

Privacy-preserving Machine Learning Based Data Analytics on Edge Devices

Jianxin Zhao and Richard Mortier and Jon Crowcroft and Liang Wang

The Computer Laboratory, University of Cambridge
William Gates Building, 15 JJ Thomson Avenue, Cambridge CB3 0FD
{firstname.lastname}@cl.cam.ac.uk

Abstract

Emerging Machine Learning (ML) techniques, such as Deep Neural Network, are widely used in today's applications and services. However, with social awareness of privacy and personal data rapidly rising, it becomes a pressing and challenging societal issue to both keep personal data private and benefit from the data analytics power of ML techniques at the same time. In this paper, we argue that to avoid those costs, reduce latency in data processing, and minimise the raw data revealed to service providers, many future AI and ML services could be deployed on users' devices at the Internet edge rather than putting everything on the cloud. Moving ML-based data analytics from cloud to edge devices brings a series of challenges. We make three contributions in this paper. First, besides the widely discussed resource limitation on edge devices, we further identify two other challenges that are not yet recognised in existing literature: lack of suitable models for users, and difficulties in deploying services for users. Second, we present preliminary work of the first systematic solution, i.e. Zoo, to fully support the construction, composing, and deployment of ML models on edge and local devices. Third, in the deployment example, ML service are proved to be easy to compose and deploy with Zoo. Evaluation shows its superior performance compared with state-of-art deep learning platforms and Google ML services.

Introduction

Machine Learning (ML) techniques have begun to dominate data analytics applications and services. Recommendation systems are the driving force of online service providers such as Amazon, Netflix and Spotify. Finance analytics has quickly adopted ML to harness large volume of data in such areas as fraud detection, risk-management, and compliance. Deep Neural Network (DNN) is the technology behind voice-based personal assistance (Schuster 2010), self-driving cars (Bojarski et al. 2016), automatic image processing (Google 2017a), *etc.* By deploying ML technologies to cloud computing infrastructures, they are benefiting numerous aspects of our daily life.

However, the surge of ML is accompanied by a public concern of personal data privacy. The basic business model of ML-based data analytics is a closed circle: with more data, more accurate model can be trained, which means more

useful services and more users, and they finally lead to more data being collected. At the same time, people are increasingly aware of the data privacy issue. Ubiquity of sensing via mobile and IoT devices has caused a surge in personal data generation and use. They contain our photos, browsing history, and voice records, *etc.*, some of which one might not want to share with data analytics service providers. But these data are also perfect targets for them to collect in order to provide personalised services. So even with legal regulatory frameworks such as EU's General Data Protection Regulation (European-Commission 2016), the aforementioned business model cycle is still difficult to break.

Existing solutions that aim at solving this data privacy issue mostly focus on improving the cloud-based services, including making databases more secure (Cuzzocrea, Mastroianni, and Grasso 2016), hiding trained ML models (Abadi et al. 2016), enable users to choose which part of data to remove before uploading to cloud servers (Xu et al. 2017), *etc.*

However, we argue that patching existing cloud-based data analytics is not enough. Take the speech recognition application as an example. Current approaches require user to upload audio data to server side for inference, which returns the results back to user via the Internet. Aside from the data privacy issue, this workflow also leads to network communication cost and response latency, especially when users are limited by time or bandwidth budgets. In this paper, "users" could be individuals who are looking for a one-off speech recognition tool, business such as banks that seeks to keep record of internal conferences, or researchers who want to preprocess their large audio dataset before some experiments. In all these cases, the priority is data privacy, integrity or service response time.

One solution is to deploy ML services on edge devices. Moving services from cloud to users' edge devices can keep the data private, and effectively reduce the communication cost and response latency. Some research begin to emerge that aims to solve accompanied challenges. They recognise that mobile devices cannot afford to support most of today's intelligent systems because of the large amount of computation resource and memory required. As a result, many current end-side services only support simple ML models. These solutions mostly focus on reducing model size (Chen et al. 2015) (Howard et al. 2017) (Kang et al. 2017).

In this paper, we further identify two other main challenges to do ML based data analytics on edge devices. The first one is the lack of suitable models for users. Training a model requires large datasets and rich computing resources, which are often not available to most users. That's one of the reasons that they are bounded to use the models/services provided by large companies. Towards this end we propose the idea "composable services", where the users can construct new services based on existing ones and contribute them to the community. The second new challenge is the deployment of services. Users may lack required knowledge of ML to understand how a model works and how to properly deploy it to any local devices. We present the design of a novel system Zoo to support the construction, compose, and easy deployment of ML models on edge and local devices. In the deployment example, ML services are proved to be easy to compose and deploy with Zoo. Our evaluations show its performance compared with state-of-art deep learning platforms and Google ML services.

Related Work

Social Awareness for Privacy

Ever since Edward Snowden revealed the secret large scale government-level surveillance programmes, social awareness of privacy and personal data is quickly arising. The Internet Trends 2016 report (Meeker 2016) points out that, according to its survey, 45% respondents feel more worried about their online privacy than one year ago, and 74% have limited their online activity because of privacy concerns.

Many online service providers, while collecting large-scale data from users, may be prone to data breaches. Users rely on the promises from big companies to keep their data private and safe (Google 2017b). However, these promises are not always infallible. Yahoo, Tumblr, and Ashley Madison (an online dating service for married people) are only a tiny bit of the current frequent enterprise data breach cases (Cheng, Liu, and Yao 2017). As a result, billions of users' private information is endangered. All too often, these leakages lead to a series of social and ethical aftermaths. As part of the effort to restrain this trend, Regulations such as EU's General Data Protection Regulation (European-Commission 2016) are implemented to "give citizens back control over of their personal data, and to simplify the regulatory environment for business".

Privacy-preserving Analytics on Cloud

Many popular data analytics are conducted on cloud, therefore its data privacy issue has attracted a lot of research interests. Making databases private and secure is one of the solutions. (Babuji et al. 2016) introduces a cloud-based framework that enables secure management and analysis of large, and potentially sensitive, datasets. It claims to ensure secure data storage that can only be accessed by authorised users. (Joshi et al. 2016) has developed an ontology so that big data analytics consumers can write data privacy policies using formal policy languages and build automated systems for compliance validation. (Cuzzocrea, Mastroianni, and

Grasso 2016) presents a short review of existing research efforts along this line.

Hiding data alone is not enough. The model itself can also reveal private information. (Fredrikson, Jha, and Ristenpart 2015) has shown that, with access to the face recognition ML model, the authors can use that to recover recognisable images. (Abadi et al. 2016) develops a method that can prevent these kinds of model inversion attacks against a strong adversary who has full knowledge of the training mechanism and access to the model parameters. Similarly, (Papernot et al. 2016) proposes a "teacher-student" approach that is based on the knowledge aggregation and transfer technique, so as to hide the models trained on sensitive data in a black box.

Instead of hiding data or models, some research suggest user to choose which part of data to upload. (Xu et al. 2017) proposes a mechanism to allow users to clean their data before uploading them to process. It allows for prediction of the desired information, while hiding confidential information that client want to keep private. RAPPOR (Erlingsson, Pihur, and Korolova 2014) enables collecting statistics data from end-user in privacy-preserving crowdsourcing.

However, all of the aforementioned work focus on the traditional cloud side solution. Users' data are still collected to central server for processing, which are prone to issues such as increased service response latency, communication cost, and single point failure.

Data Analytics on Edge Devices

Computing on edge and mobile devices has gained rapid growth. Recently HUAWEI has identified speed and responsiveness of native AI processing on mobile devices as the key to a new era in smartphone innovation (Huawei-News 2017). Applications in this field are of great interest to academia and industry. (Mortier et al. 2016) provides a collection of software components that enable individual data subjects to manage, log and audit access to their data by other parties. (Pan et al. 2017) presents an environment monitoring smartphone app that allow users to takes photos of the outdoor to recognise air quality. Intrusion Detection System (IDS) is important in securing computer networks. Using ML for IDS is especially suitable in stopping attacks that do not have know signatures. (Sankaran and Calix 2016) designs specific hardware for ML based IDS tasks. It achieves high detection accuracy and low energy consumption.

More and more advanced smart services are being pushed to edge devices. Intel's Movidius Neural Compute Stick (Movidius 2017) is a tiny deep learning device that one can use to accelerate AI programming and DNN inference application deployment at the edge. Kvasir (Wang et al. 2016b) is a semantic recommendation system built on top of latent semantic analysis and other state-of-the-art data analytics technologies. Specifically, it seamlessly integrates an automated and proactive content provision service into web browser on users' end.

Many challenges arise when moving ML analytics from cloud to edge devices. One is that compared with resource-rich computing clusters, edge and mobile devices only have quite limited computation power and working memory. To

accommodate heavy ML computation on edge devices, one solution is to train suitable small models to do inference on mobile devices (Chun and Maniatis 2009). This method leads to unsatisfactory accuracy and user experience.

Some techniques are recently proposed to enhance this method. To reduce the memory and disk usage of speech recognition application, (Lei et al. 2013) uses a compressed n-gram language model to do on-the-fly model rescoring. (Chen et al. 2015) presents HashNets, a network architecture that can reduce the redundancy of neural network models to decrease model sizes, while keeping little impact on prediction accuracy. (Howard et al. 2017) from Google reduces model size by a different technique: factorisation of convolution operation.

One of our previous research work (Rodríguez et al. 2018) explores the method of training personalised model on local devices from an initial shared model. Instead of moving data from user to cloud, our method provides for model training and inference in a system where computation is moved to the data. Specifically, we take an initial model learnt from a small set of users and retrain it locally using data from a single user. It is proved to both be robust against adversarial attacks and can improve accuracy.

Besides deploying data analytics purely on cloud or edge, putting computation dynamically on both ends or somewhere in between is also a popular trend. In (Kang et al. 2017), the authors propose partitioning a DNN into two parts, half on edge devices and the other half on cloud, and then develop the *Neurosurgeon* system, aiming at reducing total latency and energy consumption. A key observation is that, in a DNN, output size of each node decreases from front-end to back-end, while the change of computation latency is the opposite. (Wang et al. 2016a) recognises that countless network services reside in ISP networks instead of on the cloud to provide low-latency access. It further investigates computation congestion control strategies to effectively distribute service load within a neighbourhood.

System Design and Implementation

The basic idea of *Composable Services* is users should not have to construct new ML services every time new application requirements arise. In fact, many services can be composed from basic ML services: image recognition, speech-to-text, recommendation, etc. One example is that a service to recognise multiple objects from one image can be composed from two services: image segmentation (which partitions a digital image into multiple segments/regions) and image recognition on each segment. The Zoo system aims at providing user-centric, ML-based services that enables service pulling, composing, sharing, and compatibility checking. The system architecture of Zoo is shown in Figure 1.

We define that a ML service consists of two parts: *development* and *deployment*. *Development* is the design of interaction workflow and the computational functions of different ML services. We provide primitives to construct complex services from simple ones. For example, one of the most important primitives is sequential connection, where the output of one service is used as input of another service.

In Figure 1, a user first designs the target service on computer C (Step ①) in a configuration file, and executes the compose script with Zoo. The square, hexagon, *etc.* here represent different existing ML services. They are connected with the sequential connection primitive in the new structure design. The local repository is first checked to see if models of these required services are cached locally. If not, they will be first pulled from remote repositories. In our current implementation, all published models are stored and accessed from Github Gist (represented by Server A), but this approach could be extended to support pulling services from other peer devices such as machine B in the figure (Step ②).

Deployment deals with service interface and location definition. It is separated from the logic of service construction. Users can move services from being local to remote and vice versa, without changing the structure of constructed service. Deployment is not limited to edge devices (machine E), but can also be on cloud servers (server D), or a hybrid of both cases, to minimise the data revealed to the cloud and the associated communication costs (Step ③). Thus by this design a data analytics service can easily distributed to multiple devices. Finally, user can contribute newly created services to the model repositories so that it can be accessed by others (Step ④).

The Zoo system is implemented on Owl (Wang 2017), an open-source numerical computing system in OCaml language. The reason we choose Owl to support the implementation of Zoo is some of its nice features. Owl provides a full stack support for numerical methods, scientific computing, and advanced data analytics on OCaml. Built on the core data structure of matrix and n-dimensional array, Owl supports a comprehensive set of classic analytics such as math functions, statistics, linear algebra, as well as advanced analytics techniques, namely optimisation, algorithmic differentiation, and regression. On top of them, Owl provides Neural Network and Natural Language Processing modules. Zoo relies on these modules to construct basic models. It has static type checking, and Owl’s ML modules have shown great expressiveness and code flexibility. Moreover, Owl can use the parallel and distributed engine at lower level to support distributed numerical computing and data analytics. It supports different protocols and multiple synchronisation techniques (Wang, Catterall, and Mortier 2017), which is a crucial building block for collaboration between network nodes, and thus important for scaling up data analytics in heterogeneous edge networks.

Deployment Example

Next we show an example to use Zoo to deploy an image classification services based on InceptionV3 neural network architecture. The service is deployed on a representative resource-constrained personal device: a Raspberry Pi 3 Model B (RaspberryPi 2016).

InceptionV3 (Szegedy et al. 2015) is one of Google’s latest effort to do image recognition. It is trained for the ImageNet Large Visual Recognition Challenge (Russakovsky et al. 2015). This is a standard task in computer vision, where models try to classify an image into 1000 classes, like “Ze-

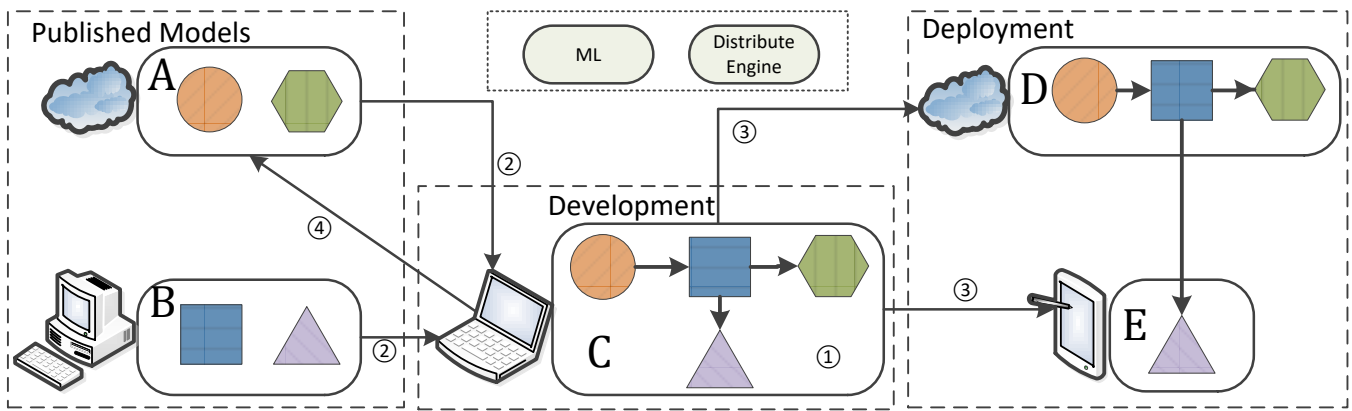


Figure 1: Zoo System Architecture

bra” or “Dishwasher”. Compared with previous DNN models, InceptionV3 has one of the most complex network architectures in computer vision.

The service is composed of two services: an InceptionV3 network that output a vector representing the recognised image class, and an decoding service for ImageNet that translates the previous vector to human-readable format. These two services are sequentially connected. By using Zoo, we can deploy this new services to local Raspberry Pi devices with only one line of command. With minor change of configuration, the service can also be deployed to our server on Digital Ocean, a cloud infrastructure provider. Currently we have used this localised deployment instance for internal automatic image tagging tasks. For the latter instance, we use it to construct an online demo of this service¹.

Performance Evaluation

In this section we present preliminary experiment results with Zoo system. We want to answer two questions here. First, we choose to use Owl to base the Zoo system on. To what degree does its expressiveness has shown? Does it bring any performance trade-off? Second, we deploy services on local devices, so what are their performance compared with popular cloud-based analytic solutions such as Google ML API?

Let’s start with the first question. As stated before, Owl provides extraordinary expressiveness and flexibility. One example to show this is that using Owl to constructs the InceptionV3 network only takes 150 lines of code (LoC), compared with the 400 or more LoC used by Tensorflow. Another one is that, we insert instrumentation code into Owl to enable collecting forward computation latency of each node (or *layer* according to some literature, *et pass.*) in a neural network when doing inference. Adding this feature only takes 50 LoC in the source code of Owl.

Since the Zoo relies on inference using Owl’s Neural Network module, we want to compare the inference time on Owl and the other state-of-art deep learning platforms. We choose three representative DNN models that vary greatly in both architecture complexity and parameter sizes: 1) one




Figure 2: Inference time comparison on different platforms.

small neural network (LeNet-5 (LeCun et al. 1998)) that only consists of 8 nodes and contains about 240KB parameters (each parameter in a model is represented by a 32-bit float number, *et pass.*) for the MNIST handwriting recognition task; 2) a VGG16 (Simonyan and Zisserman 2014) model that has a simple architecture with 38 nodes but a large amount of parameters (500MB) for real-world image recognition task; 3) an InceptionV3 model also for image recognition, with less parameters (100MB), but a far more complex architecture (313 nodes). We compare the time it takes for each model to finish its inference task on different platforms: Owl, TensorFlow, and Caffe2. Each measurement is repeated 20 times.

The results are shown in Figure 2. Regardless of great diversities in these models’ architectures and sizes, Owl takes less time to do inference than Tensorflow and Caffe2. It means that Owl can achieve both expressiveness and good performance. The superior performance of Owl on large models is attributed to its efficient math operations.

Next, we investigate the performance of Zoo system compared with Google ML API. The Google Cloud Vision API (Google 2017a) encapsulates machine learning models in a REST API. It can classify images into thousands of categories as well as detects individual objects and faces within images, and finds and reads printed words contained within images. Its workflow is simple: a user creates service token

¹<http://138.68.155.178/>

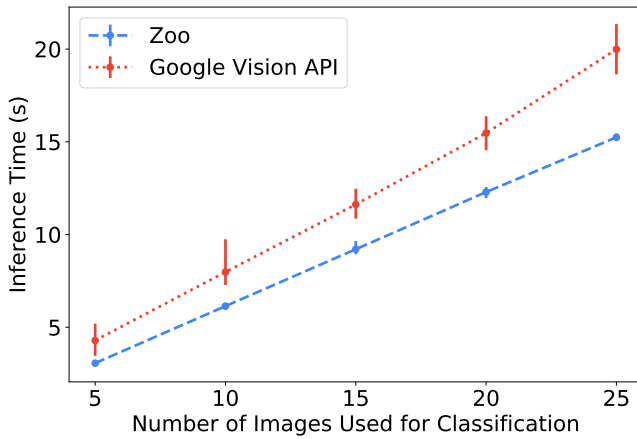


Figure 3: Performance comparison of local image classification services deployed Zoo with Google’s Machine Learning Vision service.

on the Google Cloud Platform (GCP), and passes the token and an image to Google’s server for processing, and then the processed results will be returned to the user as response in the form of JSON. In this evaluation we compare the image classification development.

We deploy Vision API service on GCP. The network connection bandwidth is 34 Mbps measured on `speedtest.net`. In the previous section we have shown a deployment example of our image classification service on local devices with Zoo. For this evaluation, we deploy this service on a Thinkpad T460s laptop with a Intel Core i5-6200U CPU. First, we have collected 100 animal images as a dataset². These images are of different sizes, ranging from 7KB to 1243KB, to better simulate users’ requests in real world. Specifically, we compare the time required for both methods to process different number of images.

The results are shown in Figure 3. It has shown that our Zoo service achieves lower response latency. When the number of input images from a user gradually increases from 5 images to 25 images, the response time of the Zoo service increases linearly, which means the process time of each image basically keeps constant despite the size difference of input images, while the increase speed Google service’s response time seems to grow with more input images. The measurement at each point is repeated for 10 times. Note that the Zoo service response time keep stable (about 0.6s per image) with very small deviation, and thus predictable. The same cannot be said of Google Vision service. It shows relatively large deviation, and with the change of network connection, the response time could easily fluctuate and become unpredictable.

Discussion

In this section we briefly discuss the threat model that is applied in this paper. We assume the companies that provide ML based data analytics services have enough incentives to collect users’ data as much as possible. These data are all

stored in the companies’ servers, and thus face various security and privacy threats. First, the companies can make mistakes and may published or even delete users’ data unintentionally. For example, in February this year, Gitlab lost six hours of database data by accident. Soon after, AWS, one of the most popular cloud computing services provider, also reported an accidental data loss cause by misoperation of an authorized team member. Second, as part of the companies’ assets, these data may of enormous economic value and thus are often targets of hackers. Not all of the attacks can be successfully evaded. This year a plastic surgery was blackmailed by a gang of hackers, who threatened to release patient list and photographs. Large scale data breaches of companies like Yahoo can have detrimental effect on users’ privacy. Last but not the least, some individuals may not be comfortable with the idea that one entity hold their personal information such as sexuality or years of browsing history, even these data are not yet used for any analytics. Moving ML services and data analytics from cloud to the edge devices on users’ side can effectively mitigate these threats.

However, we do not claim this approach can solve all security and privacy threats on user data. One factor we do not consider in this paper is that sometimes personal devices are also vulnerable to hackers. For example, in the WannaCry ransomware attack this year, the hackers encrypted data on individual’s PC and demanded ransom payment. Addressing this threat often requires users’ effort to use strong password, upgrade system frequently, *etc.* Another one is, multiple users may collude to infer users data from the shared model. It is a good topic for further investigation.

Conclusion

Machine Learning based data analytics have been widely adopted in today’s online applications and services. However, with social awareness of privacy and personal data rapidly rising, they also bring about a pressing societal issue: data privacy. Most existing solutions to this issue focus on enhancing the currently popular cloud-based analytics approach, where users’ data need to be collected to central server for processing. Besides data privacy, this approach is also prone to issues such as increased service response latency, communication cost, and single point failure.

In this paper, we argue that to avoid those costs, reduce latency in data processing, and minimise the raw data revealed to service providers, many future AI and ML services could be deployed on users’ devices at the Internet edge rather than putting everything on the cloud. This brings many challenges. Besides the widely discussed factor of resource limitation on edge devices, we identify two other challenges that are not yet recognised in existing literature: lack of suitable models for users, and difficulties in deploying services for users. We present the design of our Zoo system to support the construction, compose, and easy deployment of ML models on edge and local devices. Then we show how services can be composed and deployed with Zoo, and its superior performance compared with existing deep learning platform and cloud-based services. We believe this area of research is only just beginning to gain momentum.

²Available at <https://goo.gl/BJqiBD>

Acknowledgements This work is funded in part by the EPSRC Databox project (EP/N028260/1), NaaS (EP/K031724/2) and Contrive (EP/N028422/1).

References

- [Abadi et al. 2016] Abadi, M.; Chu, A.; Goodfellow, I.; McMahan, H. B.; Mironov, I.; Talwar, K.; and Zhang, L. 2016. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 308–318. ACM.
- [Babuji et al. 2016] Babuji, Y. N.; Chard, K.; Gerow, A.; and Duede, E. 2016. Cloud kotta: Enabling secure and scalable data analytics in the cloud. In *Big Data (Big Data), 2016 IEEE International Conference on*, 302–310. IEEE.
- [Bojarski et al. 2016] Bojarski, M.; Del Testa, D.; Dworakowski, D.; Firner, B.; Flepp, B.; Goyal, P.; Jackel, L. D.; Monfort, M.; Muller, U.; Zhang, J.; et al. 2016. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*.
- [Chen et al. 2015] Chen, W.; Wilson, J.; Tyree, S.; Weinberger, K.; and Chen, Y. 2015. Compressing neural networks with the hashing trick. In *International Conference on Machine Learning*, 2285–2294.
- [Cheng, Liu, and Yao 2017] Cheng, L.; Liu, F.; and Yao, D. D. 2017. Enterprise data breach: causes, challenges, prevention, and future directions. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 7(5).
- [Chun and Maniatis 2009] Chun, B.-G., and Maniatis, P. 2009. Augmented smartphone applications through clone cloud execution. In *HotOS*, volume 9, 8–11.
- [Cuzzocrea, Mastroianni, and Grasso 2016] Cuzzocrea, A.; Mastroianni, C.; and Grasso, G. M. 2016. Private databases on the cloud: Models, issues and research perspectives. In *Big Data (Big Data), 2016 IEEE International Conference on*, 3656–3661. IEEE.
- [Erlingsson, Pihur, and Korolova 2014] Erlingsson, Ú.; Pihur, V.; and Korolova, A. 2014. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*, 1054–1067. ACM.
- [European-Commission 2016] European-Commission. 2016. Protection of personal data. http://ec.europa.eu/justice/data-protection/reform/index_en.htm. Accessed: 2017-10-07.
- [Fredrikson, Jha, and Ristenpart 2015] Fredrikson, M.; Jha, S.; and Ristenpart, T. 2015. Model inversion attacks that exploit confidence information and basic countermeasures. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, 1322–1333. ACM.
- [Google 2017a] Google. 2017a. Google cloud vision api. <https://cloud.google.com/vision>. Accessed Nov. 11, 2017.
- [Google 2017b] Google. 2017b. Google privacy. <https://privacy.google.com>. Accessed Nov. 11, 2017.
- [Howard et al. 2017] Howard, A. G.; Zhu, M.; Chen, B.; Kalenichenko, D.; Wang, W.; Weyand, T.; Andreetto, M.; and Adam, H. 2017. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *CoRR* abs/1704.04861.
- [Huawei-News 2017] Huawei-News. 2017. HUAWEI Reveals the Future of Mobile AI at IFA 2017. <http://consumer.huawei.com/en/press/news/2017/ifa2017-kirin970/>. [Online; accessed 10-Nov-2017].
- [Joshi et al. 2016] Joshi, K. P.; Gupta, A.; Mittal, S.; Pearce, C.; Joshi, A.; and Finin, T. 2016. Semantic approach to automating management of big data privacy policies. In *Big Data (Big Data), 2016 IEEE International Conference on*, 482–491. IEEE.
- [Kang et al. 2017] Kang, Y.; Hauswald, J.; Gao, C.; Rovinski, A.; Mudge, T.; Mars, J.; and Tang, L. 2017. Neurosurgeon: Collaborative intelligence between the cloud and mobile edge. In *Proceedings of the Twenty-Second International Conference on Architectural Support for Programming Languages and Operating Systems*, 615–629. ACM.
- [LeCun et al. 1998] LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86(11):2278–2324.
- [Lei et al. 2013] Lei, X.; Senior, A. W.; Gruenstein, A.; and Sorensen, J. 2013. Accurate and compact large vocabulary speech recognition on mobile devices. In *Interspeech*, volume 1.
- [Meeker 2016] Meeker, M. 2016. Internet trends 2016 report. <http://www.kpcb.com/internet-trends>. Accessed: 2017-10-01.
- [Mortier et al. 2016] Mortier, R.; Zhao, J.; Crowcroft, J.; Wang, L.; Li, Q.; Haddadi, H.; Amar, Y.; Crabtree, A.; Colley, J.; Lodge, T.; et al. 2016. Personal data management with the databox: What’s inside the box? In *Proceedings of the 2016 ACM Workshop on Cloud-Assisted Networking*, 49–54. ACM.
- [Movidius 2017] Movidius. 2017. Movidius neural compute stick. <https://developer.movidius.com/>. Accessed Nov. 11, 2017.
- [Pan et al. 2017] Pan, Z.; Yu, H.; Miao, C.; and Leung, C. 2017. Crowdsensing air quality with camera-enabled mobile devices. In *AAAI*, 4728–4733.
- [Papernot et al. 2016] Papernot, N.; Abadi, M.; Erlingsson, Ú.; Goodfellow, I.; and Talwar, K. 2016. Semi-supervised knowledge transfer for deep learning from private training data. *arXiv preprint arXiv:1610.05755*.
- [RaspberryPi 2016] RaspberryPi. 2016. Raspberry pi. <https://www.raspberrypi.org/products/raspberry-pi-3-model-b/>. Accessed February 15, 2017.
- [Rodríguez et al. 2018] Rodríguez, S. S.; Wang, L.; Zhao, J. R.; Mortier, R.; and Haddadi, H. 2018. Privacy-preserving personal model training. *Internet-of-Things Design and Implementation (IoTDI), The 3rd ACM/IEEE International Conference on*.

- [Russakovsky et al. 2015] Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; Berg, A. C.; and Fei-Fei, L. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)* 115(3):211–252.
- [Sankaran and Calix 2016] Sankaran, R., and Calix, R. A. 2016. On the feasibility of an embedded machine learning processor for intrusion detection. In *Big Data (Big Data), 2016 IEEE International Conference on*, 1082–1089. IEEE.
- [Schuster 2010] Schuster, M. 2010. Speech recognition for mobile devices at google. In *Pacific Rim International Conference on Artificial Intelligence*, 8–10. Springer.
- [Simonyan and Zisserman 2014] Simonyan, K., and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *CoRR* abs/1409.1556.
- [Szegedy et al. 2015] Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; and Wojna, Z. 2015. Rethinking the inception architecture for computer vision. *CoRR* abs/1512.00567.
- [Wang et al. 2016a] Wang, L.; Almeida, M.; Blackburn, J.; and Crowcroft, J. 2016a. C3po: Computation congestion control (proactive). In *ICN*, 231–236.
- [Wang et al. 2016b] Wang, L.; Tasoulis, S.; Roos, T.; and Kangasharju, J. 2016b. Kvasir: Scalable provision of semantically relevant web content on big data framework. *IEEE Transactions on Big Data* 2(3):219–233.
- [Wang, Catterall, and Mortier 2017] Wang, L.; Catterall, B.; and Mortier, R. 2017. Probabilistic Synchronous Parallel. *ArXiv e-prints*.
- [Wang 2017] Wang, L. 2017. Owl: A general-purpose numerical library in ocaml. *CoRR* abs/1707.09616.
- [Xu et al. 2017] Xu, K.; Cao, T.; Shah, S.; Maung, C.; and Schweitzer, H. 2017. Cleaning the null space: A privacy mechanism for predictors. In *AAAI*, 2789–2795.