, regardless of the encryption technique and a specific user

Our Approach

195.154.82.180_443_10.0.2.15_56113_TCP:
[(40, 0.0), (178, 0.0011), (512, 0.1427), (1012,0.1472),…]

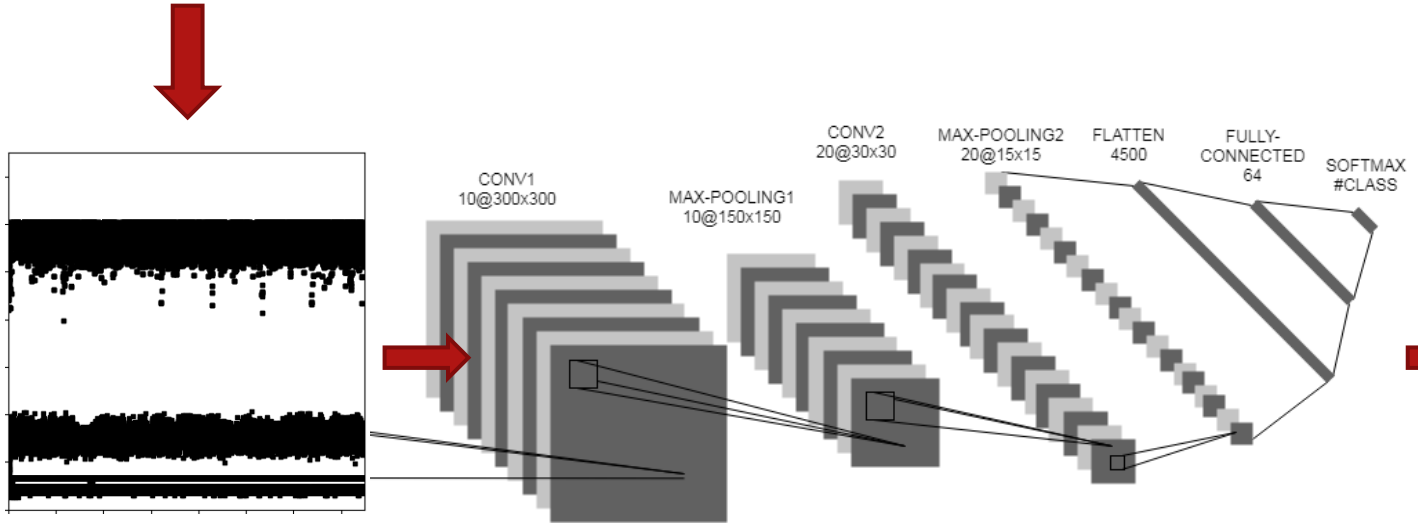

Class is Video

FlowPic Construction

- Extract records from each flow, which comprised of a list of pairs, {IP packet size, time of arrival} for each packet in the flow
- Split each unidirectional flow to equal blocks (15/60 seconds)
- Generate 2D-histogram. For simplicity, we set the 2D-histogram to be a square image
  - We disregard all packets with size greater than the MTU (1500)
  - For the X-axis, first we normalize all time of arrival values to be between 0 and 1500

195.154.82.180_443_10.0.2.15_56113_TCP:
[(40, 0.0), (178, 0.0011), (512, 0.1427), (1012,0.1472),…]

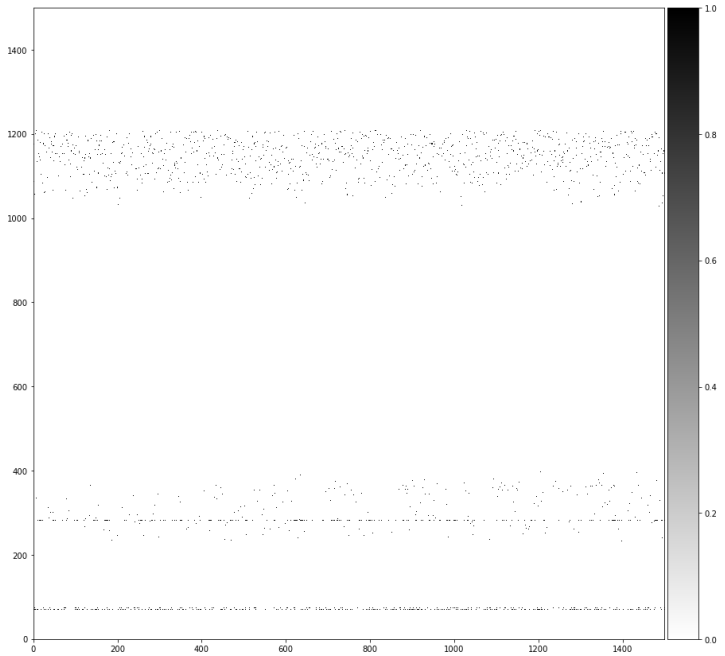

FlowPic Exploration – Traffic Categories

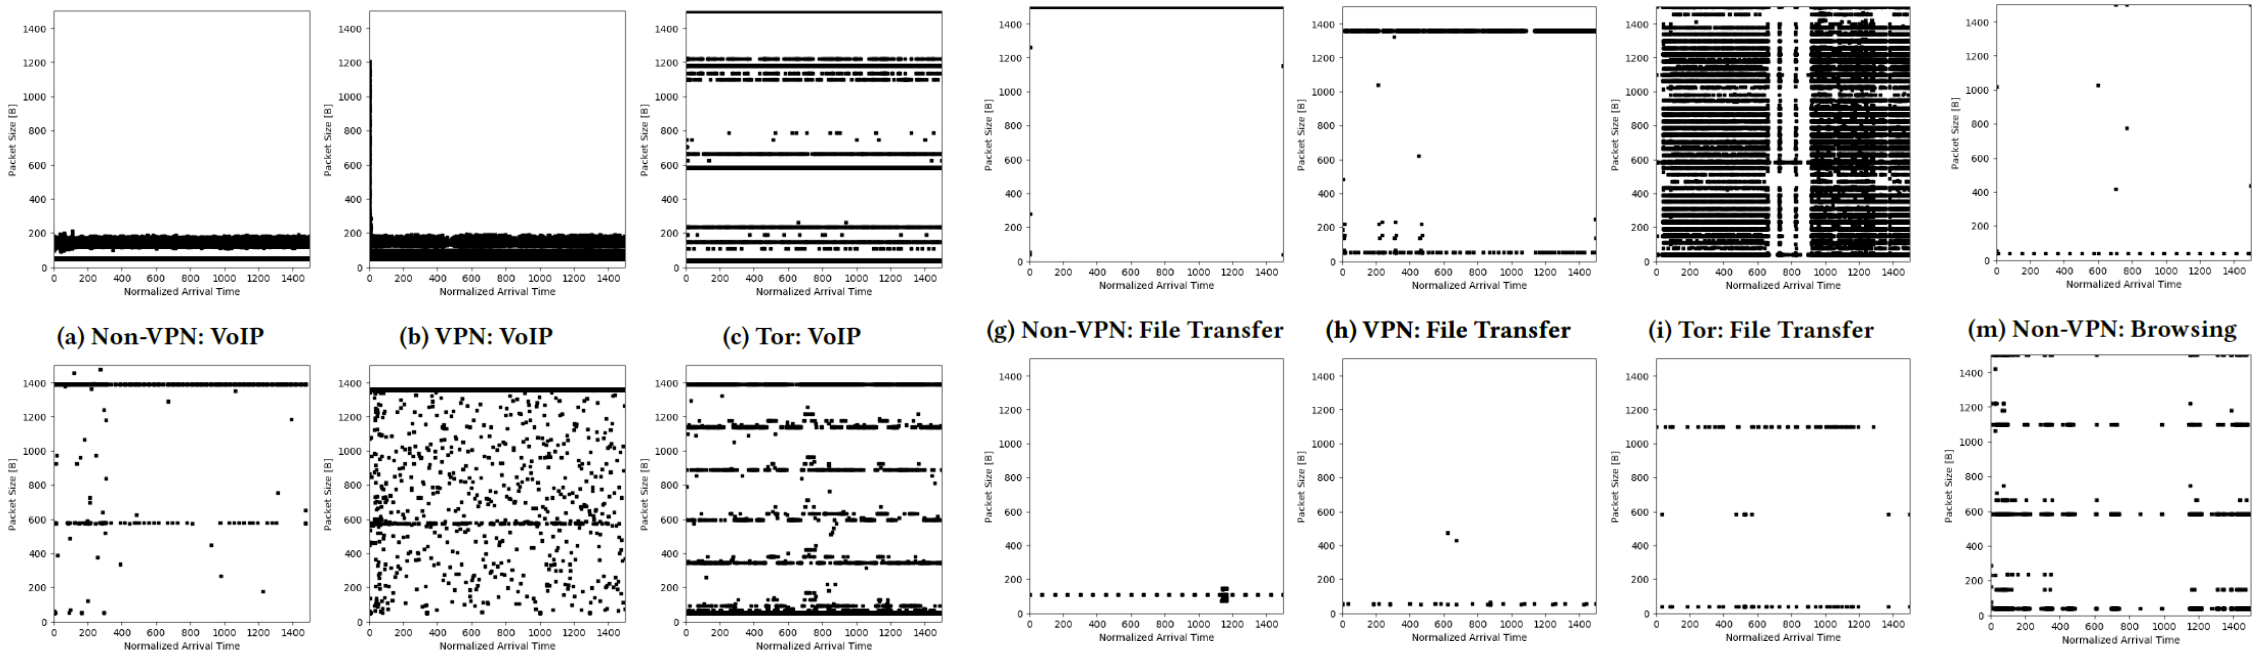

(a) Non-VPN: VoIP (b) VPN: VoIP (c) Tor: VoIP

(g) Non-VPN: File Transfer (h) VPN: File Transfer (i) Tor: File Transfer (m) Non-VPN: Browsing

(d) Non-VPN: Video (e) VPN: Video (f) Tor: Video

(j) Non-VPN: Chat (k) VPN: Chat (l) Tor: Chat (n) Tor: Browsing

FlowPic Exploration – Video Applications

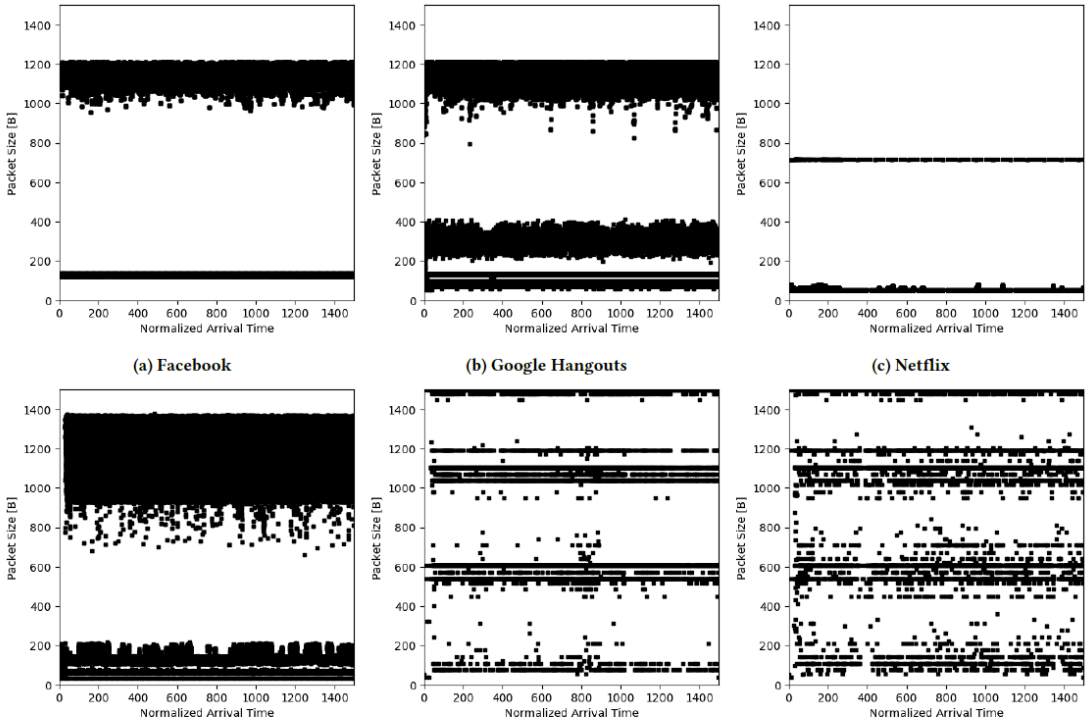

(a) Facebook (b) Google Hangouts (c) Netflix

(d) Skype (e) Vimeo (f) YouTube

FlowPic: Encrypted Internet Traffic Classification is as Easy as Image Recognition

Our CNN Architecture

- A *LeNet-5** style architecture
  - #params=309,094 + 65m
- Dropout

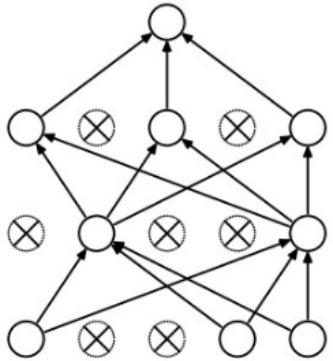

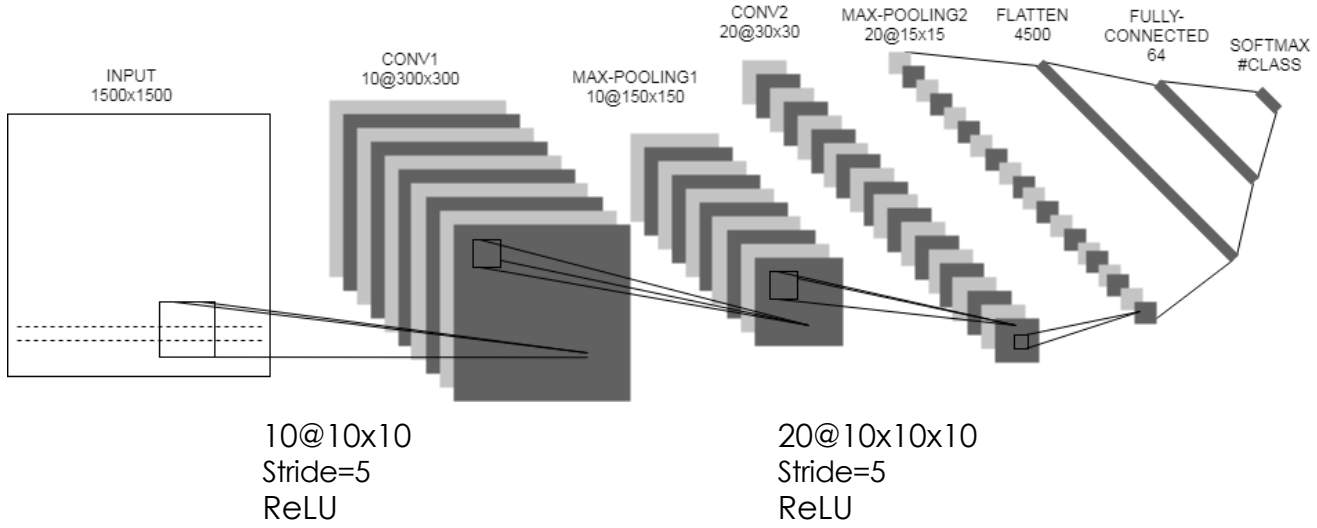

- Softmax final layer $\sigma(\mathbf{z})_j = \frac{e^{z_j}}{\sum_{k=1}^{K} e^{z_k}}$ for $j = 1, \ldots, K$.

***[Lecun et. al 1998]** *"Gradient-based learning applied to document recognition"*

Specifications

- *Categorical cross entropy* loss function
- Using the *Adam* gradient-based optimizer with default hyper-parameters
- We build and run our networks using Keras *and Tensorflow frameworks*
- We set our batch size to 128, and run our network for 40 epochs of 10 batches each

$$\begin{bmatrix}0.1\\0.5\\0.4\end{bmatrix} \quad D(\hat{\mathbf{y}},\mathbf{y}) = -\sum_j y_j \ln \hat{y}_j \quad \begin{bmatrix}0\\1\\0\end{bmatrix}$$

Dataset

- Our dataset consists of the two datasets from Uni. of New Brunswick (UNB):
  - "ISCX VPN-nonVPN traffic dataset" (ISCX-VPN)*
  - "ISCX Tor-nonTor dataset" (ISCX-Tor) **
- In addition we generate our own small packet capture

*[**Gil et al. 2016]** "*Characterization of Encrypted and VPN Traffic using Time-related Features*"
[Lashkari et al. 2017]** "*Characterization of Tor Traffic using Time based Features*"

Dataset

	Non-VPN	VPN	Tor
VoIP	Google Hangouts, Facebook, VoipBuster and Skype	Google Hangouts, VoipBuster and Skype	Google Hangouts, Facebook and Skype
Video	Google Hangouts, Facebook, Netflix, Vimeo, YouTube and Skype	Netflix, Vimeo and YouTube	Vimeo and YouTube
File Transfer	FTPS, SCP, SFTP and Skype	FTPS, SFTP and Skype	FTP, SFTP and Skype
Chat	Google Hangouts, Facebook, AIM Chat, Skype, ICQ and WhatsApp Web	Google Hangouts, Facebook, AIM Chat, Skype and ICQ	Google Hangouts, Facebook, AIM Chat, Skype and ICQ
Browsing	Firefox and Chrome	-	Firefox and Chrome

	Non-VPN	VPN	Tor
VoIP	3304	2872	2978
Video	1553	302	754
File Transfer	1174	242	1126
Chat	635	1061	422
Browsing	3191	-	2026

Experiments

- Multiclass classification experiments :
  - **Traffic categories**: VoIP, Video, Chat, File Transfer, Browsing over 4 datasets: non-VPN, VPN, Tor and merged dataset
  - **Encryption techniques**: non-VPN, VPN, Tor
  - **Application identification**: 10 classes of VoIP and Video applications
- **Class vs. all** classification experiments :
  - 5 classes: VoIP, Video, Chat, File Transfer, Browsing
  - 4 types of datasets: non-VPN, VPN, Tor and merged
- Classification of an Unknown Application

Summary of Traffic categorization Results

- Former results:
  - **The UNB group** gained a best average precision of:
    - 84.0% for non-VPN traffic categorization
    - 89.0% for VPN traffic categorization
    - 84.3% for Tor traffic categorization
  - **Wang et al.** achieved a best accuracy of:
    - 83.0% for non-VPN traffic categorization
    - 98.6% for VPN traffic categorization
- Furthermore, we test the network that was **trained over the merged** set over non-VPN, VPN and Tor set experiment achieves an accuracy of 88.2%, 98.4% and 67.8%, respectively

Multiclass Dataset	Top-1 Acc. (%)	Top-2 Acc. (%)
Traffic Categories over Non-VPN	85.0	99.4
Traffic Categories over VPN	98.4	100.0
Traffic Categories over Tor	67.8	81.0
Traffic Categories over Merged Dataset	83.0	95.1

Class vs. All Results

- We get an average accuracy of:
  - **97.0%** for non-VPN
  - **99.7%** for VPN
  - **85.7%** for Tor
- There are no former results

Class	Accuracy (%)			
VoIP	Training/Test	Non-VPN	VPN	Tor
	Non-VPN	**99.6**	99.4	48.2
	VPN	95.8	**99.9**	58.1
	Tor	52.1	35.8	**93.3**
Video	Training/Test	Non-VPN	VPN	Tor
	Non-VPN	**99.9**	98.8	83.8
	VPN	54.0	**99.9**	57.8
	Tor	55.3	86.1	**99.9**
File Transfer	Training/Test	Non-VPN	VPN	Tor
	Non-VPN	**98.8**	79.9	60.6
	VPN	65.1	**99.9**	54.5
	Tor	63.1	35.8	**55.8**
Chat	Training/Test	Non-VPN	VPN	Tor
	Non-VPN	**96.2**	78.9	70.3
	VPN	71.7	**99.2**	69.4
	Tor	85.8	93.1	**89.0**
Browsing	Training/Test	Non-VPN	VPN	Tor
	Non-VPN	**90.6**	-	57.2
	VPN	-	-	-
	Tor	76.1	-	**90.6**

Class	Accuracy (%)
VoIP	99.7
Video	99.9
File Transfer	77.4
Chat	89.3
Browsing	90.2

Classification of an Unknown Application

- Video vs. all dataset, where we exclude all Facebook's videos
  - We achieves an accuracy of **99.9%**
- VoIP vs. all dataset, where we exclude all Facebook's VoIPs
  - We achieves an accuracy of **96.3%**
- There are no former results

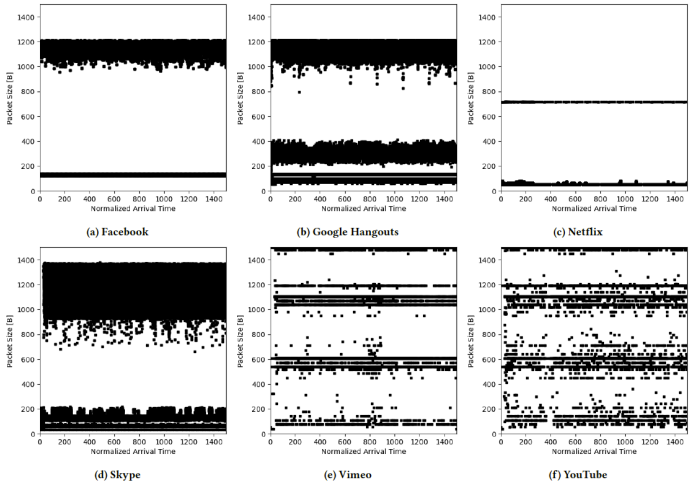

Applications Identification Results

- Our network achieves an accuracy of **99.7%** classifying 10 VoIP and Video applications

- Former Results:
  - **Yamansavascilar et al.** constructed 111 flow features and used k-NN algorithm to achieve an overall accuracy of 93.9% classifying 14 classes of applications
  - For example, they achieved an accuracy of 45.1%, 80.10% and 86.30% classifying Skype-VoIP, Vimeo and YouTube, respectively

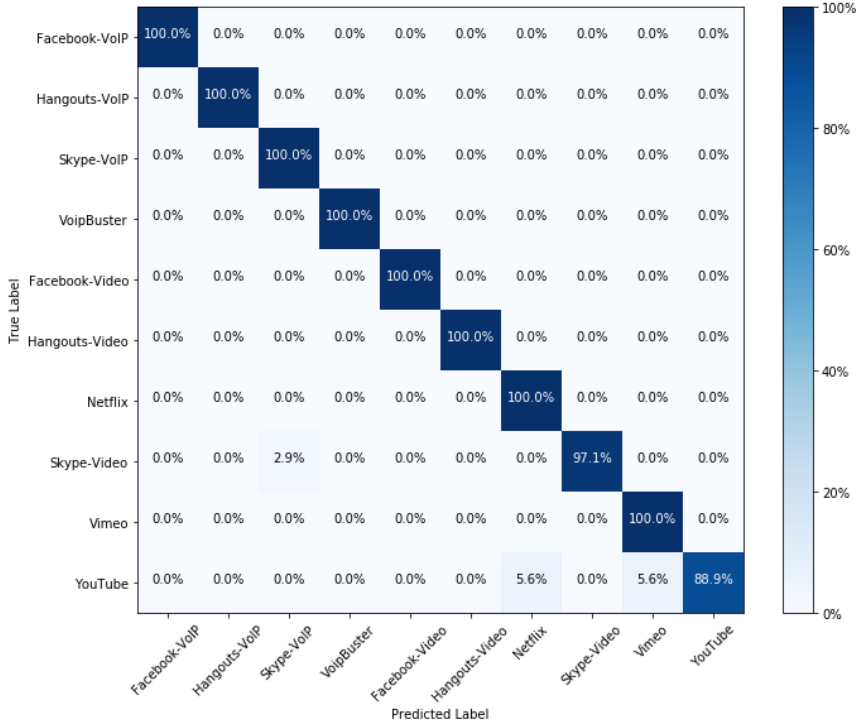

Future Research

- Optimize our CNN architecture or examine other well known architectures
- Reduce run-time and memory consumption by:
  - Using binary images
  - Reducing FlowPic size
  - Using blocks with a shorter time duration
  - Relying only on a sub-sampling of the flow

Summary

- A **generic approach** for Internet traffic classification
  - We use the exact same CNN architecture for all the experiments
  - **Dealing with all kinds of Internet traffic classification problems**; traffic categorization, application identification, encryption techniques
- We are the first to show:
  - **Identifying traffic category of an unfamiliar application**, by learning samples of other applications of the same traffic category
  - Classifying different Internet traffic categories that pass through different encryption techniques by learning traffic that pass throw other encryption techniques
- **Fast & simple**: Classification is done after the first **15 seconds of a unidirectional flow**
- **No need for manual extraction of features**
- Good results for traffic over VPN and Tor traffic

Thank you for your attention. Questions?

TEL AVIV UNIVERSITY אוניברסיטת תל אביב

School of Electrical Engineering