A. Impact of Fundamental Models on Earth Observation

Over the past few decades, remote sensing and Earth observation have had a transformative effect on a wide range of applications, including military, insurance, market prediction, and climate science, among others. Although this substantial impact cannot be directly ascribed to deep learning or large pre-trained networks, it forms part of a broader discussion that goes beyond the scope of this section. The focus here is on examining the role of fundamental models in enhancing Earth observation.

A.1. Contributions to Climate Mitigation and Adaptation

The use of machine learning on remote sensing data has become prevalent in devising solutions for an array of problems related to climate change [4, 9, 12, 16]. These solutions are primarily designed by curating datasets for specific tasks, which necessitates considerable resources. In addition, these solutions are usually tailored to particular regions as extending the methods to new geographies continues to be a significant challenge, largely due to the scarcity of labelled data [16]. Regions with less economic development, while equally vulnerable to the effects of climate change, often suffer from a deficit of effective remote sensing-based solutions [4]. Fundamental models for Earth observation can potentially tackle many of these concerns, thereby significantly hastening and facilitating the creation of novel remote sensing solutions for climate change.

A.2. Promoting Accessibility

By diminishing the requirement to curate a large labelled dataset for each individual task, we could democratize the development of machine learning models for remote sensing, particularly for groups or entities operating on limited budgets [1, 10]. Fundamental models could be particularly beneficial for non-profit organizations, academic institutions, startups, and developing countries. They could also pave the way for applications that were previously not profitable. We posit that the wider availability of these models will primarily have a net positive impact, although we recognize that this access could lead to unforeseen applications with potentially negative effects [3]. Moreover, it's important to note that these models may have dual-use implications, where they could, for instance, aid oil and gas industries in their operations in a way that either increases or decreases overall emissions.

A.3. Emissions from Large Pre-trained Models

Recent studies have examined the emissions of large neural networks [8, 11, 13–15]. Notably, training a large transformer can result in the emission of 284

when run on computers primarily powered by fossil fuel energy (US national average) [15]. When juxtaposed with individual actions, such emissions are substantial - a round-trip passenger flight from San Francisco to London results in 2.8

, which is roughly 100 times smaller. Yet, the wide applicability of pre-trained models and their potential in aiding efforts to mitigate climate change [12] prompts a shift in perspective.

Assessing new tools and systems necessitates a consideration of the probable net impact on emissions, both in terms of the tool's creation and its eventual deployment. For instance, testing the performance of airborne methane sensing tools at emission levels typically found in oil and gas operations can lead to the emission of about 7 metric tonnes of methane, roughly equivalent to 600

over a 20-year global warming potential [5]. Nevertheless, in a single day of operation, such an instrument can survey hundreds of sites, often identifying leaks that require repair and which emit considerably more than 7 metric tonnes of methane per day [7]. Similarly, fundamental models could significantly advance our capacity to utilize large amounts of passively collected satellite data, leading to massive reductions in emissions, enhancing our understanding of climate science qualitatively, and bolstering our ability to adapt to climate change.

In summary, the potential advantages for climate change mitigation through improved Earth observation methodologies likely outweigh the emissions associated with fundamental models. Furthermore, several actions can be undertaken to reduce and mitigate emissions linked to the training of your model [8]:

- Choose data centers that are certified as carbon neutral or predominantly powered by renewable energy, with efficient power usage (PUE). Such steps can drastically reduce emissions by about 50 times [8].
- Configure your code development process to minimize the need for computationally-intensive runs, for example, by using modular development and testing when possible.
- Improve the efficiency of your code and sparsify your network where feasible [11]. This could reduce emissions by up to tenfold.
- Opt for more energy-efficient hardware, such as TPUs or GPUs.
- Monitor [13] and report your emissions [8]. Better communication about climate change is vital for systemic changes. Improved documentation will assist other developers to continue from where you left off, possibly avoiding some computationally intensive runs.
- Offset the cumulative emissions of your projects.

A.4. Fairness and Biases

It's well known that large language models can amplify and perpetuate biases [2]. While this can lead to serious societal problems, we believe that biases in remote sensing models are likely to have a considerably lesser impact. However, we do foresee potential biases and fairness issues.

Data Coverage and Resolution Certain satellites provide standard spatial resolution and revisit rate coverage for the entire Earth (e.g., Sentinel-2 offers global coverage at a resolution of 10-60 m/pixel every five days). This ensures that imagery is freely and uniformly available across the planet. Other satellite data providers, such as Maxar, provide images on demand and have a higher spatial resolution (up to 0.3m per pixel), but have lower revisit rates and higher costs. Some countries, such as New Zealand, freely offer aerial imagery with a resolution of up to 0.1m per pixel[1]. Finally, it's worth mentioning that cloudy seasons in certain climates may limit data availability for some countries. Overall, while coverage is relatively uniform, some regions have much higher coverage than others, and financial constraints can limit access to data. This can lead to some degree of biases and fairness issues.

B. Hyper-parameters

The training and fine-tuning in our experiments follows the original MAE [6] training paradigm.

All models were pre-trained using the same hyper-parameters:

- **Effective batch size**: 2048 (32 per GPU × 64 GPUs)
- **Base learning rate (blr)**: 1.32×10^{-4}
- **Gradient clipping norm**: 1
- **Number of epochs**: 100
- **Warmup epochs**: 10
- **Weight decay**: 0.0457

For fine-tuning, we use an effective batch size of 32 with a weight decay of 0.05 and 5 warmup epochs.

The specific base learning rates (blr) used for each task were found through random search, see 1.

Table 1. Base Learning Rates (blr) for Fine-tuning Tasks Using Different Models

Task	Sentinel	Satellogic	Sentinel + Satellogic
m-bigearthnet	9.76×10^{-4}	3.12×10^{-4}	7.17×10^{-4}
m-so2sat	6.53×10^{-4}	2.12×10^{-4}	4.57×10^{-4}
m-brick-kiln	1.86×10^{-4}	1.69×10^{-5}	1.81×10^{-4}
m-forestnet	5.09×10^{-4}	5.64×10^{-4}	1.81×10^{-4}
m-eurosat	7.57×10^{-4}	4.60×10^{-4}	1.65×10^{-4}
m-pv4ger	5.56×10^{-4}	2.72×10^{-4}	4.05×10^{-4}

Table 2. Comparison of reconstruction losses under different training datasets and masking schemes. Satellogic data is generally harder to reconstruct due to its higher resolution. Random masking tends to be easier for the model to reconstruct, as it can leverage different bands and timesteps to recover missing information.

Training Data	Masking Schema	Reconstruction Loss
Satellogic	tunnel	0.561
Satellogic	random	0.458
Sentinel	tunnel	0.285
Sentinel	random	0.284
Sentinel + Satellogic	tunnel	0.284

[1] https://data.linz.govt.nz/

References

[1] Hamed Alemohammad. The case for open-access ML-ready geospatial training data. In *International Geoscience and Remote Sensing Symposium*. IEEE, 2021. 1

[2] Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 610–623, 2021. 1

[3] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021. 1

[4] Marshall Burke, Anne Driscoll, David B Lobell, and Stefano Ermon. Using satellite imagery to understand and promote sustainable development. *Science*, 371(6535), 2021. 1

[5] EPA. Greenhouse Gas Emissions: Understanding Global Warming Potentials. Technical report, US Environmental Protection Agency, Feb. 2017. 1

[6] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. 2

[7] Forrest Johnson, Andrew Wlazlo, Ryan Keys, Viren Desai, Erin Wetherley, Ryan Calvert, and Elena Berman. Airborne methane surveys pay for themselves: An economic case study of increased revenue from emissions control. preprint, Environmental Monitoring, July 2021. 1

[8] Alexandre Lacoste, Alexandra Luccioni, Victor Schmidt, and Thomas Dandres. Quantifying the carbon emissions of machine learning. *arXiv preprint arXiv:1910.09700*, 2019. 1

[9] Lei Ma, Yu Liu, Xueliang Zhang, Yuanxin Ye, Gaofei Yin, and Brian Alan Johnson. Deep learning in remote sensing applications: A meta-analysis and review. *ISPRS journal of photogrammetry and remote sensing*, 152:166–177, 2019. 1

[10] M Maskey, H Alemohammad, KJ Murphy, and R Ramachandran. Advancing ai for earth science: A data systems perspective. *Eos*, 101, 2020. 1

[11] David Patterson, Joseph Gonzalez, Quoc Le, Chen Liang, Lluis-Miquel Munguia, Daniel Rothchild, David So, Maud Texier, and Jeff Dean. Carbon emissions and large neural network training. *arXiv preprint arXiv:2104.10350*, 2021. 1

[12] David Rolnick, Priya L Donti, Lynn H Kaack, Kelly Kochanski, Alexandre Lacoste, Kris Sankaran, Andrew Slavin Ross, Nikola Milojevic-Dupont, Natasha Jaques, Anna Waldman-Brown, et al. Tackling climate change with machine learning. *arXiv preprint arXiv:1906.05433*, 2019. 1

[13] Victor Schmidt, Kamal Goyal, Aditya Joshi, Boris Feld, Liam Conell, Nikolas Laskaris, Doug Blank, Jonathan Wilson, Sorelle Friedler, and Sasha Luccioni. CodeCarbon: Estimate and Track Carbon Emissions from Machine Learning Computing. 2021. 1

[14] Roy Schwartz, Jesse Dodge, Noah A Smith, and Oren Etzioni. Green ai. *Communications of the ACM*, 63(12):54–63, 2020. 1

[15] 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, 2019. 1

[16] Xiao Xiang Zhu, Devis Tuia, Lichao Mou, Gui-Song Xia, Liangpei Zhang, Feng Xu, and Friedrich Fraundorfer. Deep learning in remote sensing: A comprehensive review and list of resources. *IEEE Geoscience and Remote Sensing Magazine*, 5(4):8–36, 2017. 1