

Reconsidering CO2 emissions from Computer Vision

Andre Fu

University of Toronto

andre.fu@mail.utoronto.ca

Mahdi S. Hosseini

University of New Brunswick

mahdi.hosseini@unb.ca

Konstantinos N. Plataniotis

University of Toronto

kostas@ece.utoronto.ca

Abstract

Climate change is a pressing issue that is currently affecting and will affect every part of our lives. It's becoming incredibly vital we, as a society, address the climate crisis as a universal effort, including those in the Computer Vision (CV) community. In this work, we analyze the total cost of CO2 emissions by breaking it into (1) the architecture creation cost and (2) the life-time evaluation cost. We show that over time, these costs are non-negligible and are having a direct impact on our future. Importantly, we conduct an ethical analysis of how the CV-community is unintentionally overlooking its own ethical AI principles by emitting this level of CO2. To address these concerns, we propose adding "enforcement" as a pillar of ethical AI and provide some recommendations for how architecture designers and broader CV community can curb the climate crisis.

1. Introduction

Our Society is experiencing an exponential growth in developing, training, and using deep-learning pipelines, particularly within Computer Vision (CV), the rapid growth has even surpassed human-level cognition [54]. Fueled by advances in computing, the state-of-the-art (SOTA) CV models are consistently beating benchmarks, with the most computationally intensive models achieving SOTA [40]. Consequently, developing and training SOTA models demand an inordinate computational burden, which corresponds to energy and environmental costs. Currently, there exist two methods for developing CV models, (1) Heuristic hand-crafted model development; and (2) CV Neural Architecture Search (CV-NAS) algorithms.

In order to design SOTA models, using either hand-crafted or NAS methods, some form of searching for an optimal architecture is required. With hand-crafted networks, the search phase is the iterative training for optimization and with NAS, the search is embedded directly into the architecture generation pipeline. After architectures get frozen they are released to the public which leads to wide-scale adoption of popular architectures such as ResNet [30], VGG




Figure 1: The two phases of CO2 emissions: model development through search and life-time evaluation. As a model (with high computational footprint i.e. FLOPS) gets world-wide adoption, the CO2 emissions increase.

[47], and GoogLeNet [49] just to name a few. Then models experience an evaluation phase, over the life-time of these architectures, they are trained and evaluated a great number of times. Both these phases require enormous computational costs and equivalently enormous CO2 emissions. As such, an important area to consider is the CO2 impacts of the CV community, the ethical considerations of those impacts, and how we can curb the global climate crisis.

Within CV no such analysis of the CO2 has been conducted, and therefore this paper aims to do so; guided by the following questions:

- During the model-searching phase, how much CO2 is being emitted, and how can we curb these emissions?
- Over the lifetime of a model, can we quantify the total amount of CO2 emitted and recommend solutions to curb these emissions?
- Who are the marginalized communities affected by CO2 emissions, and how can we better serve them?

2. Methods & Results

We break the total environmental cost of the CV pipeline into two sub-components: (1) the cost of searching for an architecture and (2) the evaluation cost of an architecture over its lifetime. Within this work, the evaluation cost refers to CV researchers & practitioners training a pre-established model. We attempt to quantify both types of costs and produce a quantitative metric for how our actions today are generating CO2 emissions.

Searching Phase's CO2 emissions: In order to determine the amount of CO2 emissions from CV-NAS search-

ing phase, we used a Strubell et al. [48] inspired methodology. We begin by collecting the 45 major CV-NAS papers in the last three years (2018-2020) [11, 7, 61, 13, 37, 10, 42, 15, 21, 55, 60, 8, 56, 24, 25, 31, 38, 14, 59, 26, 53, 17, 45, 44, 46, 16, 57, 20, 22, 29, 36, 50, 35, 27, 43, 12, 51, 41, 40, 19, 18, 32, 58, 39, 62] finding 157 models. For every model, we extract the Top-1 Accuracy, Parameters, FLOPS, GPU hours and GPU type. Using the GPU type, we found the total power draw for that specific GPU [5] denoted as p_g . If no GPU type was found, we assumed the most common GPU type, the Nvidia V100 was used. Often CPU type and DRAM weren't included in the papers, so we disregarded DRAM and assumed the Intel Core i7-10750H was used as the CPU power draw, denoted p_c . We then multiply the combined power draw by the GPU hours and by the Power Usage Effectiveness, 1.59 [6] to obtain a total power metric. We then denote this total power draw as P_t in Watt-hours by:

$$P_t \text{ [Wh]} = 1.59 \times \text{GPU hours} \times (p_g + p_c) \quad (1)$$

We then convert Watt-hours to CO2 emissions through the US Environmental Protection Agency's (EPA) kWh to CO2 measurement. This value was taken as a national weighted average of all CO2 emissions across the United States, giving 0.707×10^{-3} metric tonnes/kWh [4].

$$\text{CO2 [kg]} = P_t \times 0.707 \times 10^{-3} \quad (2)$$

Evaluation Phase's CO2 emissions: As models get used over time, their FLOPS become a crucial factor in their computational cost. We can assess a model's evaluation cost over time, as the number of times a model is being trained compared to the CO2-emissions. We denote a model's FLOPS as f and quantify the Watt-to-FLOPS ratio as $\omega = \text{Watt}/f$. We then denote the GPU Watt-to-FLOPS as ω_g and the CPU's Watt-to-FLOPS as ω_c . We obtain a model's power draw, P_m over training by multiplying the model's flops, f by the Watt-to-FLOPS then multiply the total GPU hours to train, seen below:

$$P_m \text{ [Wh]} = f \times (\omega_g + \omega_c) \times \text{GPU hours} \quad (3)$$

Similar to Equation 2 we can then compute the emitted CO2 by multiplying the power draw from a model by the EPA's Wh to CO2 measurement, $\text{CO2} = P_m \cdot 0.707 \times 10^{-3}$.

2.1. Results

Searching Phase: Using Eq. 1 and Eq. 2 we can quantify CV-NAS' searching phase's CO2 cost. These initial costs can then be amortized over time to observe how the searched architectures perform compared to the most popular hand-crafted networks like ResNet [30]. With hand-crafted networks, the heuristics used are often derived over

multiple trials and iterative training methods until satisfactory. These experiments are often never reported, making it difficult to quantify the hand-crafted network's CO2 cost.

CV-NAS as an automated framework must have a trade-off between the top-1 accuracy compared to the CO2-emissions and relative FLOPS of a model. We highlight a few important models below in Figure 2, namely: NAT-M4 [40], DARTS [38], Once-For-All (OFA) [12] and FB-Net A [53]. These models demonstrate how CV-NAS can trade off the initial CO2 cost for either accuracy or computational complexity in terms of FLOPS.

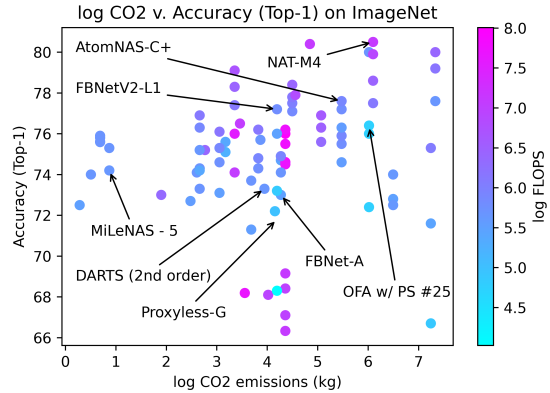


Figure 2: CO2 emissions of CV-NAS optimized models compared to their initial CO2 searching cost.

In Figure 2 we can see how the average FLOP models like MileNAS-5 have lower CO2 emissions but also average top-1 accuracy. Observe how OFA has lower computational complexity but also has a large initial CO2 cost, but doesn't beat the SOTA. Therefore, we can then trade-off the computational complexity to achieve SOTA by reducing our FLOP constraint, yielding NAT-M4. Low FLOP models like FBNet-A [53] and ProxyleSS-G [14] generally have average initial CO2 cost but only average accuracy.

Evaluation Phase: When a new model is released and becomes SOTA, it quickly becomes popular and can amass thousands of citations. For example, the ResNet architecture [30] published in CVPR 2016 has currently amassed over 73k citations [3] and is considered the standard for research. Assuming each paper ran their ResNet-backbone a modest 50 times to determine heuristic optimizations for their model, ResNet has been trained ≈ 3.6 million times. Therefore, model adoption has a huge impact on the total computational burden required to train a model, and by consistently using sub-optimal networks, we are incurring a large CO2 cost from evaluation. These large CO2 costs can be seen in Figure 3, which was calculated by assuming 250 epochs at 40 mins/epoch for ResNet and 60 mins/epoch for CV-NAS models. Then, we used Eq. 3 to derive a model's power draw and equivalent CO2 emission. The models cho-

sen to compare to ResNet in Figure 3 were to demonstrate the trade-offs between CO2-Accuracy-FLOPS.

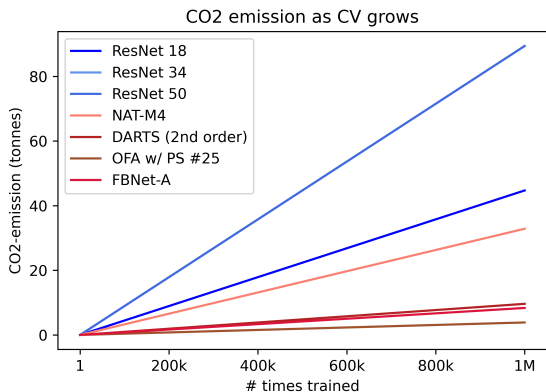


Figure 3: CO2 emissions of CV-NAS compared to heuristically optimized networks per 1 million times trained on ImageNet. The initial searching cost is negligible compared to the enormous evaluation cost over time.

Furthermore, we can also compare how hand-crafted networks compare to each other over their lifetimes so far. We do so by multiplying Eq 3 by the number of citations and 50 times trained per citation yielding Figure 4 below.




Figure 4: CO2 emissions of hand-crafted models over their lifetimes so far. The common architectures are ordered from left-to-right as lowest-to-highest number of citations

3. Discussion

To quantify the total amount of CO2 emissions emitted, we isolated two components: (1) the initial searching phase and (2) the evaluation phase. By addressing how each phase of a model’s life-cycle can be optimized to reduce CO2 emissions, we can lower the overall CO2-cost. We first identify how CO2 emissions tie into the global climate crisis, and identify marginalized groups who will be adversely affected by our actions. We then suggest some possible solutions to the CV-NAS and broader CV community.

The Current Climate Crisis: According to the Intergovernmental Panel on Climate Change (IPCC) [1] we are currently encountering a climate crisis that must be capped

at a 1.5°C change. They highlight that while 1.5° C is the ideal goal, we are currently not on track to reach a limit of 1.5°C. In 2016, we emitted 52 GtCO2 and by 2030 we will have 52-58 GtCO2, while we should be at 25-30 GtCO2 within the same timeframe. This crisis highlights how vital it is that we, within the CV-community enact change as any reduction in CO2 emissions can help curb the climate crisis. Additionally, the IPCC [1] highlights the unprecedented level of cooperation required between inter-government and inter-personal to reduce CO2 output to desirable levels.

The natural question underpinning this paper is the impact the CV community has on the CO2 emissions and if we can even make a substantive change to our environmental future. We believe that this question is valid, but mistakes the core idea behind curbing climate emissions, that everyone has a role to play, even the CV community.

Stakeholders: CO2 emissions are becoming a severe issue within society with considerable affect on the future, as demonstrated by the IPCC report. While everyone today and in the future is affected by the climate crisis, we scope this discussion to a rarely discussed marginalized group: our future generations. In identifying this group we discover subgroups who will have concrete impacts due to our CO2 emissions. Two such subgroups that will be impacted are those in the coastal and equatorial regions. In equatorial Africa, between 75 and 250 million people will be experiencing increased water-stress and rain-fed agriculture will be reduced by up to 50%. [1] Those in coastal regions will experience between 1 to 8 feet of sea level increase, engulfing their homes and livelihoods. These subgroups of future generations require immediate changes to our current CO2 emissions to minimize the effects of the climate crisis on them, and our role within as researchers within the CV community, should be to attempt to curb our emissions where possible.

Ethical considerations: Within the Ethical AI community, we’ve adopted [33, 28, 52] Principlism’s [9] four pillars: Respect for autonomy, Beneficence, Non-maleficence and Justice to analyze ethical concerns that have risen from computer-vision based dilemmas such as dataset bias and identifying marginalized communities. Through principlism, we can also analyze how stakeholders are unintentionally being affected by our actions today. In particular how we are overlooking all four pillars for the future generations, in effect showing that the ethical AI frameworks within CV are ill-enforced. Firstly, by emitting this level of atmospheric CO2, we overlook the future generation’s capacity to be self-determining by forcing them to cope with incredible amounts of agricultural stress and increased risk for extreme weather events. Secondly, we are currently disregarding beneficence by not acting in the benefit of future generations. Our unintended actions aren’t only non-benevolent but also maleficent. By causing harm to future generations

despite the warnings, we are inadvertently acting in direct contraposition to the ethical framework we’ve established. Finally, by emitting CO2 today that will have severe negative implications for future generations, we’re taking on the benefit while shifting the environmental cost to the future generations thus violating Justice. As such, we can see that despite establishing principalism as our working ethical framework, our unforeseen actions in emitting CO2 are going against those values.

3.1. Future Design Considerations

While the AI community has adopted principalism as a foundational tenet of ethical AI, the four pillars don’t necessarily reflect current societal and environmental considerations. As researchers, we have an obligation to maintain ethical considerations for marginalized groups. Therefore, we propose a *fifth pillar* for ethical AI, *enforcement*. Enforcement is the idea that our obligation to the marginalized groups should be included within our research, and these obligations should permeate our work. By holding ourselves accountable, we can act responsibly and attempt to minimize the global climate crisis.

CV-NAS community: Within the CV-NAS community we can adopt *enforcement* by incorporating CO2 considerations, providing models that not only minimize CO2 during search, but also ones that minimize over lifetime of the model. During initial searching, these models trade off accuracy at any computational and environmental cost.

By incorporating both types of CO2 cost into the optimization, newer models can balance CO2-emissions to the generated architecture’s optimization. We further highlight that while CV-NAS models have an initial searching cost (Fig. 2) it is often dwarfed by the life-time evaluation cost (Fig. 3). Therefore, by optimizing on CO2 for searching, but prioritizing low FLOP models, CV-NAS can have a considerable impact on CO2 emissions as a whole. In particular we note how the 3 branches of CV-NAS can better optimized: (a) *NAS-RL*, building on hardware-aware NAS to include a CO2/FLOP objective, (b) *Gradient-based*, including a CO2/FLOP constraint on the chosen-operations level of the bi-level [38] and (c) *Evolution*, adding a CO2/FLOP regularization parameter to the fitness function.

CV-NAS, like most of deep-learning, suffers from a lack of explainability. Often, the models are generated through black-box approaches, which may optimize for different criteria. We suggest that further research be done into CV-NAS explainability, as a glass-box approach to model optimization would allow the community to directly manipulate optimization parameters, meaning a greater focus on FLOP reduction which curbs long-term CO2 emissions. Furthermore, by incorporating FLOPS as a constraint we can also reduce the total cost over a lifetime seen in Figure 3. We can *enforce* optimizing on FLOPS by refusing to use high FLOP

networks, thereby communally penalizing researchers publishing these types of models. Optimizing on FLOPS has the added bonus of reduced evaluation computational power making them ideal for mobile, edge and low-power devices.

	CO2 (tons)	Cars Driven	Homes Powered
<i>ResNet</i> [30]	326.6	70.6	55.3
<i>VGG</i> [47]	181.7	39.8	30.8
<i>GoogLeNet</i> [49]	65.1	14.1	11

Table 1: To contextualize CO2 costs from the 3 most popular hand-crafted models (numeric values from Fig. 4), we contrast with cars driven & homes powered over 1 year. [2].

Broader CV community: In Figure 3 and Table 1 we can see that over time, choosing a sub-optimal model to conduct experiments has a severe and quantitative impact on the total CO2 produced. In line with *enforcement* we propose that the broader CV community adopt searched architectures from CV-NAS as baseline models in order to curb the amortized CO2 cost. If reviewers question why high FLOP models, like ResNet, was used as the optimized model over a lower FLOP cost model, our community can *enforce* the ethical AI principles. By adopting lower FLOP models such as FBNet-A or OFA, we can reduce the CO2 cost over a lifetime, allowing the broader CV community to also curb CO2 emissions.

One of the main reasons behind high-FLOP network development is image-scaling. Popular datasets like ImageNet [23] dominate research as the standard for generalization, but it come with high-FLOPS. Therefore, in order to curb the life-time evaluation CO2 cost, we encourage the CV-community to consider downscaled benchmarks for training & development such as Tiny-ImageNet [34].

4. Conclusion

Through this paper, we investigated a vitally important problem that is threatening our lives: Climate Change, and how the CV-NAS and broader CV community are contributing to the climate crisis. We then investigated the two main components behind CV CO2 emissions, (1) the initial architecture creation and (2) the evaluation of the network over its lifetime. We then provide a novel analysis of the CO2 emissions from CV-NAS and the CV-community, demonstrating that emissions are non-negligible. We then discuss which communities are at most risk and discuss how our community is inadvertently disregarding the Ethical AI principles. Therefore, we proposed a set of recommendations for the CV-NAS and broader CV community in order to curb the climate crisis. Further work can be done in this domain, linking environmental racism to CO2 output and the unintended consequences of the climate crisis.

References

- [1] Climate change and land: an ipcc special report on climate change, desertification, land degradation, sustainable land management, food security, and greenhouse gas fluxes in terrestrial ecosystems. <https://www.ipcc.ch/srcl>. 3
- [2] Epa greenhouse gases calculator, Oct 2018. 4
- [3] Nature index: Most influential papers research citations, Jul 2020. 2
- [4] Epa greenhouse gases equivalencies references, Feb 2021. 2
- [5] Techpowerup, recent database updates, Feb 2021. 2
- [6] Rhonda Ascierio. Uptime institute global data center survey 2020. *Uptime Institute*. 2
- [7] Anubhav Ashok, Nicholas Rhinehart, Fares Beainy, and Kris M. Kitani. N2n learning: Network to network compression via policy gradient reinforcement learning. In *International Conference on Learning Representations*, 2018. 2
- [8] Pouya Bashivan, Mark Tensen, and James J DiCarlo. Teacher guided architecture search, 2019. 2
- [9] Jonathan Beever and Andrew O. Brightman. Reflexive principlism as an effective approach for developing ethical reasoning in engineering. *Science and Engineering Ethics*, 22(1):275–291, 2015. 3
- [10] Gabriel Bender, Pieter-Jan Kindermans, Barret Zoph, Vijay Vasudevan, and Quoc V. Le. Understanding and simplifying one-shot architecture search. In *ICML*, pages 549–558, 2018. 2
- [11] Andrew Brock, Theo Lim, J.M. Ritchie, and Nick Weston. SMASH: One-shot model architecture search through hyper-networks. In *International Conference on Learning Representations*, 2018. 2
- [12] Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. Once-for-all: Train one network and specialize it for efficient deployment. In *International Conference on Learning Representations*, 2020. 2
- [13] Han Cai, Jiacheng Yang, Weinan Zhang, Song Han, and Yong Yu. Path-level network transformation for efficient architecture search. In *ICML*, pages 677–686, 2018. 2
- [14] Han Cai, Ligeng Zhu, and Song Han. ProxylessNAS: Direct neural architecture search on target task and hardware. In *International Conference on Learning Representations*, 2019. 2
- [15] Jianlong Chang, xinbang zhang, Yiwen Guo, GAOFENG MENG, SHIMING XIANG, and Chunhong Pan. Data: Differentiable architecture approximation. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. 2
- [16] Hanlin Chen, Baochang Zhang, Song Xue, Xuan Gong, Hong Liu, Rongrong Ji, and David Doermann. Anti-bandit neural architecture search for model defense. In *European Conference on Computer Vision*, pages 70–85. Springer, 2020. 2
- [17] Yukang Chen, Gaofeng Meng, Qian Zhang, Shiming Xiang, Chang Huang, Lisen Mu, and Xinggang Wang. Renas: Reinforced evolutionary neural architecture search. In *CVPR*, pages 4787–4796, 2019. 2
- [18] Xiangxiang Chu, Bo Zhang, Jixiang Li, Qingyuan Li, and Ruijun Xu. Scarlet-nas: bridging the gap between stability and scalability in weight-sharing neural architecture search. *arXiv preprint arXiv:1908.06022*, 2019. 2
- [19] Xiangxiang Chu, Bo Zhang, and Xudong Li. Noisy differentiable architecture search. *arXiv preprint arXiv:2005.03566*, 2020. 2
- [20] Xiangxiang Chu, Tianbao Zhou, Bo Zhang, and Jixiang Li. Fair DARTS: Eliminating Unfair Advantages in Differentiable Architecture Search. In *16th European Conference On Computer Vision*, 2020. 2
- [21] J. Cui, P. Chen, R. Li, S. Liu, X. Shen, and J. Jia. Fast and practical neural architecture search. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6508–6517, 2019. 2
- [22] Xiyang Dai, Dongdong Chen, Mengchen Liu, Yinpeng Chen, and Lu Yuan. Da-nas: Data adapted pruning for efficient neural architecture search. *arXiv preprint arXiv:2003.12563*, 2020. 2
- [23] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009. 4
- [24] Xuanyi Dong and Yi Yang. Network pruning via transformable architecture search. *arXiv preprint arXiv:1905.09717*, 2019. 2
- [25] Xuanyi Dong and Yi Yang. One-shot neural architecture search via self-evaluated template network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3681–3690, 2019. 2
- [26] Xuanyi Dong and Yi Yang. Searching for a robust neural architecture in four gpu hours. In *CVPR*, pages 1761–1770, 2019. 2
- [27] Jiemin Fang, Yuzhu Sun, Qian Zhang, Yuan Li, Wenyu Liu, and Xinggang Wang. Densely connected search space for more flexible neural architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10628–10637, 2020. 2
- [28] Luciano Floridi, Josh Cowls, Monica Beltrametti, Raja Chatila, Patrice Chazerand, Virginia Dignum, Christoph Luetge, Robert Madelin, Ugo Pagallo, Francesca Rossi, and et al. Ai4people—an ethical framework for a good ai society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4):689–707, 2018. 3
- [29] Chaoyang He, Haishan Ye, Li Shen, and Tong Zhang. Mile-nas: Efficient neural architecture search via mixed-level reformulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11993–12002, 2020. 2
- [30] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 1, 2, 4
- [31] Hanzhang Hu, John Langford, Rich Caruana, Saurajit Mukherjee, Eric Horvitz, and Debadeepta Dey. Efficient forward architecture search. *CoRR*, abs/1905.13360, 2019. 2
- [32] Andrew Hundt, Varun Jain, and Gregory D Hager. sharp-darts: Faster and more accurate differentiable architecture search. *arXiv preprint arXiv:1903.09900*, 2019. 2

- [33] Anna Jobin, Marcello Ienca, and Effy Vayena. The global landscape of ai ethics guidelines. *Nature Machine Intelligence*, 1(9):389–399, 2019. 3
- [34] Ya Le and Xuan Yang. Tiny imagenet visual recognition challenge. *CS 231N*, 7:7, 2015. 4
- [35] Changlin Li, Jiefeng Peng, Liuchun Yuan, Guangrun Wang, Xiaodan Liang, Liang Lin, and Xiaojun Chang. Blockwisely supervised neural architecture search with knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1989–1998, 2020. 2
- [36] Guohao Li, Guocheng Qian, Itzel C Delgadillo, Matthias Muller, Ali Thabet, and Bernard Ghanem. Sgas: Sequential greedy architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1620–1630, 2020. 2
- [37] Hanxiao Liu, Karen Simonyan, Oriol Vinyals, Chrisantha Fernando, and Koray Kavukcuoglu. Hierarchical representations for efficient architecture search. In *International Conference on Learning Representations*, 2018. 2
- [38] Hanxiao Liu, Karen Simonyan, and Yiming Yang. DARTS: Differentiable architecture search. In *International Conference on Learning Representations*, 2019. 2, 4
- [39] Zhichao Lu, Kalyanmoy Deb, and Vishnu Naresh Boddeti. Muxconv: Information multiplexing in convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12044–12053, 2020. 2
- [40] Zhichao Lu, Gautam Sreekumar, Erik Goodman, Wolfgang Banzhaf, Kalyanmoy Deb, and Vishnu Naresh Boddeti. Neural architecture transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 1, 2
- [41] Renqian Luo, Xu Tan, Rui Wang, Tao Qin, Enhong Chen, and Tie-Yan Liu. Neural architecture search with gbdt. *arXiv preprint arXiv:2007.04785*, 2020. 2
- [42] Renqian Luo, Fei Tian, Tao Qin, Enhong Chen, and Tie-Yan Liu. Neural architecture optimization. In *NeurIPS*, pages 7827–7838, 2018. 2
- [43] Jieru Mei, Yingwei Li, Xiaochen Lian, Xiaojie Jin, Linjie Yang, Alan Yuille, and Jianchao Yang. Atomnas: Fine-grained end-to-end neural architecture search. *arXiv preprint arXiv:1912.09640*, 2019. 2
- [44] Houwen Peng, Hao Du, Hongyuan Yu, Qi Li, Jing Liao, and Jianlong Fu. Cream of the crop: Distilling prioritized paths for one-shot neural architecture search. *Advances in Neural Information Processing Systems*, 33, 2020. 2
- [45] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V. Le. Regularized evolution for image classifier architecture search. In *AAAI*, pages 4780–4789, 2019. 2
- [46] Han Shi, Renjie Pi, Hang Xu, Zhenguo Li, James Kwok, and Tong Zhang. Bridging the gap between sample-based and one-shot neural architecture search with bonas. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1808–1819. Curran Associates, Inc., 2020. 2
- [47] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 1, 4
- [48] Emma Strubell, Ananya Ganesh, and Andrew McCallum. Energy and policy considerations for deep learning in NLP. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3645–3650, Florence, Italy, July 2019. Association for Computational Linguistics. 2
- [49] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015. 1, 4
- [50] Alvin Wan, Xiaoliang Dai, Peizhao Zhang, Zijian He, Yuandong Tian, Saining Xie, Bichen Wu, Matthew Yu, Tao Xu, Kan Chen, et al. Fbnetv2: Differentiable neural architecture search for spatial and channel dimensions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12965–12974, 2020. 2
- [51] Dilin Wang, Meng Li, Chengyue Gong, and Vikas Chandra. Attentionenas: Improving neural architecture search via attentive sampling. *arXiv preprint arXiv:2011.09011*, 2020. 2
- [52] Jess Whittlestone, Rune Nyrupe, Anna Alexandrova, and Stephen Cave. The role and limits of principles in ai ethics. *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, 2019. 3
- [53] Bichen Wu, Xiaoliang Dai, Peizhao Zhang, Yanghan Wang, Fei Sun, Yiming Wu, Yuandong Tian, Peter Vajda, Yangqing Jia, and Kurt Keutzer. Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. In *CVPR*, pages 10734–10742, 2019. 2
- [54] Ren Wu, Shengen Yan, Yi Shan, Qingqing Dang, and Gang Sun. Deep image: Scaling up image recognition, 2015. 1
- [55] Sirui Xie, Hehui Zheng, Chunxiao Liu, and Liang Lin. SNAS: stochastic neural architecture search. In *International Conference on Learning Representations*, 2019. 2
- [56] Yunyang Xiong, Ronak Mehta, and Vikas Singh. Resource constrained neural network architecture search: Will a submodularity assumption help? In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. 2
- [57] Xiang Wu Yibo Hu and Ran He. Tf-nas: Rethinking three search freedoms of latency-constrained differentiable neural architecture search. In *Proc. Eur. Conf. Computer Vision (ECCV)*, 2020. 2
- [58] Jiahui Yu, Pengchong Jin, Hanxiao Liu, Gabriel Bender, Pieter-Jan Kindermans, Mingxing Tan, Thomas Huang, Xiaodan Song, Ruoming Pang, and Quoc Le. Bignas: Scaling up neural architecture search with big single-stage models. In *European Conference on Computer Vision*, pages 702–717. Springer, 2020. 2
- [59] Chris Zhang, Mengye Ren, and Raquel Urtasun. Graph hypernetworks for neural architecture search. In *International Conference on Learning Representations*, 2019. 2
- [60] Xiawu Zheng, Rongrong Ji, Lang Tang, Baochang Zhang, Jianzhuang Liu, and Qi Tian. Multinomial distribution

learning for effective neural architecture search. *CoRR*, abs/1905.07529, 2019. [2](#)

[61] Zhao Zhong, Junjie Yan, Wei Wu, Jing Shao, and Cheng-Lin Liu. Practical block-wise neural network architecture generation. In *CVPR*, pages 2423–2432, 2018. [2](#)

[62] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V. Le. Learning transferable architectures for scalable image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. [2](#)