
Quantification in-the-wild: data-sets and baselines

Oscar Beijbom

Judy Hoffman

Evan Yao

Trevor Darrell

University of California, Berkeley, California

{obeijbom, jhoffman, trevor}@eecs.berkeley.edu {evanyao}@berkeley.edu

Alberto Rodriguez - Ramirez

Manuel Gonzalez - Rivero

Ove Hoegh - Guldberg

Global Change Institute, The University of Queensland, Australia

{alberto.rodriguez, m.gonzalezrivero, oveh}@uq.edu.au

Abstract

Quantification is the task of estimating the class-distribution of a data-set. While typically considered as a parameter estimation problem with strict assumptions on the data-set shift, we consider quantification in-the-wild, on two large scale data-sets from marine ecology: a survey of Caribbean coral reefs, and a plankton time series from Martha’s Vineyard Coastal Observatory. We investigate several quantification methods from the literature and indicate opportunities for future work. In particular, we show that a deep neural network can be fine-tuned on a very limited amount of data (25 - 100 samples) to outperform alternative methods.

1 Introduction

The term ‘quantification’ was introduced by Forman [3] and defined as the task of estimating the class-distribution of a categorical data-set. In a typical scenario, a classifier is trained on a set of labeled ‘source’ samples, and applied to a new set of ‘target’ samples. This is challenging in the presence of ‘data-set shifts’ [9] where the underlying probability function of the source data differs from that of the target data. Contrary to classification, which is studied both in the presence and absence of domain shift, quantification is only of interest under data-set shift, since otherwise, the class-distribution can be estimated directly from the labeled source data.

Quantification is important for two principled reasons. First the class-distribution itself, rather than the full labeling, is the desired end-product in many applications. This occurs for example in sampling surveys, where repeated quantification is required to assess spatial or temporal patterns [3, 18]. Second, in applications where a full labeling is required, the class-distribution of the target data can be used to re-calibrate a classifier trained on the source data [14, 16].

Quantification has been studied under the class-distribution shift assumption [16, 3, 2], meaning that the source and target class conditional distributions are the same [9]. As we shall see, this is a rather strong assumption and may not hold in practice. In another line of work, domain adaptation (DA) has been studied under the assumption that the source and target class-conditional distributions are ‘similar’ [7]. In this work, the class-distribution shift is often controlled for by experimenting on class-balanced data-sets [15].

We introduce two large-scale data-sets from marine ecology, where quantification arise naturally and where automation is imperative for ecological analysis. We evaluate efficacy of several quantification methods from the literature.

1.1 Problem statement

Let $x \in \mathbb{R}^d$ be d-dimensional input samples, and $y \in 1, \dots, c$ class labels. We assume that a large number of samples and classes, $\{(x_i, y_i)\}_{i=1}^n$, are available in a source domain defined by some joint probability function $p(x, y)$. We also assume that a large number of unlabeled samples, $\{x_i\}_{i=1}^{n'}$, are available in a target domain, defined by different probability function $p'(x, y) \neq p(x, y)$. The general goal of quantification is to estimate the probability distribution over classes in the target

domain: $p'(y) \equiv q \in \mathbb{R}^c$. However, the problem statement as defined above is intractable if (a) the domain shift is arbitrary and (b) there are no labeled samples in the target domain. It is therefore typically studied under one of two relaxations:

Definition 1.1. Unsupervised Quantification: In unsupervised quantification, the data-set shift is assumed to be a pure class-distribution shift, i.e. $p(y) \neq p'(y)$, but $p(x|y) = p'(x|y)$. Alternatively, the data-set shift is assumed to be ‘small’, and the unlabeled set of target samples, $\{x_i\}_{i=1}^{n'}$, is used to align the internal feature representation of a machine learning algorithm.

Definition 1.2. Supervised Quantification: In supervised quantification, no explicit assumptions are made on the data-set shift, but it is assumed that a small amount of samples are available in the target domain, $\{(x_i, y_i)\}_{i=1}^b \in p'(x, y)$.

For supervised quantification, we only consider methods where the labeled target samples are selected randomly, leaving the design of active sampling methods to future work.

1.2 Related work

For unsupervised quantification, a straight-forward method is classify & count [3] where a classifier, f is trained using the source data and then used to estimate $\hat{q}_c = \frac{1}{n'} \sum_{i=1}^{n'} \mathbb{1}(f(x_i), c)$. The classifier, f can also be adapted using unsupervised adaption methods [19, 5]. Further, unsupervised quantification has been studied extensively under class-distribution shift [9, 3, 16]. This work follow one of two main strategies. The first, introduced by Saerens et al. [16], and refined by [2], derive an EM algorithm for maximizing the likelihood of the target data, $p'(x)$ by iterative updates of $\hat{p}'(y)$ and $\hat{p}'(y|x)$. The second, discussed by several authors [3, 16, 17, 1], rely on the miss-classification rates (confusion matrix) estimated on the source data to adjust the estimated counts on the target data. Forman extended this method with several heuristic demonstrating significant performance increase for binary quantification [3].

For supervised quantification, simple random sampling [18] can be utilized to achieve an unbiased estimate the class-distribution: $\hat{q}_c = \frac{1}{b} \sum_{i=1}^b \mathbb{1}(y_{s_i}, c)$ where s is a vector of randomly permuted indices and b the annotation budget. Simple random sampling doesn’t utilize the classifier, f , but it can be incorporated using auxiliary sampling designs, through offset [1] or ratio [13] estimators. In these methods, some property (e.g. bias) of the classifier is estimated from the labeled subset, and then used to adjust the prediction on the whole target set. Other methods include adapting the classifier to operate in the target domain, typically achieved by modifying the internal feature representation of the classifier [19].

2 Data-sets

We introduce two large-scale, image data-sets from marine ecology. In both, repeated sets of collected survey images require quantification, and manual annotation is unfeasible due to the vast amounts of collected data. In both, a large set of labeled source images are available to train a classifier, and several randomly selected smaller sets are available for evaluation. Each such set is denoted a test ‘cell’ and the goal is to achieve accuracy quantification across all test cells. Domain shifts occur naturally in these data-sets, so that the class appearance, $p(x|y)$ and class-distributions, $p(y)$, varies across the test cells. However, the extent of these variations differ between the two data-sets as discussed below. A data-set overview is given in Table 1, and all data can be downloaded from www.eecs.berkeley.edu/~obeijbom/quantification_dataset.html.

2.1 Plankton population survey

Background: The Imaging Flow Cytobot (IFCB), is an *in-situ* instrument for measuring plankton populations [10, 11]. The IFCB is installed at an offshore tower at 4 m below water level at the Martha’s Vineyard Coastal Observatory, and collects images by automatically drawing seawater from the environment. From this stream of images ($\sim 100k \text{ day}^{-1}$) the ecologists need to quantify the daily plankton class-distribution. Manual annotation of the complete data stream is unfeasible, and is currently restrained to two randomly chosen hours each month. While insufficient for a complete ecological analysis, these randomly selected, fully annotated, image-sets are ideal for evaluation of quantification methods.

Details: We formalize the Plankton Survey quantification benchmark as follows. All IFCB labeled data from 2006-2013 is considered as pertaining to the source domain. The 21 randomly selected

Table 1: Data-set summary. n = number, avg. = average

Data-set	n train samples	n test cells	avg. n samples cell ⁻¹	n classes
Plankton Survey	3.3m	21	14248	33
Coral Survey	325k	15	1480	32

hours of annotated data from 2014 are the test cells. Only classes with > 1000 total samples are included in this benchmark, leaving around 3.3 million total labeled samples across 33 classes. The Plankton Survey data-set is dominated by a class-distribution shift (Figs. S7, S6).

2.2 Coral reef survey

Background: The XL Catlin Seaview Survey (XL CSS) is an ambitious project to monitor the world’s coral reefs [4]. Using underwater scooters, 2 kilometer of reef-scape is imaged each dive, with approximately one image meter⁻¹. The XL CSS has surveyed the Great Barrier Reef, the Coral Sea, the Caribbean, the Coral Triangle, the Maldives, and Chagos, and captured over 1 million photographs. From this images set, ecologists are interested in quantifying the percent cover of key benthic substrates for each set of 30 consecutive images [4]. Percent cover for each image is estimated by classifying 50 patches extracted at random row & column locations in each image as pertaining to one of 32 classes [12], for a total of ~ 1500 patches across the 30 images. Similarly to the Plankton Survey, we can think of these sets of patches as cells, each requiring quantification. Manual quantification of all cells is unfeasible: the images from the Caribbean alone would require 30 person-years to annotate [4].

Details: We formalize the Coral Survey quantification benchmark as follows. A training-set of 324732 annotated patches extracted from 1505 images constitute the source domain. In addition, 15 randomly selected sets of 30 consecutive images, each with 50 annotated patches, are the test cells (Fig. S2). Both class appearance and class-distribution varies across the test cells (since they are drawn from different location across the Caribbean), meaning that the data-set shifts are more complex than for the IFCB data-set (Figs. S3, S4).

2.3 Performance evaluation

Let $q^{(m)} = P(y)$ be the normalized ground-truth class-distribution: $q^{(m)} \in \mathbb{R}^c$, $\sum_{j=1}^c q_j^{(m)} = 1 \forall m$ for sampling cell m , and \hat{q} the estimated distribution. We measure the distance between $q^{(m)}$ and $\hat{q}^{(m)}$ using the Bray-Curtis distance, which is commonly used in ecology. For normalized class counts, it reduces to the l1-norm: $h^{BC}(q^{(m)}, \hat{q}^{(m)}) = \frac{|q^{(m)} - \hat{q}^{(m)}|_1}{2}$. The utility of a quantification method is evaluated by the average Bray-Curtis distance across the sampling cells.

3 Experiments

For all experiments, we train a Convolutional Neural Network, f , of the Alexnet network architecture [8], on the source data using Caffe [6]. More details are given in [11, 1].

Unsupervised quantification: We evaluate four unsupervised quantification methods. Applying Classify & Count is straight-forward, and creates a natural baseline (Fig. 1). The EM-algorithm of [16] is also evaluated (Fig. S8) along with the confusion-matrix (CM) correction method of [3, 17, 16]. The latter method requires inverting the confusion-matrix, which can be problematic for multi-class problems, since inversion requires full rank. We therefore apply the abundance correction for each class, i independently, by mapping the CM to a binary 2×2 matrix before inverting and estimating \hat{q}_i . We then normalized \hat{q} so that $\sum_{j=1}^c \hat{q}_j = 1$. Finally, we use a recent unsupervised domain-adaptation method which adapts the source net to the unlabeled data from each cell [19].

Supervised quantification: We investigate quantification performance for budgets, $10 < b < 150$ samples, which well covers the feasible range of what is economical for the respective surveys. Random sampling is used to establish an upper bound on the error, and the offset estimator is included as a simple improvement [1]. We also experimented with a ratio estimator [18], but this is impractical for small sample sizes [1]. Further, we used a DA baseline (‘DA mix’). In this method, the classifier, f was further fine-tuned on a mixture of 75% source and 25% target data (drawn from the b samples) for ~ 3 epochs. Finally, the supervised DA method of [19] was evaluated.




Figure 1: **Unsupervised results.** Quantification errors displayed as mean \pm SE for Classify & Count [3], the unsupervised Deep Transfer DA method of [19], distribution matching using the EM algorithm [16], and for correction using Confusion Matrix inversion [3, 17, 16].

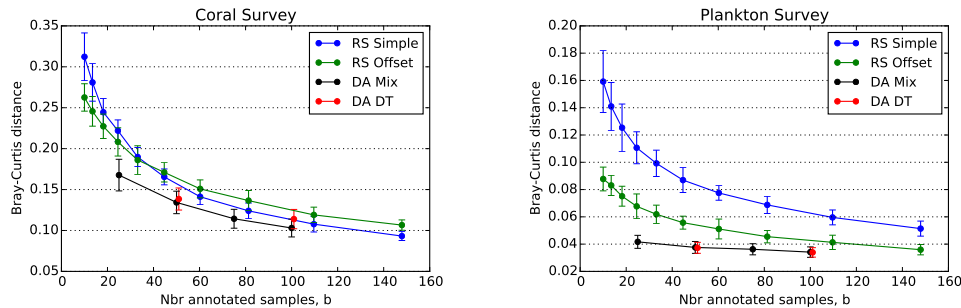


Figure 2: **Supervised results.** Quantification errors displayed as mean \pm SE for Simple random sampling [18], Offset sampling [1], DA mix, and Deep Transfer DA [19]

Results: For unsupervised quantification, our results indicate that the appropriate method depends on the nature of the data-set shift. For the Plankton Survey, the EM and CM correction methods work well, significantly lowering the estimation errors (Fig. 1). The EM method [16] outperformed the CM method [3], suggesting that this approach is more appropriate for high number of classes. However, for the Coral Survey, where the class-distribution shift assumption is violated, the CM and EM corrections corrupt the results, and simply counting the raw classification is preferable. The Deep Transfer DA method [19] was able to capture the data-set shift for the Plankton Survey, but produced inferior quantification compared to the CM and EM methods. Also note that the quantification results are, in general, stronger for the Plankton Survey, since it is an easier classification task with smaller data-set shift.

For supervised quantification, the two DA methods outperformed the random sampling baselines, in particular for smaller annotation budgets (Fig. 2). The adaptation method of [19] performed on par with DA mix. Among the random sampling methods, the offset estimator clearly outperformed simple random sampling on the Plankton Survey, but performed on-par for the Coral Survey. This is expected, as the hybrid estimator performs better when the classification errors are small, which they are in the Plankton Survey (Fig. 2; [1]).

Discussion: In pure class-distribution shift situations, as with the Plankton Surveys, the EM algorithm of [16] worked well, achieving a mean Bray-Curtis distance of $4.7 \pm 3.2\%$. Achieving such accurate quantification through simple random sampling would require $b \approx 150$ samples. The DA mix method, which achieved $3.7 \pm 0.5\%$ at $b = 50$ and $4.1 \pm 0.4\%$ at $b = 25$, makes better use of the supervision. It is a compelling alternative overall since it performed well also on the more challenging Coral Survey data. This is important since, in a real-world situation, one may not now *a-priori* what type of data-set shift to expect for the new target data.

The fact that fine-tuning of a deep neural network can be achieved with such small amount of target data (25 samples) is surprising, and deserves further investigation. Further, while the EM method presented here didn't perform very strongly, Forman suggested several improvements for binary classification [3]. We were unable to generalize these to the multi-class case, but it deserves attention. Finally, we think active sampling methods offer much promise, with collected samples utilized either to correct the raw classification counts, or for fine-tuning model parameters.

Acknowledgments: This work was supported by the National Oceanic and Atmospheric Administration grant No. NA10OAR4320156 and by the XL and Catlin Group Limited, Global Change Institute. We gratefully acknowledge the support of NVIDIA for their hardware donations.

References

- [1] O. Beijbom. Random sampling in an age of automation: Minimizing expenditures through balanced collection and annotation. *arXiv:1410.7074*, 2014.
- [2] M. C. Du Plessis and M. Sugiyama. Semi-supervised learning of class balance under class-prior change by distribution matching. *Neural Networks*, 50:110–119, 2014.
- [3] George Forman. Quantifying counts and costs via classification. *Data Mining and Knowledge Discovery*, 17(2):164–206, 2008.
- [4] M. González-Rivero, P. Bongaerts, O. Beijbom, O. Pizarro, A. Friedman, A. Rodríguez-Ramírez, B. Upcroft, D. Laffoley, D. Kline, C. Bailhache, et al. The catlin seaview survey—kilometre-scale seascape assessment, and monitoring of coral reef ecosystems. *Aquatic Conservation: Marine and Freshwater Ecosystems*, 24(S2):184–198, 2014.
- [5] H. Daumé III and D. Marcu. Domain adaptation for statistical classifiers. *Journal of Artificial Intelligence Research*, 2006.
- [6] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the ACM International Conference on Multimedia*, pages 675–678. ACM, 2014.
- [7] J. Jiang. A literature survey on domain adaptation of statistical classifiers. *Tech report*, 2008.
- [8] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [9] J. G. Moreno-Torres, T. Raeder, R. Alaiz-Rodríguez, N. V. Chawla, and F. Herrera. A unifying view on dataset shift in classification. *Pattern Recognition*, 45(1):521–530, 2012.
- [10] R. J. Olson and H. M. Sosik. A submersible imaging-in-flow instrument to analyze nano-and microplankton: Imaging flowcytobot. *Limnology and Oceanography: Methods*, 5(6):195–203, 2007.
- [11] E. C. Orenstein, O. Beijbom, E. Peacock, and H. M. Sosik. WHOI-Plankton- a large scale fine grained visual recognition benchmark dataset for plankton classification. *arXiv:1510.00745*, 2015.
- [12] E. Pante and P. Dustan. Getting to the point: Accuracy of point count in monitoring ecosystem change. *Journal of Marine Biology*, 2012.
- [13] R. M. Royall and W. G. Cumberland. An empirical study of the ratio estimator and estimators of its variance. *Journal of the American Statistical Association*, 76(373):66–77, 1981.
- [14] A. Royer and C. H. Lampert. Classifier adaptation at prediction time. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1401–1409, 2015.
- [15] K. Saenko, B. Kulis, M. Fritz, and T. Darrell. Adapting visual category models to new domains. *ECCV*, 2010.
- [16] M. Saerens, P. Latinne, and C. Decaestecker. Adjusting the outputs of a classifier to new a priori probabilities: a simple procedure. *Neural computation*, 14(1):21–41, 2002.
- [17] A. Solow, C. Davis, and Q. Hu. Estimating the taxonomic composition of a sample when individuals are classified with error. *Marine ecology. Progress series*, 216:309–311, 2001.
- [18] S. K. Thompson. *Sampling*. John Wiley & Sons, Inc., 2012.
- [19] E. Tzeng, J. Hoffman, T. Darrell, and K. Saenko. Simultaneous deep transfer across domains and tasks, 2015.

Supplementary Information: Quantification in-the-wild: data-sets and baselines

A Supplementary data-set information

A.1 Coral Survey

Image from the Catlin Seaview Survey can be accessed from catlinseaviewsurvey.com and <http://globalreefrecord.org/>. The particular sites studied in this paper are indicated in Fig. S2. From each image, 30 patches are cropped and classified. Some samples patches are shown in Fig. S1. The data-set shift in the Coral Survey data-set is shown in Fig. S4 and Fig. S3.




Figure S1: Sample images from the Coral Survey

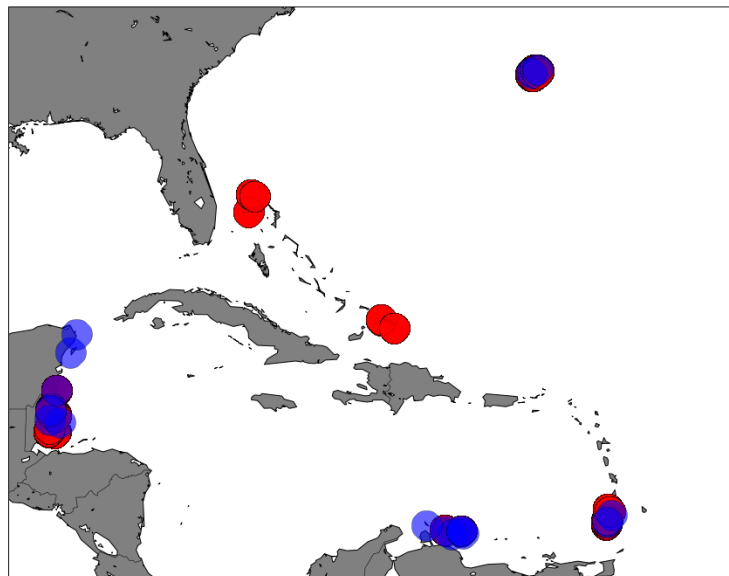


Figure S2: Caribbean locations used as for the Coral Survey benchmark. Blue circles indicate location of the test-cells and red circles the training images.




Figure S3: Class-distribution for the Coral Survey test cells. Total number of cell members is indicated above each cell




Figure S4: Area under ROC curve (mean \pm SD) for the 9 most abundance classes in the Coral Survey. The large difference indicate a significant shift in appearance, $p(x|y)$, between the cells.

A.2 Plankton Survey

The IFCB data-set can be downloaded at <https://github.com/hsosik/WHOI-Plankton> from which an interactive image viewer can be reached. Some sample images are shown in Fig. S5.




Figure S5: Sample images from the Plankton Survey




Figure S6: Area under ROC curve (mean \pm SD) for the 9 most abundance classes in the Plankton Survey. Note that the smaller y-axis range compared to Fig. S4.




Figure S7: Class-distribution for the Plankton Survey test cells, sorted by overall class abundance. Total number of cell members is indicated above each cell. Top: full distribution illustrating the dominance of the 'mix' class, bottom: highlight of the range between 80% and 100%

B Supplementary results

We include a convergence plot for the EM algorithm of [16] in Fig. S8




Figure S8: Convergence of the EM algorithm [16]. Since the class-distribution shift assumption of this method is not satisfied for the Coral Survey, it converges to a unsatisfactory state.

References

- [1] O. Beijbom. Random sampling in an age of automation: Minimizing expenditures through balanced collection and annotation. *arXiv:1410.7074*, 2014.
- [2] M. C. Du Plessis and M. Sugiyama. Semi-supervised learning of class balance under class-prior change by distribution matching. *Neural Networks*, 50:110–119, 2014.
- [3] George Forman. Quantifying counts and costs via classification. *Data Mining and Knowledge Discovery*, 17(2):164–206, 2008.
- [4] M. González-Rivero, P. Bongaerts, O. Beijbom, O. Pizarro, A. Friedman, A. Rodríguez-Ramírez, B. Upcroft, D. Laffoley, D. Kline, C. Bailhache, et al. The catlin seaview survey–kilometre-scale seascape assessment, and monitoring of coral reef ecosystems. *Aquatic Conservation: Marine and Freshwater Ecosystems*, 24(S2):184–198, 2014.
- [5] H. Daumé III and D. Marcu. Domain adaptation for statistical classifiers. *Journal of Artificial Intelligence Research*, 2006.
- [6] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the ACM International Conference on Multimedia*, pages 675–678. ACM, 2014.
- [7] J. Jiang. A literature survey on domain adaptation of statistical classifiers. *Tech report*, 2008.
- [8] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [9] J. G. Moreno-Torres, T. Raeder, R. Alaiz-Rodríguez, N. V. Chawla, and F. Herrera. A unifying view on dataset shift in classification. *Pattern Recognition*, 45(1):521–530, 2012.
- [10] R. J. Olson and H. M. Sosik. A submersible imaging-in-flow instrument to analyze nano-and microplankton: Imaging flowcytobot. *Limnology and Oceanography: Methods*, 5(6):195–203, 2007.
- [11] E. C. Orenstein, O. Beijbom, E. Peacock, and H. M. Sosik. WHOI-Plankton- a large scale fine grained visual recognition benchmark dataset for plankton classification. *arXiv:1510.00745*, 2015.
- [12] E. Pante and P. Dustan. Getting to the point: Accuracy of point count in monitoring ecosystem change. *Journal of Marine Biology*, 2012.
- [13] R. M. Royall and W. G. Cumberland. An empirical study of the ratio estimator and estimators of its variance. *Journal of the American Statistical Association*, 76(373):66–77, 1981.

- [14] A. Royer and C. H. Lampert. Classifier adaptation at prediction time. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1401–1409, 2015.
- [15] K. Saenko, B. Kulis, M. Fritz, and T. Darrell. Adapting visaul category models to new domains. *ECCV*, 2010.
- [16] M. Saerens, P. Latinne, and C. Decaestecker. Adjusting the outputs of a classifier to new a priori probabilities: a simple procedure. *Neural computation*, 14(1):21–41, 2002.
- [17] A. Solow, C. Davis, and Q. Hu. Estimating the taxonomic composition of a sample when individuals are classified with error. *Marine ecology. Progress series*, 216:309–311, 2001.
- [18] S. K. Thompson. *Sampling*. John Wiley & Sons, Inc., 2012.
- [19] E. Tzeng, J. Hoffman, T. Darrell, and K. Saenko. Simultaneous deep transfer across domains and tasks, 2015.