Efficient Image Gallery Representations at Scale Through Multi-Task Learning

Benjamin Gutelman

benjamin.gutelman@booking.com

Booking.com

Tel Aviv, Israel

Pavel Levin

pavel.levin@booking.com

Booking.com

Tel Aviv, Israel

Figure 1: Example of a retrieval task based on gallery representations: accommodations are globally ranked according to the proximity of their photo gallery embeddings to the gallery of “seed” images. Multi-task training ensures that the embeddings capture business-relevant product attributes and not just image features.

ABSTRACT

Image galleries provide a rich source of diverse information about a product which can be leveraged across many recommendation and retrieval applications. We study the problem of building a universal image gallery encoder through multi-task learning (MTL) approach and demonstrate that it is indeed a practical way to achieve generalizability of learned representations to new downstream tasks. Additionally, we analyze the relative predictive performance of MTL-trained solutions against optimal and substantially more expensive solutions, and find signals that MTL can be a useful mechanism to address sparsity in low-resource binary tasks.

CCS CONCEPTS

• **Computing methodologies** → **Multi-task learning; Image representations; Neural networks.**

KEYWORDS

multi-task learning, neural networks, image set representations

ACM Reference Format:

Benjamin Gutelman and Pavel Levin. 2020. Efficient Image Gallery Representations at Scale Through Multi-Task Learning. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '20)*, July 25–30, 2020, Virtual Event, China. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3397271.3401433>

1 INTRODUCTION

Today’s web experience is becoming increasingly more visual. A product description in a typical e-commerce marketplace is usually accompanied by a related image gallery. Whether those images are uploaded by the sellers themselves, or sourced by the e-commerce platform, they tend to considerably enrich user experience and provide customers with a valuable informational resource at various stages of their decision-making process [5, 19].

A person looking through such a gallery can usually infer a great deal about the product, often even more than they can from textual descriptions or formatted specifications. For instance, in travel domain accommodation features like the quality of a swimming pool (Figure 2), or of the room view can be subjective and difficult to represent directly in a generalizable way. The problem of finding a meaningful representation for the entire photo gallery clearly becomes even more complex.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).
SIGIR '20, July 25–30, 2020, Virtual Event, China
© 2020 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8016-4/20/07.
<https://doi.org/10.1145/3397271.3401433>

Figure 2: Swimming pools from two different accommodations. While many qualifying details and nuances can be easily inferred from the photos by website visitors, those details are challenging to represent directly in a structured and usable format.

This work describes a deep learning-based solution to the problem of finding meaningful representations of image galleries in a large-scale e-commerce setting. The solution is designed to satisfy three important constraints: (i) invariance of the representations under gallery re-orderings, (ii) their universality and flexibility to new downstream tasks, (iii) feasibility to deploy the solution in a large-scale e-commerce setting. The first requirement is essentially a consistency requirement which says that the information we obtain from a set of images is independent of the order in which we process these images. We enforce this by explicitly designing the gallery encoder as a symmetric function. The second constraint is addressed by training the gallery encoder on multiple independent tasks through a multi-task learning approach. The third requirement reflects the fact that this research problem emerged from real-world information retrieval and recommendation needs, and it constrains us to favour faster and more scalable architectures.

2 BACKGROUND

2.1 Image representations

Recent advances in deep learning, particularly convolutional neural networks (CNN) enabled impressive performance gains in a number of challenging visual recognition tasks [9, 17, 28]. These successes led to a widespread industrial adoption. While the most direct application of a CNN classification model is to use its output to categorize or label photos, we can also use its intermediate layer activations to represent images in a semantically meaningful way. These representations are commonly referred to as *image embeddings*. Intuitively, a deep neural classifier loads raw pixel intensities and then extracts increasingly abstract representations through its layers, which eventually lead to a classification attempt in the space of high-level concepts. Therefore, it is forced to learn meaningful features without any direct supervision on what those features should be, and it is reasonable to expect them to be useful for related tasks.

It is important to note, that these representations are always limited to the domain in which the original model was trained. For example, representations obtained from a model trained on general domain ImageNet [3] data alone, would not be perfectly suitable to travel-related queries (e.g. price, location, “family-friendliness”, distance from landmarks, etc). To address such issues we often need to adapt, or fine-tune, the original model to be more suitable to the target domain. A common way to address such domain adaptation

problem is by freezing part of the model parameters (usually lower-level layers), and allowing the rest of the network to be trained on in-domain data. That way we do not need to re-learn many of the lower-level features (which may stand for edges, gradients, basic shapes, etc.), but instead focus on learning relevant higher-level abstractions [22].

2.2 Representing sets

Since image collections are just sets of individual images, any function processing them is by definition a set function. In practical terms, we want such encoding function to be permutation invariant, i.e. different re-orderings of the same set of input images should produce the same output.

Recently there has been a surge of work related to machine learning on set-valued inputs. Despite a significant portion of that work being motivated by specific applications, such as point cloud data, many of the theoretical results and proposed architectures are of general scope. The theory behind learning set representations is closely related to statistical theory of exchangeability [16, 33]. Attention-based architectures have also been extended to accommodate set learning setting [18, 31]. It is worth noting that although attention-based architectures have already proven very powerful in several domains (particularly NLP [4, 30]), we leave them outside the scope of this work mainly due to the practical requirement of fast model retraining and online retrieval times. Additionally, Zaheer et al. [33] provide some theoretical reassurance that conceptually simpler approaches can in principle provide sufficient expressive power.

2.3 Multi-task learning

Multi-task learning (MTL; [2, 34]) refers to an umbrella of machine learning strategies which incorporate training signals from multiple feedback sources. MTL is attractive in many practical settings because of its inductive bias to generalize to new tasks. This in turn boosts computational—and often statistical—efficiency of machine learning solutions. MTL methods in deep learning setting vary according to how they share learned knowledge between tasks, and can be roughly split into two broad categories: soft and hard parameter sharing models [24]. Soft parameter sharing approaches keep separate models for each tasks, but regularize models to be similar to one another in some sense (e.g. by adding an L^2 distance between corresponding parameters of different models to the overall loss function [6]). Alternatively, hard parameter sharing techniques directly reuse the same model parameters for different tasks. Because our goal is to build a single universal encoder, soft parameter sharing techniques are not directly applicable in our setting, and we focus on hard parameter sharing approaches instead. Within the hard parameter sharing approaches we also must limit ourselves to the architectures which naturally support a single encoder, thus some of the newer interesting techniques which leverage several encoders [27] are also not relevant for us. This work will focus exclusively on an architecture with a single encoder with multiple decoders/heads.

In order to train an MTL system, the overall loss function needs to combine individual tasks’ losses. The way it is typically achieved

is by defining the overall loss as a weighted average of individual losses [8]. Weighting individual losses is a difficult and largely unsolved problem in the industry [12]. One promising direction is to use adaptive weighting scheme according to each task’s homoscedastic uncertainty [13], a strategy which we empirically evaluate in this paper.

2.4 Related work

There have been a number of applications involving symmetric gallery encoders in recent years. Permutation-invariant encoders using memory networks [32] have recently been successfully applied in medical imaging domain, classifying cancer types from collections of histopathology images [11]. Zaheer et al. [33] applied their sum-decomposition method to several applications involving collections of images: vision-based arithmetics and outlier detection.

Multi-task learning has also been used extensively in computer vision applications. For example, R-CNN models [23] use multi-task learning to jointly predict object classes and locations of the bounding bosses, thus mixing both classification and regression losses. Kendall et al. [13] introduced uncertainty-based loss weighting strategy which was used to jointly learn semantic and instance segmentation as well as per-pixel depth, thus also successfully combining different types of losses. For more examples of MTL applications in computer vision domain see [7, 15, 21, 29, 34].

To the best of our knowledge there is no prior work which applies multi-task learning to obtain generalizable representations of sets of images.

3 METHODOLOGY

This section describes our data, model specifications, and our training procedure. On the high level, in order to encode an image gallery we first turn each image into a vector using an image encoder (shared across all input images). Then the set of image representation vectors is passed into a permutation-invariant function computing the overall gallery representation. Gallery representations are then fed into task-specific “heads” defined by shallow models with associated loss functions. Finally, we train the entire structure jointly in a multi-task setting (see Figure 3).

3.1 Dataset

In our work we use a proprietary dataset of 520,000 accommodations sampled from Booking.com listings, randomly split into train and development sets (9:1 split). Each accommodation comes with a photo gallery and 31 additional features which we use as training signals in our multi-task learning. Table 1 summarizes accommodation features which we use as outputs for model training.

The lengths of accommodations’ photo galleries range between 5 and 44, depending on the number of available images. The additional features, which we use as training tasks, are a mix of 17 binary classifications, 7 non-binary classifications (number of classes range between 6 and 1981) and 7 regression tasks with scales spanning distances in km, prices in euros, and values of review scores.

Figure 3: MTL approach to learning image gallery representation. Shared part of the model includes image encoder and the embeddings aggregation mechanism (VGG16 encoder and simple mean pooling in our case). The resulting gallery representation is then fed into task-specific 2-layer neural networks. Individual losses are combined into the total loss through a MTL approach. All trainable parameters (shown in yellow) are learned jointly.

3.2 Model architecture

The overall model consists of shared and task-specific components. The shared part is a VGG16 encoder [26] for individual images followed by the mean pooling function applied across all the image embeddings. Given an input gallery of N images (decoded to their RGB representations and resized using bilinear interpolation to $224 \times 224 \times 3$ dimension), the VGG16 encoder produces N vectors of size 512 which are reduced to an averaged vector via mean pooling. Note that symmetry under gallery re-orderings can be preserved using much more sophisticated mechanisms than just average pooling (see e.g. [11, 31, 33]), however those encoders come at a significant computational cost in a large-scale industrial system. We leave the detailed study of the relative benefits of using more advanced permutation-invariant encoders for future research.

The non-shared part of the model is a collection of task-specific learners. Each learner is a shallow neural network which takes 512 inputs from the gallery encoder described above, passes them through a single fully-connected layer with ReLU activation, followed by an output layer with appropriate activation (sigmoid, softmax or linear, for binary, multiclass or regression tasks respectively).

The overall model contains 12.8M trainable parameters: 7.1M belong to VGG16 image encoder and 5.7M to individual task learners. Note that VGG16 encoder consists of five VGG-blocks of convolutional layers, which have 14.7M parameters in total, however we freeze the first four blocks (roughly half of the weights) and only leave the last block as trainable, which is the reason for 7.1M trainable parameters in the image encoder component.

Task	Description	Cardinality
C1	City (for top destinations)	1981
C2	World region	9
C3	Country	107
C4	Type (e.g. Hotel, B&B, etc.)	28
C5	Size (bucketized # of rooms)	9
C6	Hotel segment	6
C7	Vacation rental type	6
B1	Is part of a hotel chain	Binary
B2	Is vacation rental	Binary
B3	Offers sea view	Binary
B4	Offers city view	Binary
B5	Offers mountain view	Binary
B6	Is in a block of flats	Binary
B7	Has a balcony	Binary
B8	Has a kitchen	Binary
B9	Has a bar	Binary
B10	Has a garden	Binary
B11	Access to the beach	Binary
B12	Has a picnic area	Binary
B13	Has a communal lounge	Binary
B14	Offers hiking options	Binary
B15	Has a swimming pool	Binary
B16	Has a fitness center	Binary
B17	Is in the city center	Binary
R1	Distance from city center	Number (km)
R2	Review score	Number (0-10)
R3	Average daily rate	Number (euros)
R4	Location score	Number (0-10)
R5	Content score	Number (0-1)
R6	Average booking window	Number (days)
R7	Average length of stay	Number (days)

Table 1: Accommodation features which we use as learning targets in our MTL set up.

3.3 Loss function

As we are in a multi-task setting, we need to take special care of the loss function. Given n distinct tasks, the joint loss can be written as a vector:

$$\mathcal{L}(\hat{y}_1, \dots, \hat{y}_n, y_1, \dots, y_n) \\ = [\mathcal{L}_1(\hat{y}_1(\mathbf{w}_{shared}, \mathbf{w}_1), y_1), \dots, \mathcal{L}_n(\hat{y}_n(\mathbf{w}_{shared}, \mathbf{w}_n), y_n)]$$

The overall loss vector contains n components, each corresponding to a task. $\mathcal{L}_i$ is either a categorical cross-entropy or mean squared error (MSE) function depending on the nature of the task. For any given data point $\hat{y}_i$ denotes the predicted value of task i under model parameters $(\mathbf{w}_{shared}, \mathbf{w}_i)$, while y_i is the ground truth. In order to find optimal model parameters $\mathbf{W} = (\mathbf{w}_{shared}, \mathbf{w}_1, \dots, \mathbf{w}_n)$, we reduce the loss vector to a scalar-valued loss function $\mathcal{L}^{tot}$ using one of the two weighting procedures described below.

Our first method is based on a common uniform scaling heuristic which assigns equal weight to each task’s loss. While we assign equal weights of 1 to each classification loss, we modify the heuristic to add a constant scaling factor to the regression MSE losses as

a simple way to adjust for the different scales of variation. This factor is chosen to be the inverse of the sample variance of the given regression variable in the training dataset. The minimization objective under this method (which we refer to as pseudo-uniform scaling) is the following:

$$\mathcal{L}_{tot} = \sum_{i \in cat, bin} \mathcal{L}_i + \sum_{j \in reg} \frac{1}{\sigma_j^2} \mathcal{L}_j \\ = - \sum_{i \in cat} y_i \log \hat{y}_i - \sum_{j \in bin} y_j \log \hat{y}_j + (1 - y_j) \log (1 - \hat{y}_j) \\ + \sum_{k \in reg} \sigma_k^{-2} |\hat{y}_k - y_k|^2$$

where σ_k^2 is the sample variance of a regression label k , and bin , cat , reg denote the sets of binary, multiclass and regression tasks respectively.

The second approach follows Kendall et al. [13] and uses homoscedastic uncertainty to assign relative weights to each output. In this method the scaling factors are dynamic, and are learned jointly with the model parameters:

$$\mathcal{L}_{tot} = \sum_{i \in cat, bin} \frac{1}{\sigma_i^2} \mathcal{L}_i + \sum_{j \in reg} \frac{1}{2\sigma_j^2} \mathcal{L}_j + \sum_{k \in cat, bin, reg} \log \sigma_k \\ = - \sum_{i \in cat} \sigma_i^{-2} y_i \log \hat{y}_i - \sum_{j \in bin} \sigma_j^{-2} [y_j \log \hat{y}_j \\ + (1 - y_j) \log (1 - \hat{y}_j)] + \frac{1}{2} \sum_{k \in reg} \sigma_k^{-2} |\hat{y}_k - y_k|^2 \\ + \sum_{l \in cat, bin, reg} \log \sigma_l$$

Unlike in the pseudo-uniform scaling case, each σ_k^2 is not directly calculated from data, but is instead a learned parameter.

3.4 Model training

The training pipeline is implemented with Tensorflow v1.12.0 [1] and executed on 32 vCPU’s and a single NVIDIA P100 GPU. The vCPU’s are used in parallel to load and preprocess the binary image data stored in TFRecords format, while the GPU is used to execute machine learning training passes over the transformed data batches. We use Adam optimizer [14] with $\beta_1 = 0.9$, $\beta_2 = 0.999$ and learning rate 1.0×10^{-4} for model training, which is run until the total loss in our validation set stops decreasing. This took roughly four epochs (1.6×10^6 iterations) for both versions of the total loss functions as described in Section 3.3. Each data point may include up to 44 images, which limited us to the maximum batch size of 4.

4 EVALUATION

Our work is motivated by many real-world recommendation and retrieval tasks which require some form of product representation. We hypothesize that image gallery alone provides a rich source of information, and that we can use MTL to learn how to compress visual gallery into useful features for downstream tasks. As such, our main evaluation criterion is how well our gallery encoding approach can extract features useful to new downstream tasks. To evaluate this we look at the gains in computational and statistical efficiency in a “hold out” task of predicting property star rating

Figure 4: 2D projections of property representations given by image gallery embeddings. The visualizations demonstrate that the encoder learned to extract meaningful domain-related features from photo galleries.

from image gallery (Section 4.1). Section 4.2 analyses relative performance of the two MTL approaches.

Figure 1 and Figure 4 provide limited qualitative evaluation of the embeddings. The former is a visualization of two examples of embeddings-based accommodation retrieval based on a collection of “seed” photos. The fast similarity search was performed using bucketed random projections version of the locality-sensitive hashing [10]. Figure 4 offers a 2D mapping of representations [20] learned by the gallery encoder on a random sample of accommodations.

4.1 New task: star rating prediction

We judge the quality of learned gallery embeddings by how generalizable they are to new downstream tasks. To evaluate this, we consider a “hold out” task of predicting star ratings of accommodations. The star rating feature was present in 250k properties in our dataset, but we did not use it in any way when training the overall MTL model¹. Table 2 compares the accuracy of our model

against two baselines. The first baseline (“Imagenet transfer”) substitutes the representations from our gallery encoder with averaged photo embeddings, where photo embeddings are obtained from a VGG16 model fully pretrained on ImageNet [3] data. The second baseline (“End-to-end model”) is a model of the same configuration as described in Section 3.2, but trained end-to-end on one specific objective of classifying property star rating.

As we can see in Table 2, our solution consistently outperforms both alternative approaches, even on relatively small data sets. The somewhat surprising fact that even end-to-end solution lags behind MTL embeddings shows that MTL helps with learning relevant features which generalize beyond just the training tasks. Figure 2 can provide some intuition as to why ImageNet transfer does not perform as well on our task: if the original model’s objective is only concerned with object classification, then it will ignore many of the characteristics which are relevant to describing the object’s quality.

In terms of statistical efficiency, the end-to-end model only catches up to MTL embeddings-based model trained on 50k data

¹Because we saw pseudo-uniform scaling outperforming Kendall et al. [13] on most training task in terms of individual in-task performance (Section 4.2), our evaluation

in this section as well as visualizations (Figures 1 and 4) are based on pseudo-uniform MTL

	MTL embeddings	ImageNet transfer	End-to-end model
50k (20% of data)	0.574	0.515	0.543
100k (40% of data)	0.597	0.523	0.558
150k (60% of data)	0.616	0.546	0.575
200k (80% of data)	0.622	0.532	0.586
250k (a full epoch)	0.623	0.534	0.583
1000k (four epochs)	0.631	0.542	0.598
Mean time per instance	14 ms (on CPU)	136 ms (GPU)	
Total time (4 epochs)	4 hrs (on CPU)	40 hrs (GPU)	

Table 2: Accuracy results of the hold out star rating prediction task. Using our MTL-trained embeddings consistently outperforms training the model end-to-end or using averaged ImageNet-trained embeddings.

points when it has 3 times as much data (150k). In terms of computational efficiency, the end-to-end model trained on NVIDIA P100 GPU and took 10 times more time per iteration than the MTL transfer model trained on CPU.

4.2 Performance on individual tasks

Even though the main purpose of training multiple tasks simultaneously is to enrich our universal gallery representations to improve their potential for generalization, it is still interesting to compare the performance of each of the individual tasks against the optimal task-specific models². Figure 5 summarizes the results. We look at two MTL approaches (pseudo-uniform scaling, and Kendall et al. loss [13]) and their predictive performance relative to optimal models individually trained end-to-end for each of the tasks.

While most tasks understandably show inferior predictive performance relative to the very strong baseline of optimal task-specific models, some tasks actually seem to benefit from MTL set up. Specifically C1, a classification task with the highest label cardinality, exhibits roughly the same performance as the optimal model under pseudo-uniform scaling set up. Additionally, tasks B1, B3 and B11 perform better than the optimal models for both MTL total loss types. None of the regression tasks show better performance under MTL training. In general, we see pseudo-uniform scaling to do better than Kendall et al. [13] method in 22 out of 31 tasks (71%). This is somewhat consistent with Senter et al. [25] who also found uniform heuristic to outperform uncertainty-based weighting in some MTL settings.

As most of our tasks are binary classification tasks, we zoom in on how MTL compares to task-specific models as a function of class balance, which we define as the fraction of the minority class labels in the feature (e.g. a perfectly balanced variable will be at 0.5, while tasks with significant class imbalance will be closer to 0). Figure 6 shows a clear trend for both types of MTL strategies that we tried. In terms of statistical performance, MTL seems to bring the most value to tasks with the highest class imbalance. We hypothesise that this is happening because more skewed tasks benefit more from additional learning signals, since the data for their underrepresented classes are more sparse.

²Optimal task-specific models share the same architecture as described in Section 3.2, but instead trained end-to-end on their respective objectives

Figure 5: Relative performance of MTL against task-optimal models. We measure performance degradation in terms of error increase (accuracy error for classification and mean squared error for regression tasks). Most binary tasks are within 5% relative difference against the optimal models. Multiclass and regression tasks generally do worse. Tasks B1, B3, B11 and C1 (for pseudo-uniform scaling only) actually do better in the MTL setup.

It is important to note that in a real-world setting developing 31 task-specific models is obviously significantly more expensive, both in terms of training and deployment. In many practical situations we would still favour a more computationally economical solution even if the alternative has a marginal quality improvement. Another important comment is that even though MTL “as is” performance on the tasks is on average lower than the optimally tuned models, they still significantly outperform basic benchmarks (majority class prediction for classifications and mean value for regressions) across every single task.

5 CONCLUSION

This work proposes a scalable deep learning solution to the problem of finding product representations based on the products’ image galleries. We use multi-task learning as the main mechanism to enforce the inductive bias in favour of learning generalizable representations. We demonstrate how our method outperforms two strong baselines on a new downstream task. In addition we analyze comparative performance of the two MTL approaches which we test on the training tasks, and show that in case of binary classification tasks there is a strong relationship between label distribution skew and relative performance of the MTL against task-specific models. This finding suggests that MTL can be a useful tool for dealing with sparsity for low-resource tasks. To the best of our

Figure 6: The relationship between class balance and MTL benefit. It seems that MTL approach degrades the predictive performance less and even provides improvement in situations of higher class imbalance. The linear relationship is statistically significant: Pearson's $r = 0.78$ ($p < .001$) and $r = 0.80$ ($p < .001$) in pseudo-uniform MTL and Kendall MTL respectively.

knowledge, this work offers the first assessment of MTL approach to building gallery-based product representations in a large-scale industrial setting.

ACKNOWLEDGMENTS

The authors wish to thank Shlomi Haver, Stas Girkin and Ivana Rebic for their help, as well as Ron Rostovsky, Ron Reish, Rubens Sonntag, Noa Barbiro and Nitzan Mekel-Bobrov for their support and insightful discussions.

REFERENCES

- [1] Martin Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2015. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. <http://tensorflow.org/>. Software available from tensorflow.org.
- [2] Rich Caruana. 1997. Multitask learning. *Machine learning* 28, 1 (1997), 41–75.
- [3] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*.
- [4] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [5] Wei Di, Neel Sundaresan, Robinson Piramuthu, and Anurag Bhardwaj. 2014. Is a picture really worth a thousand words? -on the role of images in e-commerce. In *Proceedings of the 7th ACM international conference on Web search and data mining*. 633–642.
- [6] Long Duong, Trevor Cohn, Steven Bird, and Paul Cook. 2015. Low Resource Dependency Parsing: Cross-lingual Parameter Sharing in a Neural Network Parser. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*. Association for Computational Linguistics, Beijing, China, 845–850. <https://doi.org/10.3115/v1/P15-2139>
- [7] David Eigen and Rob Fergus. 2015. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *Proceedings of the IEEE international conference on computer vision*. 2650–2658.
- [8] Ting Gong, Tyler Lee, Cory Stephenson, Venkata Renduchintala, Suchismita Padhy, Anthony Ndirango, Gokek Keskin, and Oguz H Elilbol. 2019. A Comparison of Loss Weighting Strategies for Multi task Learning in Deep Neural Networks. *IEEE Access* 7 (2019), 141627–141632.
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [10] Piotr Indyk and Rameev Motwani. 1998. Approximate nearest neighbors: towards removing the curse of dimensionality. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*. 604–613.
- [11] Shivam Kalra, Mohammed Adnan, Graham Taylor, and Hamid Tizhoosh. 2019. Learning Permutation Invariant Representations using Memory Networks. *arXiv preprint arXiv:1911.07984* (2019).
- [12] Andej Karpathy. 2019. Multi-task learning in the wilderness. <https://slideslive.com/38917690/multitask-learning-in-the-wilderness>
- [13] Alex Kendall, Yarin Gal, and Roberto Cipolla. 2018. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 7482–7491.
- [14] Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*. <http://arxiv.org/abs/1412.6980>
- [15] Iasonas Kokkinos. 2017. Ubertnet: Training a universal convolutional neural network for low-, mid-, and high-level vision using diverse datasets and limited memory. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 6129–6138.
- [16] Iryna Korshunova, Jonas Degra, Ferenc Huszar, Yarin Gal, Arthur Gretton, and Joni Dambre. 2018. Bruno: A deep recurrent model for exchangeable data. In *Advances in Neural Information Processing Systems*. 7190–7198.
- [17] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*. 1097–1105.
- [18] Juho Lee, Yoonho Lee, Jungtaek Kim, Adam Kosiorek, Seungjin Choi, and Yee Whye Teh. 2019. Set Transformer: A Framework for Attention-based Permutation-Invariant Neural Networks. In *International Conference on Machine Learning*. 3744–3753.
- [19] Xiao Ma, Lina Mezghani, Kimberly Wilber, Hui Hong, Robinson Piramuthu, Mor Naaman, and Serge Belongie. 2019. Understanding Image Quality and Trust in Peer-to-Peer Marketplaces. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 511–520.
- [20] Leland McInnes, John Healy, and James Melville. 2018. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426* (2018).
- [21] Ishan Misra, Abhinav Shrivastava, Abhinav Gupta, and Martial Hebert. 2016. Cross-stitch networks for multi-task learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 3994–4003.
- [22] Maxime Oquab, Leon Bottou, Ivan Laptev, and Josef Sivic. 2014. Learning and transferring mid-level image representations using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1717–1724.
- [23] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*. 91–99.
- [24] Sebastian Ruder. 2017. An overview of multi-task learning in deep neural networks. *arXiv preprint arXiv:1706.05098* (2017).
- [25] Ozan Sener and Vladlen Koltun. 2018. Multi-task learning as multi-objective optimization. In *Advances in Neural Information Processing Systems*. 527–538.
- [26] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
- [27] Trevor Standley, Amir R Zamir, Dawn Chen, Leonidas Guibas, Jitendra Malik, and Silvio Savarese. 2019. Which Tasks Should Be Learned Together in Multi-task Learning? *arXiv preprint arXiv:1905.07553* (2019).
- [28] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1–9.
- [29] Marvin Teichmann, Michael Weber, Marius Zoellner, Roberto Cipolla, and Raquel Urtasun. 2018. Multinet: Real-time joint semantic reasoning for autonomous driving. In *2018 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 1013–1020.
- [30] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *NIPS*.
- [31] Oriol Vinyals, Samy Bengio, and Manjunath Kudlur. 2015. Order matters: Sequence to sequence for sets. *arXiv preprint arXiv:1511.06391* (2015).
- [32] Jason Weston, Sumit Chopra, and Antoine Bordes. 2014. Memory networks. *arXiv preprint arXiv:1410.3916* (2014).
- [33] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Ruslan R Salakhutdinov, and Alexander J Smola. 2017. Deep sets. In *Advances in neural information processing systems*. 3391–3401.
- [34] Yu Zhang and Qiang Yang. 2018. An overview of multi-task learning. *National Science Review* 5, 1 (2018), 30–43.