Data Analytics Service Composition and Deployment on Edge Devices

Jianxin Zhao, Tudor Tiplea, Richard Mortier, Jon Crowcroft, Liang Wang
Computer Laboratory, University of Cambridge, UK
firstname.lastname@cl.cam.ac.uk

ABSTRACT

Data analytics on edge devices has gained rapid growth in research, industry, and different aspects of our daily life. This topic still faces many challenges such as limited computation resource on edge devices. In this paper, we further identify two main challenges: the composition and deployment of data analytics services on edge devices. We present the Zoo system to address these two challenges: on one hand, it provides simple and concise domain-specific language to enable easy and type-safe composition of different data analytics services; on the other, it utilises multiple deployment backends, including Docker container, JavaScript, and MirageOS, to accommodate the heterogeneous edge deployment environment. We show the expressiveness of Zoo with a use case, and thoroughly compare the performance of different deployment backends in evaluation.

CCS CONCEPTS

• **Software and its engineering** → **Software notations and tools**; • **Human-centered computing** → *Ubiquitous and mobile computing*; • **Computing methodologies** → *Machine learning*;

1 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 and Netflix. 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, self-driving cars [3], image processing [7], *etc.* Many popular data analytics are deployed on cloud computing infrastructures. However, they require aggregating users' data at central server for processing. This architecture is prone to issues such as increased service response latency, communication cost, single point failure, and data privacy concerns.

Recently computation on edge and mobile devices has gained rapid growth, such as personal data analytics in home [16], DNN application on a tiny stick [17], and semantic search and recommendation on web browser [24]. Edge computing is also boosting content distribution by supporting peering and caching [25].

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 the author(s) 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.

Big-DAMA'18, August 20, 2018, Budapest, Hungary

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-5904-7/18/08.

<https://doi.org/10.1145/3229607.3229615>

HUAWEI has identified speed and responsiveness of native AI processing on mobile devices as the key to a new era in smartphone innovation [9].

Many challenges arise when moving ML analytics from cloud to edge devices. One widely discussed challenge is the limited computation power and working memory of edge devices. Personalising analytics models on different edge devices is also a very interesting topic [19]. However, one problem is not yet well defined and investigated: the deployment of data analytics services. Most existing machine learning frameworks such as TensorFlow and Caffe focus mainly on the training of analytics models. On the other, the end users, many of whom are not ML professionals, mainly use trained models to perform inference. This gap between the current ML systems and users' requirements is growing.

Another challenge in conducting ML based data analytics on edge devices is model composition. Training a model often requires large datasets and rich computing resources, which are often not available to normal users. That's one of the reasons that they are bounded with the models and services provided by large companies. To this end we propose the idea *Composable Service*. Its basic idea is that many services can be constructed from basic ML ones such as image recognition, speech-to-text, and recommendation to meet new application requirements. We believe that modularity and composition will be the key to increasing usage of ML-based data analytics.

This paper tries to address these two challenges. Specifically, the contribution of this paper includes:

- We identify two challenges that are not yet well explored in the literature about data analytics on edge devices: service composition and deployment.
- We present the design of the Zoo system to address the previous two challenges. It provides concise Domain-specific Language (DSL) to enable composition of different data analytics services, and also deploys services to multiple backends.
- We present a use case to demonstrate the expressiveness of the DSL, and thoroughly evaluate different deployment backend for analytics services.

2 SYSTEM DESIGN

The Zoo system is implemented on Owl [23], an open-source scientific computing library 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 N-dimensional array (ndarray), 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

Figure 1: Zoo System Architecture

regression. On top of them, Owl provides Neural Network and Natural Language Processing modules. Zoo relies on these modules to construct basic ML services. OCaml provides static type checking, and Owl’s ML modules have shown great expressiveness and code flexibility.

Initially, the Zoo system is designed to make it convenient for developers to share their OCaml code snippets. The design principle is to make the whole ecosystem open, flexible, and extensible. One typical scenario for using the basic functions of Zoo can be described as follows. Developer A creates a script, uploads it to Gist, and then share it using a string of Gist id. When developer B gets this id, he can use the functions from A’s scripts by simply using the “#zoo” directive in his code. All the OCaml files in the Gist will be imported as modules for B to use.

Based on these basic functionalities, we extend the Zoo system to address the composition and deployment challenges. First, we would like to briefly introduce the workflow of Zoo as shown in Fig. 1. The workflow consists of two parts: *development* on the left side and *deployment* on the right.

Development concerns the design of interaction workflow and the computational functions of different services. One basic component is Github Gist. A normal Gist script will be loaded as a module in OCaml. To compose functionalities from different Gists only requires a developer to add one configuration file to each Gist. This file is in JSON format. It consists of one or more name-value pairs. Each pair is a signature for a function the script developer wants to expose as a service. These Gists then can be imported and composed to make new services. When a user is satisfied with the composing result, she can save the new service as another Zoo Gist.

Deployment takes a Gist and creates models in different backends. These models can be published and deployed to edge devices. It is separated from the logic of development. Basic services and composed ones are treated equally. Besides, users can move services from being local to remote and vice versa, without changing the structure of the constructed service. Deployment is not limited to edge devices, but can also be on cloud servers, or a hybrid of both cases, to minimise the data revealed to the cloud and the associated communication costs. Thus by this design a data analytics service can easily be distributed to multiple devices.

In the rest of this section, we’ll elaborate the design and details of different parts of this workflow.

2.1 Service

Gist is a core abstraction in Zoo. It is the centre of code sharing. However, to compose multiple analytics snippets, Gist alone is insufficient. For example, it cannot express the structure of how different pieces of code are composed together. Therefore, we introduce another abstraction: *service*.

A service consists of three parts: *Gists*, *types*, and *dependency graph*. *Gists* is the list of Gist ids this service requires. *Types* is the parameter types of this service. Any service has zero or more input parameters and one output. This design follows that of an OCaml function. *Dependency graph* is a graph structure that contains information about how the service is composed. Each node in it represents a function from a Gist, and contains the Gist’s name, id, and number of parameters of this function.

Zoo provides three core operations about a service: create, compose, and publish. The *create_service* creates a dictionary of services given a Gist id. This operation reads the service configuration file from that Gist, and creates a service for each function specified in the configuration file. The *compose_service* provides a series of operations to combine multiple services into a new service. A compose operation does type checking by comparing the “types” field of two services. An error will be raised if incompatible services are composed. A composed service can be saved to a new Gist or be used for further composition. The *publish_service* makes a service’s code into such forms that can be readily used by end users. Zoo is designed to support multiple backends for these publication forms. Currently it targets Docker container, JavaScript, and MirageOS [15] as backends.

2.2 Type Checking

As mentioned in Section 2.1, one of the most important tasks of service composition is to make sure the type matches. For example, suppose there is an image analytics service that takes a PNG format image, and if we connect to it another one that produces a JPEG image, the resulting service will only generate meaningless output for data type mismatch. OCaml provides primary types such as integer, float, string, and bool. The core data structure of Owl is ndarray (or tensor as it is called in some other data analytics frameworks). However, all these types are insufficient for high level service type checking as mentioned. That motives us to derive richer high-level types.

To support it, we use generalised algebraic data types (GADTs) in OCaml. There already exist several model collections on different platforms, e.g. Caffe [22] and MxNet [18]. We observe that most current popular deep learning (DL) models can generally be categorised into three fundamental types: image, text, and voice. Based on them, we define sub-types for each: PNG and JPEG image, French and English text and voice, i.e. `png img`, `jpeg img`, `fr text`, `en text`, `fr voice`, and `en voice` types. More can be further added easily in Zoo. Therefore type checking in OCaml ensures type-safe and meaningful composition of high level services.

2.3 Backend

Recognising the heterogeneity of edge device deployment, one key principle of Zoo is to support multiple deployment methods. Containerisation as a lightweight virtualisation technology has gained enormous traction. It is used in deployment systems such as Kubernetes. Zoo supports deploying services as Docker containers. Each container provides RESTful API for end users to query.

Another backend is JavaScript. Using JavaScript to do analytics aside from front end development begins to attract interests from academia [24] and industry, such as Tensorflow.js and Facebook's Reason language. By exporting OCaml and Owl functions to JavaScript code, users can do complex data analytics on web browser directly without relying on any other dependencies.

Aside from these two backends, we also initially explore using MirageOS as an option. Mirage is an example of Unikernel, which builds tiny virtual machines with a specialised minimal OS that host only one target application. Deploying to Unikernel is proved to be of low memory footprint, and thus quite suitable for resource-limited edge devices.

2.4 DSL

Zoo provides a minimal DSL for service composition and deployment.

Composition. To acquire services from a Gist of id gid , we use $\$gid$ to create a dictionary, which maps from service name strings to services. We implement the dictionary data structure using `Hashtbl` in OCaml. The `#` operator is overloaded to represent the "get item" operation. Therefore,

$$\$gid\#sname$$

can be used to get a service that is named "sname". Now suppose we have n services: $f_1, f_2, \dots, f_n$. Their outputs are of type $t_{f_1}, t_{f_2}, \dots, t_{f_n}$. Each service s accepts m_s input parameters, which have type $t_s^1, t_s^2, \dots, t_s^{m_s}$. Also, there is a service g that takes n inputs, each of them has type $t_g^1, t_g^2, \dots, t_g^n$. Its output type is t_o . Here Zoo provides the $\$>$ operator to compose a list of services with another:

$$[f_1, f_2, \dots, f_n]\$>g$$

This operation returns a new service that has $\sum_{s=1}^n m_s$ inputs, and is of output type t_o . This operation does type checking to make sure that $t_{f_i} = t_g^i, \forall i \in 1, 2, \dots, n$.

Deployment. Taking a service s , be it a basic or composed one, it can be deployed using the following syntax:

$$s\$@ \text{backend}$$

The $\$@$ operator publish services to certain backend. It returns a string of URI of the resources to be deployed.

Note that the $\$>$ operator leads to a tree-structure, which is in most cases sufficient for our real-world service deployment. However, a more general operation is to support graph structure. This will be our next-step work.

2.5 Service Discovery

The services require a service discovery mechanism. For simplicity's sake, each newly published service is added to a public record hosted on a server. The record is a list of items, and each item contains

the Gist id that service based on, a one-line description of this service, string representation of the input types and output type of this service, e.g. "image -> int -> string -> tex", and service URI. For the container deployment, the URI is a DockerHub link, and for JavaScript backend, the URI is a URL link to the JavaScript file itself. The service discovery mechanism is implemented using off-the-shelf database.

2.6 Version Control

Developers would modify and upload their scripts several times. As such, each version of a script is assigned a unique id in Gist. Zoo supports specifying a version of a Gist.

The naming scheme of a Gist is `gid/[vid|latest]/pin`. A user can either choose a specific version id, or he can use the latest version, which means the newest version on local cache. Obviously, "latest" will introduce cache inconsistency. The latest version on one machine might not be the same on the other. To get the up-to-date version from Gist server, the download time of the latest version on a local machine will be saved as metadata. The newest version on server will be pulled to local cache after a certain period of time, if "latest" flag is set in the Gist name. Ideally, every published service should contain a specific version id, and "latest" should only be used during development.

Zoo can analyse dependency information of a Gist and save it. When the "pin" flag is set, Gist dependency graph of current script will be saved or loaded.

3 USE CASE

To illustrate the workflow above, let's consider a synthetic scenario. Alice is a French data analyst. She knows how to use ML and DL models in existing platforms, but is not an expert. Her recent work is about testing the performance of different image classification neural networks. To do that, she needs to first modify the image using the DNN-based Neural Style Transfer (NST) algorithm. The NST algorithm takes two images and outputs to a new image, which is similar to the first image in content and the second in style. This new image should be passed to an image classification DNN for inference. Finally, the classification result should be translated to French. She does not want to put academic-related information on Google's server, but she cannot find any single pre-trained model that performs this series of tasks.

Here comes the Zoo system to help. Alice find Gists that can do image recognition, NST, and translation separately. Even better, she can perform image segmentation to greatly improve the performance of NST [14] using another Gist. All she has to provide is some simple code to generate the style images she need to use. She can then assemble these parts together easily using Zoo.

```
open Zoo
(* Image classification *)
let s_img = $ "aa36e" # "infer";;
(* Image segmentation *)
let s_seg = $ "d79e9" # "seg";;
(* Neural style transfer *)
let s_nst = $ "6f28d" # "run";;
(* Translation from English to French *)
```

```

let s_trans = $ "7f32a" # "trans";;
(* Alice's own style image generation service *)
let s_style = $ alice_Gist_id # "image_gen";;

(* Compose services *)
let s = [s_seg; s_style] $> s_nst
    $> n_img $> n_trans;;
(* Publish to a new Docker Image *)
let pub = (List.hd s) $@
    (CONTAINER "alice/image_service:latest");;

```

Note that the Gist id used in the code is shorted from 32 digits to 5 due to column length limit. Once Alice creates the new service and published it as a container, she can then run it locally and send request with image data to the deployed machine, and get image classification results back in French.

4 EVALUATION

In the evaluation section we mainly compare the performance of different backends we use. Specifically, we observe three representative groups of operations: (1) map and fold operations on ndarray; (2) using gradient descent, a common numerical computing subroutine, to get *argmin* of a certain function; (3) conducting inference on complex DNNs, including SqueezeNet [10] and a VGG-like convolution network. The evaluations are conducted on a ThinkPad T460S laptop with Ubuntu 16.04 operating system. It has an Intel Core i5-6200U CPU and 12GB RAM.

The OCaml compiler can produce two kinds of executables: bytecode and native. Native executables are compiled specifically for an architecture and are generally faster, while bytecode executables have the advantage of being portable. A Docker container can adopt both options.

For JavaScript though, since the Owl library contains functions that are implemented in C, it cannot be directly supported by `js-of-ocaml`, the tool we use to convert OCaml code into JavaScript. Therefore in the Owl library, we have implemented a “base” library in pure OCaml that shares the core functions of the Owl library. Note that for convenience we refer to the pure implementation of OCaml and the mix implementation of OCaml and C as `base-lib` and `owl-lib` separately, but they are in fact all included in the Owl library. For Mirage compilation, we use both libraries.

Fig. 2(a-b) show the performance of map and fold operations on ndarray. We use simple functions such as plus and multiplication on 1-d (size < 1,000) and 2-d arrays. The log-log relationship between total size of ndarray and the time each operation takes keeps linear. For both operations, `owl-lib` is faster than `base-lib`, and native executables outperform bytecode ones. The performance of Mirage executives is close to that of native code. Generally JavaScript runs the slowest, but note how the performance gap between JavaScript and the others converges when the ndarray size grows. For fold operation, JavaScript even runs faster than bytecode when size is sufficiently large.

In Fig. 3, we want to investigate if the above observations still hold in more complex numerical computation. We choose to use a Gradient Descent algorithm to find the value that locally minimise a function. We choose the initial value randomly between [0, 10].

Table 1: Inference Speed of Deep Neural Networks

Time (ms)	VGG	SqueezeNet
owl-native	7.96 (± 0.93)	196.26(± 1.12)
owl-byte	9.87 (± 0.74)	218.99(± 9.05)
base-native	792.56(± 19.95)	14470.97 (± 368.03)
base-byte	2783.33(± 76.08)	50294.93 (± 1315.28)
mirage-owl	8.09(± 0.08)	190.26(± 0.89)
mirage-base	743.18 (± 13.29)	13478.53 (± 13.29)
JavaScript	4325.50(± 447.22)	65545.75 (± 629.10)

Table 2: Size of executables generated by backends

Size (KB)	native	bytecode	Mirage	JavaScript
base	2,437	4,298	4,602	739
native	14,875	13,102	16,987	-

For both $\sin(x)$ and $x^3 - 2x^2 + 2$, we can see that JavaScript runs the slowest, but this time the `base-lib` slightly outperforms `owl-lib`.

We further compare the performance of DNN, which requires large amount of computation. We compare SqueezeNet and a VGG-like convolution network. They have different sizes of weight and networks structure complexities. Table. 1 shows that, though the performance difference between `owl-lib` and `base-lib` is not obvious, the former is much better. So is the difference between native and bytecode for `base-lib`. JavaScript is still the slowest. The core computation required for DNN inference is the convolution operation. Its implementation efficiency is the key to these differences. Current we are working on improving its implementation in `base-lib`.

We have also conducted the same evaluation experiments on RaspberryPi 3 Model B. Fig. 2(c) shows the performance of fold operation on ndarray. Besides the fact that all backends runs about one order of magnitude slower than that on the laptop, previous observations still hold. This figure also implies that, on resource-limited devices such as RaspberryPi, the key difference is between native code and bytecode, instead of `owl-lib` and `base-lib` for this operation. The other figures are not presented here due to space limited, but the conclusions are similar.

Finally, we also briefly compare the size of executables generated by different backends. We take the SqueezeNet for example, and the results are shown in Table 2. It can be seen that `owl-lib` executives have larger size compared to `base-lib` ones, and JavaScript code has the smallest file size.

It can be seen that there does not exist a dominant method of deployment for all these backends. It is thus imperative to choose suitable backend according to deployment environment.

Besides, we want to compare the performance of the deployed services. Specifically, we compare deployed DNN services in Owl and the other state-of-art deep learning platforms, namely TensorFlow and Caffe2. We choose two representative DNN models that vary greatly in both architecture complexity and parameter sizes: 1) one small neural network LeNet-5 [12] that consists of 8 nodes and about 240KB parameters (each parameter in a model is represented by a 32-bit float number, *et pass.*) for the MNIST

Figure 2: Performance of map and fold operations on ndarray on laptop (a-b) and RaspberryPi (c).

Figure 3: Performance of gradient descent on function f to find $\arg\min(f)$

Table 3: Compare inference time of deployed DNN services

Time (ms)	Zoo/Owl	Tensorflow	Caffe2
LeNet	0.68 ± 0.19	2.00 ± 0.23	2.90 ± 0.41
InceptionV3	321.54 ± 22.74	315.78 ± 2.78	439.46 ± 14.30

handwriting recognition task; 2) an InceptionV3 [21] model for the image recognition task, with more parameters (100MB) and a far more complex architecture (313 nodes). We compare the time it takes for each deployed Owl service to finish its inference task with other platforms. Each measurement is repeated 20 times. As can be seen in Table 3, regardless of great diversities in the models' architectures and sizes, the Zoo/Owl services achieve similar or better performance than Tensorflow and Caffe2.

5 RELATED WORK

Moving ML analytics from cloud to edge devices faces many challenges. One widely recognised challenge 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 [5].

This method leads to unsatisfactory accuracy and user experience. Some techniques [4, 8, 13] are proposed to enhance this method.

To reduce the memory and disk usage of speech recognition application, [13] uses a compressed n-gram language model to do on-the-fly model rescoring. [4] 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. [8] from Google reduces model size by a different technique: factorisation of convolution operation.

Another challenge is to personalise analytics models. One of our previous research work [19] explores 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.

There exist several work on deployment of data analytics services. Clipper [6] is a general-purpose low-latency prediction serving system. It provides end users with a series of ML applications including computer vision, speech recognition, recommendation, etc. Clipper tries to maximise accuracy and throughput given certain latency budget. However, the service or model deployment here is only limited to server-side, and the users cannot deploy their own service freely. TensorFlow Serving [20] tries to simplify the deployment models that are created and trained by TensorFlow. It is similar to Zoo in its mechanism of serving a model for request from users. However, it does not support type-safe service composing, nor does it offer flexible cross platform automatic deployment solutions using multiple backends. Some deployment systems are limited to certain applications, such as Linear Regression model in LASER [1] system, and video analytics model in NoScope [11]. Serverless Architectures such as AWS Lambda [2] allow users to deploy functions cost-efficiently. Existing serverless frameworks all bound closely with cloud computing platforms such as Amazon Web Services and Google Cloud Platform.

6 CONCLUSIONS AND FUTURE WORK

In this work we identify two challenges of conducting data analytics on edge: service composition and deployment. We propose the Zoo system to address these two challenges. For the first one, it provides a simple DSL to enable easy and type-safe composition of different advanced services. We present a use case to show the expressiveness of the code. For the second, to accommodate the heterogeneous edge deployment environment, we utilise multiple backends, including Docker container, JavaScript, and MirageOS. We thoroughly evaluate the performance of different backends using three representative groups of numerical operations as workload. The results show that no single deployment backend is preferable to the others, so deploying data analytics services requires choosing suitable backend according to the deployment environment.

Currently the Zoo system is still in development phase. A stable version of Zoo is included in the source code of Owl, and we plan to open source the development version soon. Many opportunities exist to improve the implementation. As previously discussed, to enable a Gist script to be deployed as service, a user needs to specify the function names and signatures in the configuration. However, it would be more usable if the signature can be inferred automatically. Another optimisation is in the serialisation method. Currently users' data are encoded as Base64 string. However, this method is known to increase the data size in transmission across the internet. Other serialisation methods such as Protocol Buffers can be leveraged to solve this issue. Moreover, as discussed in Section 5, there exist several techniques to optimise the DNN applications. We are trying to incorporate them into Zoo.

ACKNOWLEDGMENTS

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

REFERENCES

- [1] Deepak Agarwal, Bo Long, Jonathan Traupman, Doris Xin, and Liang Zhang. 2014. Laser: A scalable response prediction platform for online advertising. In *Proc. ACM WSDM'14*. ACM, 173–182.
- [2] Amazon. 2017. AWS Lambda. <https://aws.amazon.com/cn/lambda/>. (2017). [Online; accessed 14-Mar-2018].
- [3] Mariusz Bojarski, Del Testa, et al. 2016. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316* (2016).
- [4] Wenlin Chen, James Wilson, Stephen Tyree, Kilian Weinberger, and Yixin Chen. 2015. Compressing neural networks with the hashing trick. In *International Conference on Machine Learning*. 2285–2294.
- [5] Byung-Gon Chun and Petros Maniatis. 2009. Augmented smartphone applications through clone cloud execution.. In *HotOS*, Vol. 9. 8–11.
- [6] Daniel Crankshaw, Xin Wang, Guilio Zhou, Michael J Franklin, Joseph E Gonzalez, and Ion Stoica. 2017. Clipper: A Low-Latency Online Prediction Serving System.. In *NSDI*. 613–627.
- [7] Google. 2017. Google Cloud Vision API. <https://cloud.google.com/vision>. (2017). Accessed Nov. 11, 2017.
- [8] Andrew G. Howard, Menglong Zhu, Bo Chen, et al. 2017. MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. *CoRR* abs/1704.04861 (2017). arXiv:1704.04861
- [9] Huawei-News. 2017. HUAWEI Reveals the Future of Mobile AI at IFA 2017. <http://consumer.huawei.com/en/press/news/2017/ifa2017-kirin970/>. (2017). [Online; accessed 10-Nov-2017].
- [10] Forrest N Iandola, Song Han, Matthew W Moskewicz, Khalid Ashraf, William J Dally, and Kurt Keutzer. 2016. SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and < 0.5 MB model size. *arXiv preprint arXiv:1602.07360* (2016).
- [11] Daniel Kang, John Emmons, Firas Abuzaied, Peter Bailis, and Matei Zaharia. 2017. NoScope: optimizing neural network queries over video at scale. *Proceedings of the VLDB Endowment* 10, 11 (2017), 1586–1597.
- [12] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (1998), 2278–2324.
- [13] Xin Lei, Andrew W Senior, Alexander Gruenstein, and Jeffrey Sorensen. 2013. Accurate and compact large vocabulary speech recognition on mobile devices.. In *Interspeech*, Vol. 1.
- [14] Fujun Luan, Sylvain Paris, Eli Shechtman, and Kavita Bala. 2017. Deep Photo Style Transfer. *CoRR* abs/1703.07511 (2017). arXiv:1703.07511 <http://arxiv.org/abs/1703.07511>
- [15] Anil Madhavapeddy and David J Scott. 2013. Unikernels: Rise of the virtual library operating system. *Queue* 11, 11 (2013), 30.
- [16] Richard Mortier, Jianxin Zhao, Jon Crowcroft, Liang Wang, Qi Li, Hamed Haddadi, Yousef Amar, 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*. ACM, 49–54.
- [17] Movidius. 2017. Movidius Neural Compute Stick. <https://developer.movidius.com/>. (2017). Accessed Nov. 11, 2017.
- [18] MxNet. 2017. MxNet Model Zoo. https://mxnet.incubator.apache.org/model_zoo/index.html. (2017). [Online; accessed 14-Mar-2018].
- [19] Sandra Servia Rodriguez, Liang Wang, Jianxin R. Zhao, Richard Mortier, and Hamed Haddadi. 2018. Privacy-preserving Personal Model Training. *Internet-of-Things Design and Implementation (IoTDI), The 3rd ACM/IEEE International Conference on* (2018).
- [20] TensorFlow Serving. 2017. TensorFlow Serving. <https://www.tensorflow.org/serving/>. (2017). [Online; accessed 14-Mar-2018].
- [21] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2818–2826.
- [22] Berkeley Vision and Learning Center. 2017. Caffe Model Zoo. <https://github.com/BVLC/caffe/wiki/Model-Zoo>. (2017). [Online; accessed 14-Mar-2018].
- [23] Liang Wang. 2017. Owl: A General-Purpose Numerical Library in OCaml. *CoRR* abs/1707.09616 (2017). <http://arxiv.org/abs/1707.09616>
- [24] Liang Wang, Sotiris Tasoulis, Teemu Roos, and Jussi Kangasharju. 2016. Kvasir: Scalable provision of semantically relevant web content on big data framework. *IEEE Transactions on Big Data* 2, 3 (2016), 219–233.
- [25] Kok-Kiong Yap, Murtaza Motiwala, Jeremy Rahe, Steve Padgett, Matthew Holli-man, Gary Baldus, Marcus Hines, Taeun Kim, Ashok Narayanan, Ankur Jain, et al. 2017. Taking the edge off with espresso: Scale, reliability and programmability for global internet peering. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*. ACM, 432–445.