

A Weak Supervision Learning Approach Towards an Equitable Mobility Estimation

Theophilus Aidoo,¹ Till Koebe¹, Akansh Maurya¹, Hewan Shrestha¹, Ingmar Weber¹

¹ Department of Computer Science, University of Saarland
taidoo@cs.uni-saarland.de, tkoebe@cs.uni-saarland.de, akma00001@uni-saarland.de, hesh00001@uni-saarland.de, iweber@cs.uni-saarland.de

Abstract

The scarcity and high cost of labeled high-resolution imagery have long challenged remote sensing applications, particularly in low-income regions where high-resolution data are scarce. In this study, we propose a weak supervision framework that estimates parking lot occupancy using 3m resolution satellite imagery. By leveraging coarse temporal labels—based on the assumption that parking lots of major supermarkets and hardware stores in Germany are typically full on Saturdays and empty on Sundays — we train a pairwise comparison model that achieves an AUC of 0.92 on large parking lots. The proposed approach minimizes the reliance on expensive high-resolution images and holds promise for scalable urban mobility analysis. Moreover, the method can be adapted to assess transit patterns and resource allocation in vulnerable communities, providing a data-driven basis to improve the well-being of those most in need.

Introduction

Unlike image classification tasks in natural images, where extensive labeled datasets such as ImageNet (Deng et al. 2009) and MS COCO (Lin et al. 2014) are readily available, remote sensing applications suffer from a lack of large-scale labeled data. This shortage hinders the development of broadly applicable representations for tasks like universal parking lot estimation across the globe. The scarcity of labels is largely due to the prohibitively expensive and labor-intensive nature of manual annotation (Christie et al. 2018; Uzkent et al. 2019; Berg, Pham, and Courty 2022; Yue et al. 2022; Rufener et al. 2024). Moreover, a simple survey of image availability reveals that high-resolution captures are far more common for high-income countries than for low-income ones. For example, using UP42’s API console¹, we obtained 27 images for Berlin, Germany and only 15 for Accra, Ghana, under similar conditions (cloud cover below 20% and 30cm resolution) between 2020-01-01 and 2022-12-31. This observation supports the point made by Engstrom, Hersh, and Newhouse (2022) that high-resolution satellites tend to underrepresent low-income countries.

Parking-lot occupancy serves as a scalable, proxy for human presence and movement patterns in urban environments

(Rufener et al. 2024). Since the vast majority of daily trips involve personal vehicles, fluctuations in parking lot usage directly reflect commuting flows, and commercial activity. By predicting occupancy dynamics of parking lots, we can estimate human mobility trends on a global scale.

Given these challenges in obtaining high-quality labels and high-resolution imagery for low-income countries, alternative methods that require less precise annotations and low-resolution imagery are essential towards the estimation of generalizable parking lot occupancy. Weak supervision learning is one such approach. In this framework, models learn from labels that are inexact (providing only coarse information), incomplete (available for only a subset of the data), or inaccurate (noisy annotations) (Zhou 2018; Yue et al. 2022; Ratner et al. 2017; Uzkent et al. 2019). Despite their lack of precision, these weak labels can still serve as valuable signals for model training.

Our work explores weak supervision learning for determining occupied vs empty parking lots using 3m PlanetScope imagery provided by Planet² of parking lots in Germany. We aim to contribute to growing research of learning from noisy annotations in domains where ground-truth labels are scarce and expensive to obtain.

Related Works

Weak Supervision. According to Yue et al. (2022), a model is said to be weakly supervised if any of the following conditions hold: (1) there are few labeled instances among many training samples, (2) annotation is provided on a coarse level, or (3) training samples have labels that are only partially correct or noisy. The first condition (1), often referred to as *Incomplete Supervision*, involves learning from a dataset where only a small subset of instances is labeled (He et al. 2021). In this case, the challenge lies in generalizing from the limited labeled data to the much larger unlabeled portion. This typically involves active learning (Zheng et al. 2021), where human experts are used to annotate the unlabeled data points or semi-supervised learning (Li, Guo, and Zhou 2019; Zheng et al. 2021), where the model learns from both labeled and unlabeled data points.

The second condition (2), *Inexact Supervision* occurs when only coarse labels are available, and the task is to in-

fer finer-grained structures. In remote sensing applications, this often requires interpolation or multi-instance learning techniques to derive detailed insights from coarse data, e.g. identifying individual tree species in a forest from vegetation cover maps (Carbonneau et al. 2018; Wang et al. 2022; Li et al. 2023).

The third condition (3), also known as *Inaccurate Supervision*, is a learning paradigm where the annotation of labels may not always reflect the ground-truth (Zheng et al. 2021), thus the labels are not accurate or noisy. Noise in the labels can arise from incorrect annotations or flawed assumptions used to generate the labels. This type of supervision poses a unique challenge as models must learn to distinguish between genuine patterns in the data and noise in the labels (Li, Zhang, and Zhu 2019).

Parking Lot Occupancy Estimation. Unlike other satellite imagery providers, PlanetScope provides *near-daily* images, pending cloud-free skies. This high temporal resolution makes this data a potential fit to estimate parking lot occupancy. Using PlanetScope imagery of a parking lot captured on different days, Drouyer (2020) estimated the occupancy ratio, that is, the proportion of occupied spaces for a large commercial parking lot in the US. The model takes as input a series of images of a parking lot and returns the occupancy ratio. To achieve this, Drouyer uses three pixel-level features derived from PlanetScope imagery: brightness, gradient, and Pearson correlation with surrounding pixels. These features show strong correlations with ground-truth data captured by a webcam positioned next to the parking lot, with the median image calculated from a series of images when the lot was empty. In their work on estimating occupancy ratios, Zhao et al. (2022) use a high-resolution image (30 cm) for spatial guidance, as cars are easier to identify at that scale. They also use three PlanetScope images: a main image, a second image captured closest in time to the main image, and a third image taken no more than seven weeks apart from the main image.

The pair of images that accompany the main image is used to capture the temporal changes in the images. The use of both high resolution and pair of before/after images yielded a high occupancy correlation to the ground-truth data. The proposed approach achieved a low mean absolute error in vehicle counting, outperforming the baseline models. However, it still requires high-resolution imagery for spatial guidance, which may not be available for retrospective analyses or be expensive to obtain. Another work done in this area is the use of high-resolution SkySat images provided by Planet to label corresponding PlanetScope images as a regression-based task to classify vehicles on the road and further estimate the heading direction of vehicles (Van Etten 2024).

Our Approach

Existing methods for parking lot occupancy estimation often rely on high-resolution imagery, additional temporal data, which may not always be available or cost-effective or the approach does not generalize well to other parking lots. In contrast, our approach utilizes weak supervision learning, leveraging noisy labels derived from temporal patterns in

parking lot usage. This weak supervision arises from the inherent inaccuracy in our labeling assumption—specifically, we assume that parking lots of supermarkets and hardware stores in Germany are typically full on Saturdays and empty on Sundays, as most stores remain closed by national law on Sundays. This policy effect guides our methodology, allowing us to estimate occupancy without requiring extensive labeled data or high-resolution images. While this law and its effects apply widely, there are notable exceptions. For instance, in some communities, these parking lots may host events such as second-hand product sales also known as “flea markets” on Sundays on such parking lots. Similarly, shops may open on special Sundays, known as “verkaufsoffene Sonntag” to compensate for extended holiday closures. Another thing to note is that even though parking lots are expected to be full(er) on Saturdays, due to the popularity of the place, weather condition, and season actual busyness may vary. In addition to the inaccuracies in our labels, an observer can barely differentiate between subtle features such as individual cars and other structures such as shopping cart stands, garbage containers or other variations in parking lot usage from 3m PlanetScope imagery. Despite these challenges, our goal is to extract general patterns of presence or absence of cars from the data by leveraging the weak annotations described above. Once trained we expect to apply our model ‘as-is’ to new regions without fine-tuning or re-training.

Given a pair of images captured on a Saturday and a Sunday for a large parking lot ($\geq 10,000$ sqm), it is relatively easy for a human observer to determine which image corresponds to an occupied parking lot and which represents an empty one. For instance, the upper row of Figure 1 presents a typical example of such an image pair for a parking lot with a total area of 36,866 sqm. In this case, the image on the left appears more colorful and generally darker in the middle, which can be attributed to the presence of (parked) cars. In contrast, the image on the right exhibits a more homogeneous appearance, indicating an empty parking lot. However, this visual distinction becomes increasingly challenging for smaller parking lots ($\leq 5,000$ sqm), where larger patterns as in large parking lot become less visible. As shown in bottom row of Figure 1, for a parking lot with an area of 2,971 sqm, the variation in occupancy between the Saturday and the Sunday image is not as easily discernible. This limitation highlights the need for an automated, scalable approach based on computer vision to determine whether a parking lot is occupied or empty, particularly when visual cues are subtle or ambiguous.

In this work, we set out to overcome expensive labeling efforts usually required for supervised learning tasks. For our study, we used PlanetScope images worth around EUR 3,400. In comparison, acquiring a similar training dataset (i.e. about 400 parking lots captured on average 48 times across a span of 9 years, namely 2016 to 2024) for estimating parking lot occupancy based on a traditional car detection approach (i.e. high-resolution imagery + car detection model), it would cost between EUR 450,000 and EUR 500,000 depending on the vendor and discounts. Apart from the costs, the availability of (historic) high-resolution im-

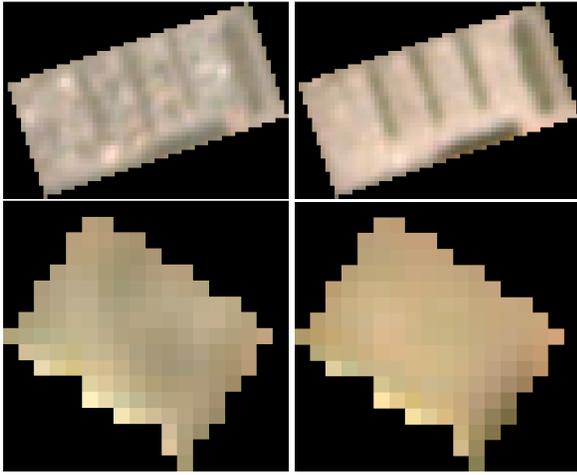


Figure 1: (**Top row**) A PlanetScope imagery showing the parking situation of a large parking lot observed on a typical Saturday (**left**), a typical Sunday (**right**). (**Bottom row**): A PlanetScope image showing the parking occupancy situation on a small parking lot observed on a Saturday (**left**), and a Sunday (**right**).

agery also limits the number of applications that can be explored. To address the above challenge, we leverage weak supervision learning to infer occupied and empty parking lots from satellite images captured on Saturdays and Sundays, respectively. Our contributions in this work include the development of a pipeline for acquiring parking lot polygons from OpenStreetMaps (OSM)³ based on Places of Interest. The full processing pipeline is published alongside the paper on GitHub⁴.

Data

As illustrated in Figure 2, we first identify Places of Interest (POI), specifically supermarkets and hardware stores (DIY), as these locations typically have dedicated parking lots. For each supermarket or hardware store, we retrieve the nearest parking lot polygon from OSM. In order to ensure that only clearly identifiable parking lots are selected, we also limit our search to rooftop and surface parking lots that are open to customers to reduce the inclusion of irrelevant parking areas, such as private or abandoned parking spaces, which could introduce further noise into the dataset. Additionally, we set a proximity threshold of 10 m around each POI to exclude distant parking lots that may not be directly associated with the identified POI. After applying these selection criteria, we identify a total of 745 parking lots across Germany. We obtain both 4-band and 8-band ortho-rectified surface reflectance PScene images using the Planet API for the summer period (April - September), from 2019 to 2023 and April - May in 2024. This season was chosen because cloud cover is generally lower. PlanetScope imagery is cap-

³<https://www.openstreetmap.org>

⁴https://github.com/Societal-Computing/equitable_mobility_estimation

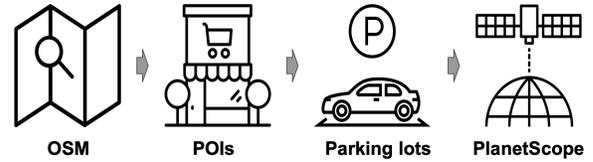


Figure 2: Data acquisition pipeline from POI identification on OSM to acquiring PlanetScope images.

ured by a constellation of small satellites at a spatial resolution of 3m, known as CubeSats, which capture images of the world on a near-daily frequency. However, due to variations in satellite coverage and stitching, some images only partially covered our parking lots, leading to inconsistencies in image dimensions. To address this issue, for a single parking lot, we removed all images that do not fully cover the area of interest.

Furthermore, we utilize the improved Usable Data Mask (UDM2), which accompanies each ortho-rectified analytic PlanetScope image, to filter out images containing cloudy pixels. Additionally, to maintain consistent brightness across images, we remove any image whose histogram distribution differs by more than 0.2 from that of the median image for its corresponding parking lot. This step helps to mitigate brightness variations introduced by different satellite sensor captures, ensuring uniformity in the dataset.

After these preprocessing steps, we retain 683 parking lots which include 355 large parking lots ($\geq 10,000$ sqm), 142 medium parking lots (between 5,000 sqm and 10,000 sqm) and 186 small parking lots ($\leq 5,000$ sqm).

Model

Our main contribution in this paper is the weak supervision learning approach from the observation made on parking lots of large supermarkets and hardware stores in Germany. We also show a baseline model which we use to evaluate our approach. The processing is illustrated in Figure 3 below. The data that goes into the model are image pairs, and their corresponding labels. We assign label 1 for a Saturday-Sunday pair for image 1 and image 2, respectively, and assign label 0 for a Sunday-Saturday pair for image 1 and image 2, respectively. Our pairwise comparison model consists of two encoders, namely **Encoder 1** and **Encoder 2** with weight sharing. Each encoder converts the input image into a 128 feature vector. Specifically, **Feature 1** for image 1 and **Feature 2** for image 2. We then do a simple difference between the feature vectors and the results serve as an input to a two-layer **MLP**. The purpose of the MLP is to project the 128 vector from the differences between Feature 1 and Feature 2 to a single output neuron. We then apply a sigmoid function to convert the logit values to probability values.

We use ResNet50 (He et al. 2016) as our encoder, based on its high accuracy on remote sensing classification tasks (Harini et al. 2024) and parking lot estimation tasks (Thakur et al. 2024), we also used a threshold value of 0.5 to sepa-

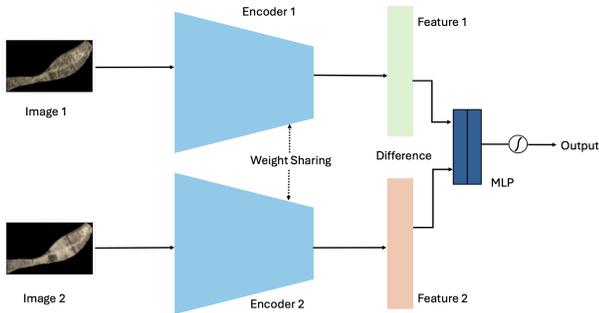


Figure 3: Given a pair of images, **Encoder 1** and **Encoder 2** are used to extract features of a lower dimension, yielding **Feature 1** and **Feature 2** for image 1 and image 2 respectively. The difference between the feature vectors is then projected to a single output with a Multi-Layer Perceptron (MLP).

rate output class 0 from 1. We will release a robust comparison with other encoders in future works as our main goal in this paper is to introduce the weak supervision learning approach. Also we consider our pairwise comparison model as a baseline to future approaches.

Experiments

We perform an 80-20 train-test split at the parking lot-level separately for large, medium, and small parking lots, using the same pairwise comparison approach. We observe from our results in Table 1, that for large and medium sized parking lots, the model is able to distinguish between occupancy difference between Saturday and Sunday images. However, for small parking lots our approach yields a low AUC score, which we attribute to the difficulty in identifying relatively occupied and empty parking lots of parking lot of small size. We consider a further analysis on small parking lots as a future work.

Parking lot size	AUC Score
large parking lots	0.92
medium parking lots	0.91
small parking lots	0.65

Table 1: AUC score of pairwise comparison for different parking lot sizes.

In our final experiment, we seek to validate our approach in the context of the on-going conflict between the Sudanese Armed Forces and the Rapid Support Forces that broke out in April 2023. As reported by Guo et al. (2024), there were fewer cars in Sudan during the post_war period (April 14–21, 2023) compared to the pre_war period (April 1–7, 2023), as reflected by a decrease in nitrogen dioxide (NO_2) concentrations, which are primarily emitted by vehicles and power plants.

To assess whether our method reflects this observation, we analyzed PlanetScope images of the Jackson Bus Terminal in Khartoum. Specifically, we select days of the week in both

the pre_war and post_war periods for which images were available, then apply our pairwise comparison approach to all possible image pairs (excluding pairs with the same date). As shown in Figure 4, images captured during the war period consistently received lower rankings than their pre_war counterparts, aligning with the observed drop in NO_2 concentrations.

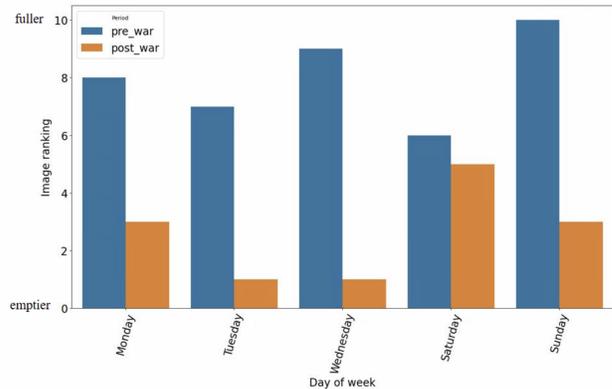


Figure 4: Pairwise comparison of day of week image rankings for Jackson Bus Terminal in Khartoum. The pre_war period spans 2023-04-01 to 2023-04-07, and the post_war period spans 2023-04-15 to 2023-04-21. Images captured during the war (orange bars) show consistently lower rankings than their pre-war counterparts (blue bars).

Conclusion

In this work, we proposed a weak-supervision learning approach to distinguish relative parking lot occupancies via pairwise comparisons. By exploiting Germany’s Sunday closure law for large supermarkets and hardware stores, we cheaply generated labels (“occupied” vs. “empty” pairs) without manual annotation. Our experiments on large, medium, and small lots show this proxy holds especially well for large and medium facilities. Notably, a model trained solely on large German lots was still able to detect mobility shift of a bus terminal in Sudan during the onset of the war. Despite the inexactness of our labels, this cost-effective strategy opens the door to many real-world applications, such as detecting population movements around Ukrainian borders, and measuring attendance drop-offs at hospital or market parking lots during epidemics. In low-income regions where high-resolution data are scarce and public transport hubs or large open-air markets stand in for parking lots, our weakly supervised framework can similarly reveal aggregate human mobility patterns. Future work will extend these case studies and explore real-time deployment for disaster response and humanitarian aid. Overall, our findings demonstrate that even “cheap” weak labels can unlock scalable, global insights into human mobility trends from satellite imagery.

Acknowledgments

This work is supported by funding from the Alexander von Humboldt Foundation and its founder, the Federal Ministry of Education and Research (Bundesministerium für Bildung und Forschung). We also express our gratitude to the European Space Agency for providing us with UP42 credits to access high-resolution images through the NoR scholarships program.

References

- Berg, P.; Pham, M.-T.; and Courty, N. 2022. Self-supervised learning for scene classification in remote sensing: Current state of the art and perspectives. *Remote Sensing*, 14(16): 3995.
- Carbonneau, M.-A.; Cheplygina, V.; Granger, E.; and Gagnon, G. 2018. Multiple instance learning: A survey of problem characteristics and applications. *Pattern Recognition*, 77: 329–353.
- Christie, G.; Fendley, N.; Wilson, J.; and Mukherjee, R. 2018. Functional map of the world. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6172–6180.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. IEEE.
- Drouyer, S. 2020. Parking occupancy estimation on planetscope satellite images. In *IGARSS 2020-2020 IEEE International Geoscience and Remote Sensing Symposium*, 1098–1101. IEEE.
- Engstrom, R.; Hersh, J.; and Newhouse, D. 2022. Poverty from space: Using high resolution satellite imagery for estimating economic well-being. *The World Bank Economic Review*, 36(2): 382–412.
- Guo, Z.; Abushama, H.; Siddig, K.; Kirui, O. K.; Abay, K.; and You, L. 2024. Monitoring indicators of economic activities in Sudan amidst ongoing conflict using satellite data. *Defence and Peace Economics*, 35(8): 992–1008.
- Harini, M.; Selvavarshini, S.; Narmatha, P.; Anitha, V.; Selvi, S. K.; and Manimaran, V. 2024. Resnet-50 integrated with attention mechanism for remote sensing classification. In *International conference on advances in distributed computing and machine learning*, 255–265. Springer.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- He, Y.; Yuan, X.; Chen, S.; and Wu, X. 2021. Online learning in variable feature spaces under incomplete supervision. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 4106–4114.
- Li, C.; Zhang, Z.; Liu, L.; Kim, J. Y.; and Sangaiah, A. K. 2023. A Novel Deep Multi-Instance Convolutional Neural Network for Disaster Classification From High-Resolution Remote Sensing Images. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*.
- Li, Y.; Zhang, Y.; and Zhu, Z. 2019. Learning Deep Networks under Noisy Labels for Remote Sensing Image Scene Classification. In *IGARSS 2019 - 2019 IEEE International Geoscience and Remote Sensing Symposium*, 3025–3028.
- Li, Y.-F.; Guo, L.-Z.; and Zhou, Z.-H. 2019. Towards safe weakly supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 43(1): 334–346.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, 740–755. Springer.
- Ratner, A.; Bach, S. H.; Ehrenberg, H.; Fries, J.; Wu, S.; and Ré, C. 2017. Snorkel: Rapid training data creation with weak supervision. In *Proceedings of the VLDB endowment. International conference on very large data bases*, volume 11, 269. NIH Public Access.
- Rufener, M.-C.; Offi, F.; Fatehkia, M.; and Weber, I. 2024. Estimation of internal displacement in Ukraine from satellite-based car detections. *Scientific Reports*, 14(1): 31638.
- Thakur, N.; Bhattacharjee, E.; Jain, R.; Acharya, B.; and Hu, Y.-C. 2024. Deep learning-based parking occupancy detection framework using ResNet and VGG-16. *Multimedia Tools and Applications*, 83(1): 1941–1964.
- Uzkent, B.; Sheehan, E.; Meng, C.; Tang, Z.; Burke, M.; Lobell, D.; and Ermon, S. 2019. Learning to interpret satellite images using wikipedia. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*.
- Van Etten, A. 2024. Vehicle Vectors and Traffic Patterns from Planet Imagery. *arXiv preprint arXiv:2406.06320*.
- Wang, X.; Xu, H.; Yuan, L.; Dai, W.; and Wen, X. 2022. A remote-sensing scene-image classification method based on deep multiple-instance learning with a residual dense attention ConvNet. *Remote Sensing*, 14(20): 5095.
- Yue, J.; Fang, L.; Ghamisi, P.; Xie, W.; Li, J.; Chanussot, J.; and Plaza, A. 2022. Optical Remote Sensing Image Understanding With Weak Supervision: Concepts, methods, and perspectives. *IEEE Geoscience and Remote Sensing Magazine*, 10(2): 250–269.
- Zhao, Q.; Xiao, J.; Wang, Z.; Ma, X.; Wang, M.; and Satoh, S. 2022. Vehicle counting in very low-resolution aerial images via cross-resolution spatial consistency and Intra-resolution time continuity. *IEEE Transactions on Geoscience and Remote Sensing*, 60: 1–13.
- Zheng, M.; Wang, F.; You, S.; Qian, C.; Zhang, C.; Wang, X.; and Xu, C. 2021. Weakly supervised contrastive learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10042–10051.
- Zhou, Z.-H. 2018. A brief introduction to weakly supervised learning. *National science review*, 5(1): 44–53.