

Poster: Estimation of Traffic Matrices via Super-resolution and Federated Learning

Roberto Amoroso
roberto.amoroso@slu.edu
Saint Louis University

Flavio Esposito
flavio.esposito@slu.edu
Saint Louis University

Maria Luisa Merani
marialuisa.merani@unimore.it
Univ. of Modena and Reggio Emilia

ABSTRACT

Network measurement and telemetry techniques are central to the management of today's computer networks. One popular technique with several applications is the estimation of traffic matrices. Existing traffic matrix inference approaches that use statistical methods, often make assumptions on the structure of the matrix that may be invalid. Data-driven methods, instead, often use detailed information about the network topology that may be unavailable or impractical to collect.

Inspired by the field of image processing, we propose a super-resolution technique for traffic matrix inference that does not require any knowledge on the structural properties of the matrix elements to infer, nor a large data collection. Our experiments with anonymized Internet traces demonstrate that the proposed approach can infer fine-grained network traffic with high precision outperforming existing data interpolation techniques, such as bicubic interpolation.

CCS CONCEPTS

• **Computing methodologies** → **Machine learning**; • **Networks** → **Network measurement**; **Network management**.

KEYWORDS

traffic engineering, machine learning

ACM Reference Format:

Roberto Amoroso, Flavio Esposito, and Maria Luisa Merani. 2020. Poster: Estimation of Traffic Matrices via Super-resolution and Federated Learning. In *The 16th International Conference on emerging Networking EXperiments and Technologies (CoNEXT '20)*, December 1–4, 2020, Barcelona, Spain. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3386367.3431677>

1 INTRODUCTION

The improvement of computer network efficiency, once low-hanging fruits have been harvested, necessarily passes through advanced analytics and non-trivial timely troubleshooting. Collecting the right amount of traffic at the right time is, however, a challenging operation, given the scale of today's networks and the "hidden" network spots. For example, a single high-speed network interface within a given region could generate hundreds of gigabytes

of (un-sampled) flow statistics per day if fully utilized, while the whole network might generate several gigabytes of simple network management protocol (SNMP) statistics per day [1, 3].

To reduce the costs of fine-grained measurements and storage needs, the researchers looked for ways to infer network patterns from the least amount of data stored and collected [6], sometimes because such data is not available, since, before the (performance) problem arose, the frequency of measurements was kept low.

A goal of this work is to explore how a computer vision technique called super-resolution can help with network monitoring and telemetry operations, in general, and to infer fine-grain traffic details only using aggregate measurements, in particular. In our approach, we consider a Low Resolution (LR) image, representing the traffic matrix in pixels (heatmap). Using a Super-Resolution (SR) technique, we can then reconstruct a more accurate, i.e., High Resolution (HR) representation of the traffic that we would have seen on a given subnetwork, with the smallest possible inference error. To do so, we apply a supervised learning problem and train a deep neural network with several high-resolution traffic images together with their low-resolution counterparts. We then use the trained neural network to infer the high-resolution version when only a low-resolution version is available.

The proposed technique is effective and could be applied to many network telemetry operations for network debugging and performance diagnosis. We are not the first to apply super-resolution techniques to networking problems. Existing work, however, is limited, and used mobile traffic geolocations [7]. In this project, we propose its use as a traffic volume inference technique. The somehow surprising result that we present in this project is the fact that, despite the lack of geoproximity, present in existing work, heatmaps resembling traffic matrices still preserve some features that we can exploit in our inference problem. In particular, we apply super-resolution to images obtained from computer network traffic matrices, proposing a *Traffic Matrix Super-Resolution (TMSR)*. In Figure 1, we illustrate an example of super-resolution when applied to a standard image (example from the DIV2K dataset [4]) and a pixel representation of the traffic matrix of a real network [5].

We assume that "enough" high-resolution images of traffic matrices are available for training. In large networks, however, large data transfers, and the security and privacy constraint may render this approach impractical. Moreover, some service providers have customers that are unwilling or incapable of sharing their local measurement data. To cope with these practical problems, we exploit federated learning [2], a fairly recent distributed machine learning technique that allows training by sending copies of a model where data resides (as opposed to sending the data to the model).

We analyze the benefits of training our Traffic-Matrix Super Resolution model with federated learning. With some surprising

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CoNEXT '20, December 1–4, 2020, Barcelona, Spain

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7948-9/20/12...\$15.00

<https://doi.org/10.1145/3386367.3431677>

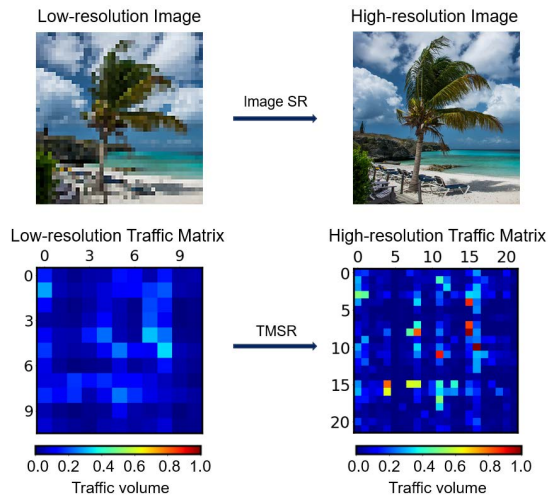


Figure 1: Comparison between images reconstructed using a Super-Resolution technique, on a standard image (top), and on a pixel representation of a Traffic Matrix (bottom).

results, our simulations using real datasets show the effectiveness of our approach to improve the accuracy performance compared to a super-resolution model trained in a centralized way. In particular, we obtained better traffic inference accuracy results with an often significantly lower number of training epochs.

2 INITIAL EVALUATION RESULTS

We started comparing the low-resolution (LR) input and the relative high-resolution (HR) ground truth (Figure 2); the traffic matrices are generated using bicubic interpolation and with our proposed *TMSR* method, both in its centralized and federated version. Each sub-figure shows snapshots of the predictions made, when working with different scale factors: $\times 2$, $\times 3$, and $\times 6$. A scale factor of 2 indicates that the low-resolution traffic matrix image has been down-sampled by a factor of 2 from its high-resolution original version. In particular, in Figure 2 we observe that our models predict traffic accurately for all scale factors. As expected, the accuracy of the inference decreases as the scale factor increases. The texture and details are, however, almost perfectly recovered for the scale factor $\times 2$. In contrast, the bicubic technique, although improving the resolution of network traffic “snapshots”, loses details and deviates more from the ground truth measurements. These initial results show how the proposed architecture is able to estimate traffic matrices with an 87% lower prediction error than bicubic interpolation. When we use a scale factor of $\times 6$, our model is able to estimate network traffic with high accuracy, having to process, save, and transmit only $(\frac{1}{6 \times 6} = \frac{1}{36})$ -th of the original traffic volume data. Moreover, to obtaining lower values of the loss function, a measure of the model accuracy, we train the system using a federated model and obtain the same results as in the centralized model, but performing 80% fewer training rounds (Figure 3).

ACKNOWLEDGMENTS

This work has been partially supported by Comcast and by NSF awards CNS-1647084, CNS-1836906, and CNS-1908574.

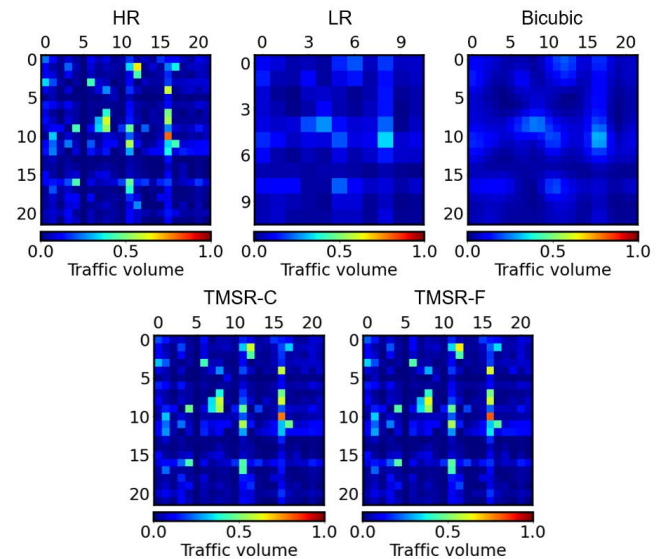


Figure 2: Qualitative comparison of our models with the Bicubic interpolation technique for scale factors $\times 2$. *TMSR-C* and *TMSR-F* refer to the *TMSR* model in its centralized and federated version respectively.

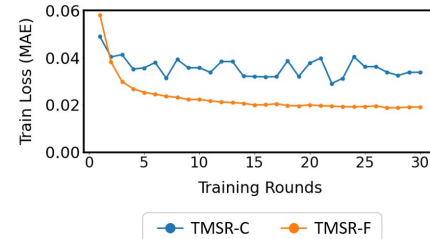


Figure 3: Mean Absolute Error during model training. Centralized *TMSR* architecture (*TMSR-C*) and federated counterpart (*TMSR-F*). Each model is trained for the scale factor $\times 6$.

REFERENCES

- [1] Nick Duffield et al. 2004. Sampling for passive internet measurement: A review. *Statist. Sci.* 19, 3 (2004), 472–498.
- [2] H Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, et al. 2016. Communication-efficient learning of deep networks from decentralized data. *arXiv preprint arXiv:1602.05629* (2016).
- [3] Erik Nygren, Ramesh K Sitaraman, and Jennifer Sun. 2010. The akamai network: a platform for high-performance internet applications. *ACM SIGOPS Operating Systems Review* 44, 3 (2010), 2–19.
- [4] Radu Timofte, Eirikur Agustsson, Luc Van Gool, Ming-Hsuan Yang, and Lei Zhang. 2017. Ntire 2017 challenge on single image super-resolution: Methods and results. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*. 114–125.
- [5] Steve Uhlig, Bruno Quoitin, Jean Lepropre, and Simon Balon. 2006. Providing public intradomain traffic matrices to the research community. *Computer Communication Review* 36, 1 (2006), 83–86. <https://doi.org/10.1145/1111322.1111341>
- [6] Minlan Yu, Lavanya Jose, and Rui Miao. 2013. Software Defined Traffic Measurement with OpenSketch. In *Proc. of NSDI*. USENIX, Lombard, IL, 29–42.
- [7] Chaoyun Zhang, Xi Ouyang, and Paul Patras. 2017. ZipNet-GAN: Inferring fine-grained mobile traffic patterns via a generative adversarial neural network. In *CoNEXT 2017 - Proceedings of the 2017 13th International Conference on emerging Networking Experiments and Technologies*, Vol. 17. 363–375. <https://doi.org/10.1145/3143361.3143393> arXiv:1711.02413