Dollar Street Dataset NeurIPS 2022 Datasets and Benchmarks Track Submission Supplementary Materials

Table of Contents

The Dollar Street Dataset Datasheet

The original questions are in **bold**. The subtext to each question is in *italics.* The answers are in plain text with no formatting. The questions were copied from the "Datasheets for Datasets" paper available online: https://arxiv.org/pdf/1803.09010.pdf

Motivation

The questions in this section are primarily intended to encourage dataset creators to clearly articulate their reasons for creating the dataset and to promote transparency about funding interests.

For what purpose was the dataset created?

Was there a specific task in mind? Was there a specific gap that needed to be filled? Please provide a description.

The Dollar Street dataset is a supervised image dataset derived from Gapminder's Dollar Street project (https://www.gapminder.org/dollar-street) that contains everyday household items from homes around the world. It was created with three goals in mind:

1. Make available a highly curated set of images with valuable metadata (e.g. country, monthly income) that is more closely representative of the geographic and socioeconomic diversity of the world when compared with existing image datasets.
2. Help combat bias in downstream applications (e.g. image classification) that utilize image datasets. Our evaluation results show that the Dollar Street dataset can add significant value to accuracy improvements when considering computer vision images that represent the geographic and socioeconomic diversity of the world.
3. Legally permit commercial usage. Concretely, this means that the dataset only contains CC-BY-licensed works.

Who created the dataset (e.g., which team, research group) and on behalf of which entity (e.g., company, institution, organization)?
Who funded the creation of the dataset?

If there is an associated grant, please provide the name of the grantor and the grant name and number

Gapminder (http://gapminder.org/) funded and collected the raw images. MLCommons (https://mlcommons.org), a non-profit organization aimed at enabling ML for everyone, is funding the hosting costs for the dataset for the ML community.

Any other comments?

No.

Composition

Most of these questions are intended to provide dataset consumers with the information they need to make informed decisions about using the dataset for specific tasks. The answers to some of these questions reveal information about compliance with the EU's General Data Protection Regulation (GDPR) or comparable regulations in other jurisdictions.

What do the instances that comprise the dataset represent (e.g., documents, photos, people, countries)?

Are there multiple types of instances (e.g., movies, users, and ratings; people and interactions between them; nodes and edges)? Please provide a description.

The instances represent photos of objects found inside homes around the world, along with associated metadata about the photo (i.e. object tags) and the homes (i.e. region, country, monthly income).

How many instances are there in total (of each type, if appropriate)?

A total of 38,479 photos were used, for a total of 101.3 GB of image data.

Does the dataset contain all possible instances or is it a sample (not necessarily random) of instances from a larger set?

The photos contain a sample of all possible objects in a home. The Dollar Street project at Gapminder selected a number of objects to photograph that appropriately captured everyday life at different income levels throughout the world. This was based on Gapminder's analytics on what appeared to matter the most to capture in images.

If the dataset is a sample, then what is the larger set? Is the sample representative of the larger set (e.g., geographic coverage)? *If so, please describe how this representativeness was validated/verified. If it is not representative of the* *larger set, please describe why not (e.g., to cover a more diverse range of instances, because instances were withheld or unavailable).*

The larger set is all possible objects in all homes in the world. The intention of this sample is to be as representative as possible of crucial differences that highlight variations in socioeconomic welfare for the same objects (such as what “stoves” look like across different regions of the world, as shown in Figure 1 of the paper), though it is difficult to validate if the sample is representative given the scope of the larger set.

What data does each instance consist of?

“Raw” data (e.g., unprocessed text or images) or features? In either case, please provide a description.

The raw data consist of image files (saved in .jpg and .png format) along with a .csv file containing metadata associated with each image.

Is there a label or target associated with each instance?

If so, please provide a description.

Yes, a column in the .csv file lists all possible objects shown in each image and acts as the “label”.

Is any information missing from individual instances?

If so, please provide a description, explaining why this information is missing (e.g., because it was unavailable). This does not include intentionally removed information, but might include, e.g., redacted text.

No.

Are relationships between individual instances made explicit (e.g., users’ movie ratings, social network links)?

If so, please describe how these relationships are made explicit.

Yes, a column in the .csv file indicates the home that the image belongs to. On average, a single home has 118 images.

Are there recommended data splits (e.g., training, development/validation, testing)?

If so, please provide a description of these splits, explaining the rationale behind them.

We recommend a 60:20:20 split for the dataset, ensuring that all content from a given family is contained within a given split in order to ensure independence between the splits. In the paper, we show a fine-tuned model using these splits. It can achieve significant accuracy improvements (over 60% as shown in Figure 6 that is included in the paper).

Are there any errors, sources of noise, or redundancies in the dataset?

If so, please provide a description.

To the best of the authors' knowledge, each image is unique and known sources of errors have been removed.

Is the dataset self-contained, or does it link to or otherwise rely on external resources (e.g., websites, tweets, other datasets)?

If it links to or relies on external resources, a) are there guarantees that they will exist, and remain constant, over time; b) are there official archival versions of the complete dataset (i.e., including the external resources as they existed at the time the dataset was created); c) are there any restrictions (e.g., licenses, fees) associated with any of the external resources that might apply to a future user? Please provide descriptions of all external resources and any restrictions associated with them, as well as links or other access points, as appropriate.

This dataset is self-contained and MLCommons (https://mlcommons.org/en/), as the source entity for the dataset, will host the dataset and pay for its hosting fees. The dataset will be available free of charge for its end users.

Does the dataset contain data that might be considered confidential (e.g., data that is protected by legal privilege or by doctor-patient confidentiality, data that includes the content of individuals' non-public communications)?

If so, please provide a description.

No.

Does the dataset contain data that, if viewed directly, might be offensive, insulting, threatening, or might otherwise cause anxiety?

If so, please describe why.

No.

Does the dataset relate to people?

If not, you may skip the remaining questions in this section.

Yes. All of our data comes from the homes of real people.

Does the dataset identify any subpopulations (e.g., by age, gender)?

If so, please describe how these subpopulations are identified and provide a description of their respective distributions within the dataset.

The dataset broadly identifies the world region (Africa, America, Asian and Europe) and country of the homes photographed. For their distribution, refer to Table 2 and Figure 5 in the main manuscript. But it does not target specific IID information.

Is it possible to identify individuals (i.e., one or more natural persons), either directly or indirectly (i.e., in combination with other data) from the dataset?

If so, please describe how.

During the collection process of the images, Gapminder reviewed and filtered all photos that could reveal the family's location or could contain sensitive information. However, images may contain individuals, and, although difficult, the demographic information (region, country, home name) could potentially be used to identify individuals. Given that these families agreed to participate in the Dollar Street project and gave the appropriate permissions for the public use of the data in this dataset, it is natural that the involved individuals who could be identified are aware of this. It is also worth noting that the images used in this dataset are already publicly available at https://www.gapminder.org/dollar-street.

Does the dataset contain data that might be considered sensitive in any way (e.g., data that reveals racial or ethnic origins, sexual orientations, religious beliefs, political opinions or union memberships, or locations; financial or health data; biometric or genetic data; forms of government identification, such as social security numbers; criminal history)?

If so, please provide a description.

During the collection process of the images, Gapminder reviewed and filtered all photos that could reveal the family's location or could contain sensitive information. However, there is still a possibility that images may contain objects or background items that could be considered sensitive by some. Given that these families agreed to participate, gave the appropriate permissions for their homes to be photographed and for these photos to be used in public, it is natural that the involved individuals are aware of this. That said, we provide a means through MLCommons for individuals to remove their images from machine learning purposes. It is also worth noting that the images used in this dataset are already publicly available at https://www.gapminder.org/dollar-street.

Any other comments?

No.

Collection process

The answers to questions here may provide information that allow others to reconstruct the dataset without access to it.

How was the data associated with each instance acquired?

Was the data directly observable (e.g., raw text, movie ratings), reported by subjects (e.g., survey responses), or indirectly inferred/derived from other data (e.g., part-of-speech tags, model-based guesses for age or language)? If data was reported by subjects or indirectly inferred/derived from other data, was the data validated/verified? If so, please describe how.

The data was acquired and curated by photographers working with Gapminder. Monthly income data was calculated from estimated consumption as described in Section 5.

What mechanisms or procedures were used to collect the data (e.g., hardware apparatus or sensor, manual human curation, software program, software API)?

How were these mechanisms or procedures validated?

Cameras were used to digitally capture the images by photographers. For further information, see Gapminder's FAQs at the following link: https://www.gapminder.org/dollar-street/about

If the dataset is a sample from a larger set, what was the sampling strategy (e.g., deterministic, probabilistic with specific sampling probabilities)?

Gapminder collected photographs of specific objects in homes, with the home owners permission, with the intention of this being a representative sample of everyday living.

Who was involved in the data collection process (e.g., students, crowdworkers, contractors) and how were they compensated (e.g., how much were crowdworkers paid)?

A team of photographers located globally collected all of the data on behalf of Gapminder. Gapminder found these photographers by advertising on international network websites, receiving recommendations from friends/colleagues, and by contacting some themselves based on previous work. Both professional and volunteer photographers collected the photos. Photographers were either compensated based on fair market value of their country or volunteered to help collect the data. Photographed families volunteered to participate and did not receive monetary compensation for their participation.

Over what timeframe was the data collected?

Does this timeframe match the creation timeframe of the data associated with the instances (e.g. recent crawl of old news articles)? If not, please describe the timeframe in which the data associated with the instances was created.

The data in this dataset was collected starting in 2014, nonetheless this has proved useful already as we demonstrate in Figure 6 of the paper. Since this is a living project (as in the data is continuously changing and evolving as life progresses), Gapminder will continue collecting data for the Dollar Street project, and MLCommons commits to keeping Dollar Street updated through periodic revisions.

Were any ethical review processes conducted (e.g., by an institutional review board)?

If so, please provide a description of these review processes, including the outcomes, as well as a link or other access point to any supporting documentation.

No. But please note that the images were collected with the consent of the home owners and Gapminder reviewed and filtered all photos that could contain sensitive information.

Does the dataset relate to people?

If not, you may skip the remainder of the questions in this section.

Yes.

Did you collect the data from the individuals in question directly, or obtain it via third parties or other sources (e.g., websites)?

Were the individuals in question notified about the data collection?

If so, please describe (or show with screenshots or other information) how notice was provided, and provide a link or other access point to, or otherwise reproduce, the exact language of the notification itself.

Yes, they were notified about the data collection process and effort. This informed consent and data collection was conducted by Gapminder (dollarstreet.org).

Did the individuals in question consent to the collection and use of their data?

If so, please describe (or show with screenshots or other information) how consent was requested and provided, and provide a link or other access point to, or otherwise reproduce, the exact language to which the individuals consented.

Yes. Please refer to the Dollar Street photo guide in the FAQs found at the following link: https://www.gapminder.org/dollar-street/about?

If consent was obtained, were the consenting individuals provided with a mechanism to revoke their consent in the future or for certain uses?
If so, please provide a description, as well as a link or other access point to the mechanism (if appropriate).

While Gapminder has collected permission from all participating homes, we will allow anyone to request that images from their home be removed from the dataset. While the Creative Commons license legally permits such use, it is understandable that people may not have given permission with this use case in mind. It is also possible that a particular piece of source imagery is in misalignment with the consent requirements under current or future legislation (e.g., GDPR). For example, new permissions may need to be gathered in a changing regulatory environment, or a subject may wish to revoke their consent under new provisions. In summary, as outlined above, MLCommons will permit easy withdrawal of content and handle legal issues with the dataset.

Has an analysis of the potential impact of the dataset and its use on data subjects (e.g., a data protection impact analysis) been conducted?

If so, please provide a description of this analysis, including the outcomes, as well as a link or other access point to any supporting documentation.

No.

Any other comments?

No.

Preprocessing/cleaning/labeling

The questions in this section are intended to provide dataset consumers with the information they need to determine whether the “raw” data has been processed in ways that are compatible with their chosen tasks. For example, text that has been converted into a “bag-of-words” is not suitable for tasks involving word order.

Was any preprocessing/cleaning/labeling of the data done (e.g., discretization or bucketing, tokenization, part-of-speech tagging, SIFT feature extraction, removal of instances, processing of missing values)?

If so, please provide a description. If not, you may skip the remainder of the questions in this section.

No.

Was the “raw” data saved in addition to the preprocessed/cleaned/labeled data (e.g., to support unanticipated future uses)?

If so, please provide a link or other access point to the “raw” data.

N/A.

Is the software used to preprocess/clean/label the instances available?

If so, please provide a link or other access point.

N/A.

Any other comments?

No.

Uses

These questions are intended to encourage dataset creators to reflect on the tasks for which the dataset should and should not be used. By explicitly highlighting these tasks, dataset creators can help dataset consumers to make informed decisions, thereby avoiding potential risks or harms.

Has the dataset been used for any tasks already?

If so, please provide a description.

Yes, and only with respect to this NeurIPS paper only. For fine-tuning and scoring image classification models. See Section 6 of the main manuscript for additional information.

Is there a repository that links to any or all papers or systems that use the dataset?

If so, please provide a link or other access point.

No.

What (other) tasks could the dataset be used for?
Is there anything about the composition of the dataset or the way it was collected and preprocessed/cleaned/labeled that might impact future uses?

For example, is there anything that a future user might need to know to avoid uses that could result in unfair treatment of individuals or groups (e.g., stereotyping, quality of service issues) or other undesirable harms (e.g., financial harms, legal risks) If so, please provide a description. Is there anything a future user could do to mitigate these undesirable harms?

The dataset users should keep in mind that the dataset was curated at a certain point in time. And as such, it is important to note that the socioeconomic income of a region is with respect to a certain point in time and conclusions should be drawn accordingly.

Are there tasks for which the dataset should not be used?

If so, please provide a description.

Data can be used for good or bad. This dataset should not be used to identify low resource income communities and intentionally take detrimental actions.

Any other comments?

No.

Distribution

Will the dataset be distributed to third parties outside of the entity (e.g., company, institution, organization) on behalf of which the dataset was created?

If so, please provide a description.

Yes. The dataset will be publicly available under CC-BY license.

How will the dataset be distributed (e.g., tarball on website, API, GitHub)?

Does the dataset have a digital object identifier (DOI)?

The dataset will be made publicly available for download through MLCommons. There is no DOI at this time. We will assign this ID after we get feedback on whether the dataset is deemed acceptable for public use by the reviewers of NeurIPS. We do not want to release something into the public that the reviewers may find objectionable.

When will the dataset be distributed?
Will the dataset be distributed under a copyright or other intellectual property (IP) license, and/or under applicable terms of use (ToU)?

If so, please describe this license and/or ToU, and provide a link or other access point to, or otherwise reproduce, any relevant licensing terms or ToU, as well as any fees associated with these restrictions.

The dataset will be distributed before NeurIPS 2022. It will be under a CC-BY license.

Have any third parties imposed IP-based or other restrictions on the data associated with the instances?

If so, please describe these restrictions, and provide a link or other access point to, or otherwise reproduce, any relevant licensing terms, as well as any fees associated with these restrictions.

No.

Do any export controls or other regulatory restrictions apply to the dataset or to individual instances?

If so, please describe these restrictions, and provide a link or other access point to, or otherwise reproduce, any supporting documentation

No.

Any other comments?

No.

Maintenance

These questions are intended to encourage dataset creators to plan for dataset maintenance and communicate this plan with dataset consumers.

Who is supporting/hosting/maintaining the dataset?
How can the owner/curator/manager of the dataset be contacted (e.g., email address)? Is there an erratum?

If so, please provide a link or other access point.

MLCommons's dataset working group handles hosting and maintenance. Please contact
dollarstreet@mlcommons.org with questions. Instead of an "erratum", we plan to publish updates to the emails that people use to request the dataset.

Will the dataset be updated (e.g., to correct labeling errors, add new instances, delete instances)?

If so, please describe how often, by whom, and how updates will be communicated to users (e.g., mailing list, GitHub)?

Yes. It will be updated on an as-needed basis, with updates sent to all emails provided by users who request data access. We will be putting together a Dollar Street dataset user mailing list to facilitate this update process.

If the dataset relates to people, are there applicable limits on the retention of the data associated with the instances (e.g., were individuals in question told that their data would be retained for a fixed period of time and then deleted)?

If so, please describe these limits and explain how they will be enforced.

No.

Will older versions of the dataset continue to be supported/hosted/maintained?

If so, please describe how. If not, please describe how its obsolescence will be communicated to users.

This depends. Broadly the answer is: No, as in the case that some data needs to be removed for legal or ethical reasons, we do not want to keep maintaining that data. But we welcome reviewer feedback regarding this (or any other) matter.

If others want to extend/augment/build on/contribute to the dataset, is there a mechanism for them to do so?

If so, please provide a description. Will these contributions be validated/verified? If so, please describe how. If not, why not? Is there a process for communicating/distributing these contributions to other users? If so, please provide a description.

If someone would like to contribute directly back to the Dollar Street project (Gapminder), we recommend emailing the Gapminder group at info@gapminder.org.

If someone would like to contribute directly to the Dollar Street machine learning dataset (MLCommons), we recommend joining the working group meetings (instructions here).

Any other comments?

No.

Data Access

The dataset will be made publicly available before NeurIPS 2022 (assuming no issues are identified during the peer review process). Public access will be hosted accessible from https://mlcommons.org/en/dollar-street

If necessary, we can provide limited access to the reviewers, assuming the link is not distributed. Please email dollarstreet@mlcommons.org for access to the dataset.

Statement of Responsibility

The authors hereby declare that they bear all responsibility for violations of rights and that this dataset is CC-BY-licensed.

Hosting and Maintenance Plan

The MLCommons organization will host the dataset and handle maintenance concerns like data removal in case of misuse of data. It is a well-funded non-profit with no concerns about its ability to deliver on these requirements. When made publicly available, individuals will be able to download the dataset from a Google Cloud Storage bucket for which MLCommons will for the download bandwidth costs on behalf of downloaders.

Persistent Dereferenceable Identifier

We will assign this ID after we get feedback on whether the dataset is deemed acceptable for public use by the reviewers of NeurIPS. We do not want to release something into the public that the reviewers may find objectionable.

Resources Used

The data download, analysis, and evaluation were run on cloud computing resources paid for by Coactive AI. We utilize cloud storage buckets for storing all related data (e.g., images, metadata, logs) and Colab notebooks using general purpose (i.e., n2-standard) virtual machines for analysis/evaluation. We used NVIDIA T4 GPU powered instances for training, which are the lowest cost GPU-based instances in the cloud for machine learning inference and small-scale training. Fine-tuning ResNet50 with the entire Dollar Street dataset, as in Figure 6, took about 2 hours for 15 epochs on NVIDIA T4 GPU instances.

Additional Data

We ran experiments of fine-tuning and evaluated the performance on different income socioeconomic quartiles. We found that socioeconomic status still impacts model accuracy. Overall, quartiles improve the most from fine-tuning on neighboring quartiles. Additional data is included in the figures below.

Figure 1a-d shows the effects of fine-tuning the data on a particular quantile (shown in the caption) and evaluating the model's performance on the other quantiles.

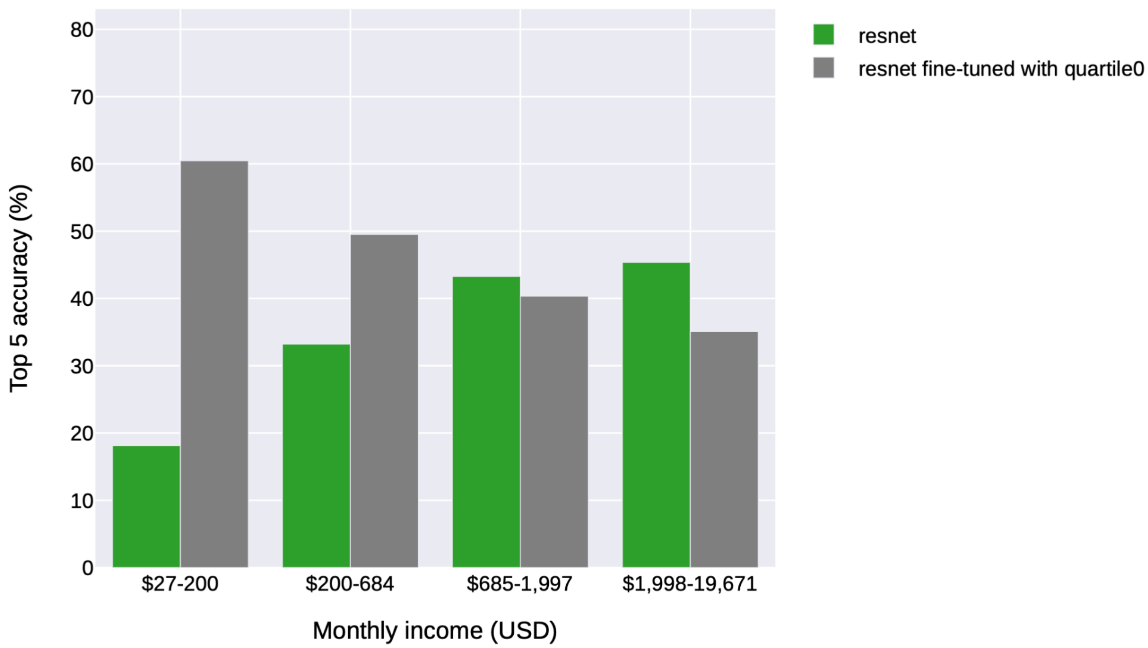

Figure 1a: Fine-tune on Quartile 0 ($27-200) and evaluate on all quartiles. We see that fine-tuning the pretrained resnet model dramatically improves performance on quartile 0, but results in monotonically decreasing performance with increasing monthly income.

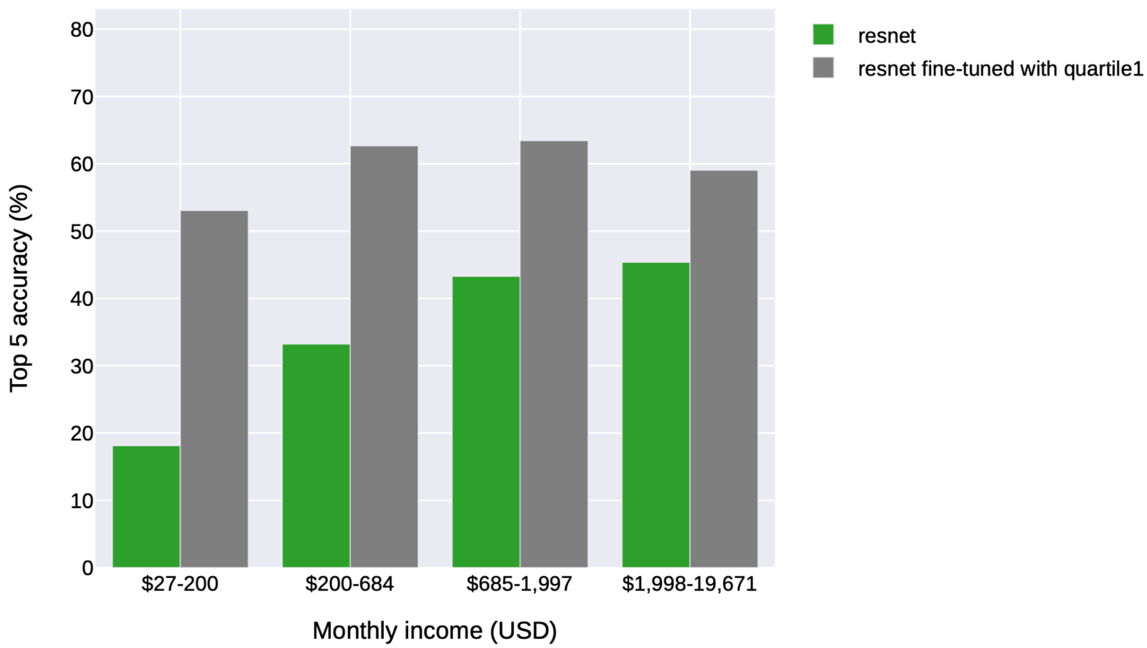

Figure 1b: Fine-tune on quartile 1 ($200-684) and evaluate on all quartiles. We see that fine-tuning the pretrained resnet model dramatically improves performance in quartile 1, but quartile 2 and quartile 3 are also improved more when compared to quartile 0.

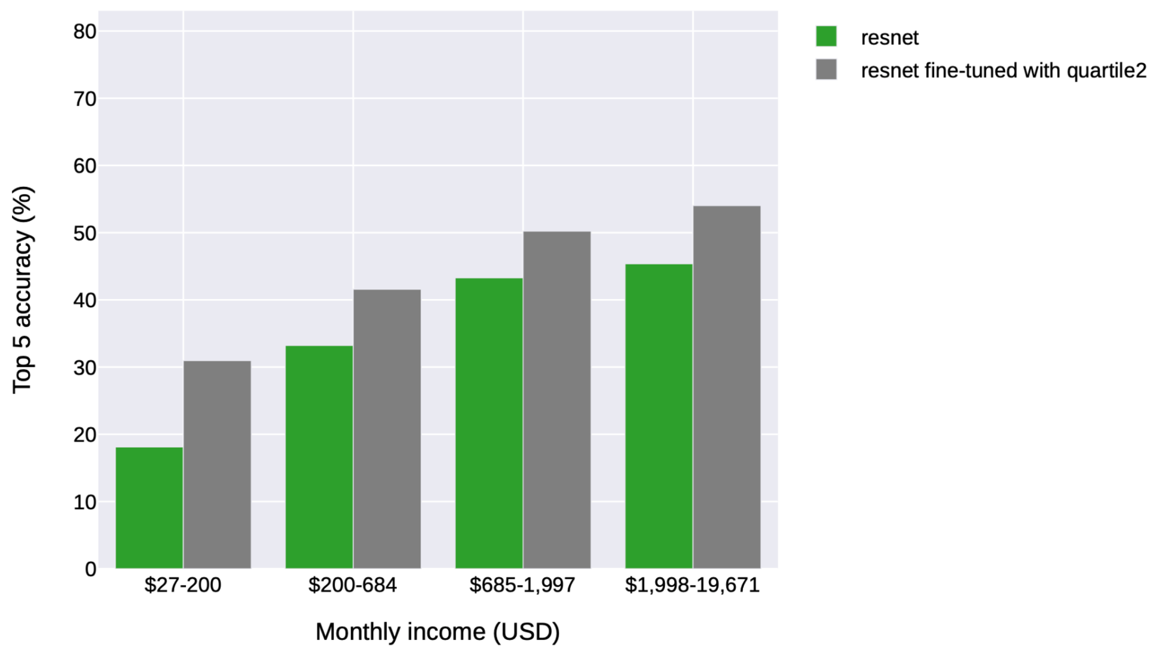

Figure 1c: Fine-tune on quartile 2 ($685-1997) and evaluate on all quartiles. We see that fine-tuning the pretrained resnet model helps improve performance in quartile 2 and quartile 3.

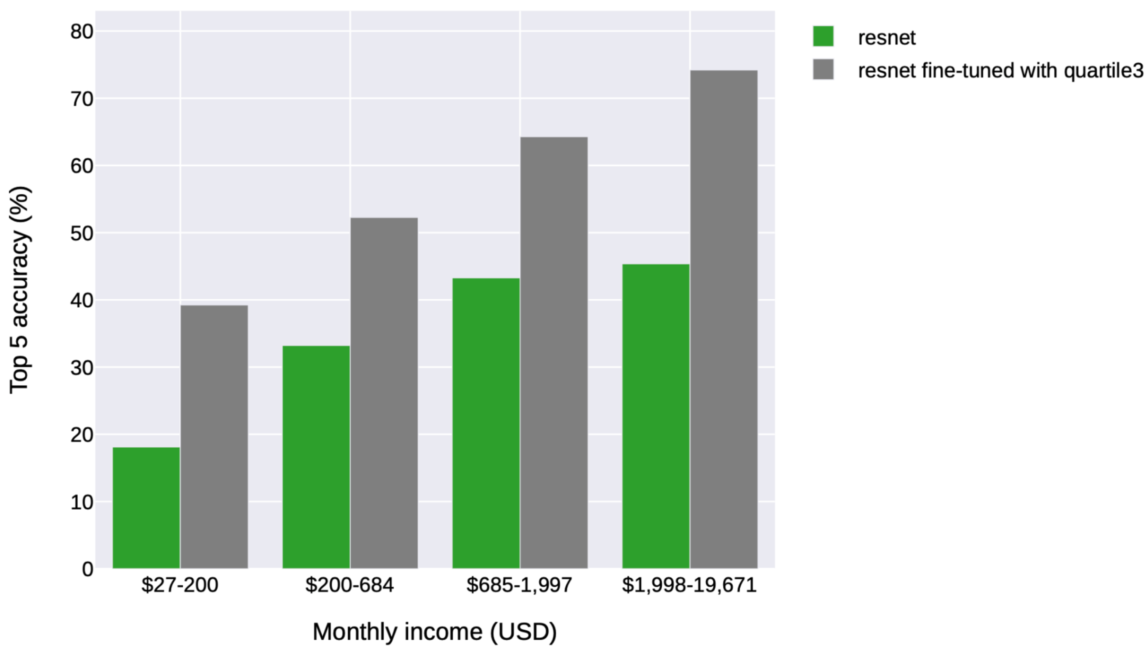

Figure 1d: Fine-tune on quartile 3 ($1998-19671) and evaluate on all quartiles. We see that fine-tuning the pretrained resnet model dramatically improves performance on quartile 3, but results in monotonically decreasing performance with decreasing monthly income.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes] See Section 1.
   (b) Did you describe the limitations of your work? [Yes] See Section 5.
   (c) Did you discuss any potential negative societal impacts of your work? [Yes] The societal impact of our work is discussed in Section 5 and Section 6.
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [No] The data was pulled from the Dollar Street project (`www.gapminder.org/dollar-street`). To avoid burdening them with additional bandwidth charges, we wanted to wait until MLCommons publicly hosts the data before releasing the code needed to reproduce the main experimental results. However, we can provide limited access to the reviewers if need be, assuming the link is not distributed and downloading the data from Gapminder is kept to a minimum.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] All details are specified in the manuscript.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] All details are specified in supplemental materials.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] All used pre-trained models were properly cited in Section 6.1.
   (b) Did you mention the license of the assets? [Yes] See section 3.2 for the image data and the pre-trained models are publicly available from torchvision [4].
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] The raw data is available from the Dollar Street project (`www.gapminder.org/dollar-street`). The downloaded and filtered dataset will be made publicly available, assuming no issues are identified during the peer review process. If necessary, we can provide limited access to the reviewers, assuming the link is not distributed.
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [Yes] Please refer to 3.4 and 3.3 for details.
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [Yes] Please refer to 3.4 and 3.3 for details.
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A] The raw data was sourced by Gapminder [11]. Please refer to 3.4 and 3.3 for details.
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A] The raw data was sourced by Gapminder [11]. Please refer to 3.4 and 3.3 for details.
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A] The raw data was sourced by Gapminder [11]. Please refer to 3.4 and 3.3 for details.

[4] `https://pytorch.org/vision/stable/models.html`

Appendix A Topics

The topics gathered from Gapminder is shown in Table 4.

Table 4: **Topic list.** Each Dollar Street dataset image is tagged from this set of 289 possible topics.

adding spices to food while cooking	earings	music idol	spices
agriculture land	eating	nature sceneries	starting stove
air cleaning equipments	electric wires	necklaces	storage room
air fresheners (scents	electricity wires	newspapers	stove/hob
alarm clock	elevators	next big thing you are planning to buy	street detail
alcoholic drinks	equipment	nicest shoes	street view
answering the phone	everyday shoes	opening and closing the freezer	surroundings
arm watch	family	opening and closing the refrigerator	switch on/off
arm watches	family eating	opening mail	table with food
armchair	family snapshots	opening the front door	tabloids
baby powder	favorite home decorations	other transport	taking a teaspoon of salt
backyard	favourite item in kitchen	oven	tattoos
bad outdoor air obstructions	favourite sports clubs	paper	teeth
baking sheets	fields	parking lot	thing i dream about having
baking tables	fishes	pen/pencils	things i wish i had
baking tools	fishing equipment	pet	throwing food trash away
bathroom privacy	floor	pet foods	toilet
bathroom/toilet	foodstores	phone	toilet paper
bed	freezer	photo guide images	tools
bed kids	front door	picking up the phone	tooth paste
bed_hq	frontdoor keys	piercings	toothbrush
bedroom	fruit trees	place where eating dinner	toothpaste on toothbrush
bike	fruits and vegetables	place where serving guests	toys
bills of money	get water	plate of food	tractors
boat	glasses or lenses	plates	transport of heavy things
books	go through mail	play area	trash/waste
bowls	goats	playgrounds	turn tv on
bread - ready	grains	playing	turning fan/ac on
bread bowls	guest bed	playing an instrument	turning heater on
bread ready	hair brush/comb	playing with most loved toy	turning lights on and off
brushing hair	hallway	plugging into and out of power outlet	tv
brushing teeth	hand back	portraits	using most loved item
car	hand open to closed	pouring drinking water	using toilet
car decorations	hand palm	pouring water	vegetable markets
car keys	hand washing	power outlet	vegetable plot
carrying water	hanging clothes to dry	preparing food	vegetables
cattle	home	preparing social drink	ventilation
ceiling	horse stables	presenting dollar street	visit
celebrity posters	horses	putting on make up	walking through home
chickens	house overview	radio	walking to get water
children room	how the most loved item is used	reading	walking towards front door
chopping food	icons	reading a book	wall
chopping ingredients	idols	reading light	wall clock
cigarettes	ingredients	refrigerator	wall decoration
cleaning after toilet	instrument	rehabilitation technology	wall inside
cleaning equipment	jewelry	replaced	wardrobe
cleaning floors	kids playing inside	roof	washing clothes/cleaning
closing the front door	kids playing outside	rug	washing detergent
clothes	kitchen	salt	washing hands
coats and jackets	kitchen sink	seeing the back of book	waste dumps
coins	knifes	shampoo	water outlet
computer	latest furniture bought	shaving	water purifier solutions
contraceptives	laying in bed - pretend to sleep	sheep	water sources
cooking	light source in kitchen	shoes	water sources for doing dishes
cooking food	light source in livingroom	shower	wedding photos
cooking pots	light sources	sitting and watching tv	what i wish i could buy
cooking utensils	lightsources by bed	sitting area	wheel barrow
cosmetics	listening to the radio	skies outside	work area
couch	living room	sleeping	worship places
cups/mugs/glasses	lock on front door	smells	worshipping
cutlery	looking over the shoulder	smog/bad air breathing protection	writing
daylight ostructions	make up	smoke and steam exit	writing "home"
diapers (or baby-pants)	markets	smoking	writing signature
dinner guests	meat markets	snack stores	youth culture
disability aid	meat or fish	snacks	
dish racks	meat storages	soap for hands and body	
dish washing brush/cloth	medication	soccer balls	
dish washing soap	menstruation pads / tampax	soccer supporter items	
dishwasher	milk cows or bulls	social drink	
doing dishes	moped/motorcycle	socializing	
drainage	mosquito protection	sofa	
drinking social drink	most loved item	source of cool	
drinking water	most loved toy	source of heat	
drinks	most played songs on the radio	source of light	
drying	music equipment	sources of drinking water	

Appendix B Fuzzy Match

In Section 6.1, we mention that we map the list of topics in our Dollar Street dataset to one of the ImageNet's 1,000 classes. The exact mapping is provided below.

```
{
  "home": "manufactured home",
  "street view": "street sign",
  "tv": "television",
  "washing clothes/cleaning": "washing machine",
  "toilet": "toilet seat",
  "kitchen sink": "washbasin",
  "drinking water": "water bottle",
  "stove/hob": "stove",
  "salt": "salt shaker",
  "bed": "day bed",
  "toys": "toyshop",
  "everyday shoes": "running shoe",
  "plate of food": "plate",
  "cooking pots": "skillet",
  "social drink": "soda bottle",
  "phone": "cellphone",
  "place where eating dinner": "dining table",
  "lock on front door": "padlock",
  "wardrobe": "wardrobe",
  "soap for hands and body": "soap dispenser",
  "ceiling": "tile roof",
  "refrigerator": "refrigerator",
  "bathroom/toilet": "toilet seat",
  "dish washing brush/cloth": "dishrag",
  "toilet paper": "toilet paper",
  "plates": "plate",
  "dish washing soap": "soap dispenser",
  "trash/waste": "trash can",
  "dish racks": "plate rack",
  "shower": "shower curtain",
  "cups/mugs/glasses": "cup",
  "armchair": "rocking chair",
  "light sources": "table lamp",
  "light source in livingroom": "table lamp",
  "books": "bookcase",
  "switch on/off": "switch",
  "light source in kitchen": "table lamp",
  "couch": "studio couch",
  "sofa": "studio couch",
  "roof": "tile roof",
  "cutlery": "wooden spoon",
  "cooking utensils": "spatula",
  "medication": "medicine cabinet",
  "source of cool": "electric fan",
  "pen/pencils": "ballpoint",
  "street detail": "street sign",
  "turning lights on and off": "switch",
  "music equipment": "speaker",
  "tools": "tool kit",
  "cleaning equipment": "dishrag",
  "bed kids": "day bed",
  "table with food": "dining table",
  "get water": "water jug",
```

```
  "paper": "paper towel",
  "radio": "radio",
  "shoes": "running shoe",
  "starting stove": "igniter",
  "freezer": "icebox",
  "source of heat": "space heater",
  "computer": "desktop computer",
  "jewelry": "necklace",
  "knifes": "paper knife",
  "wall clock": "wall clock",
  "pouring water": "water jug",
  "doing dishes": "dishwasher",
  "guest bed": "day bed",
  "mosquito protection": "mosquito net",
  "bike": "all-terrain bike",
  "pouring drinking water": "water bottle",
  "oven": "stove",
  "place where serving guests": "eating place",
  "glasses or lenses": "dark glasses",
  "necklaces": "necklace",
  "source of light": "table lamp",
  "parking lot": "parking meter",
  "waste dumps": "trash can",
  "eating": "restaurant",
  "car": "passenger car",
  "reading light": "table lamp",
  "lightsources by bed": "table lamp",
  "family eating": "eating place",
  "arm watch": "digital watch",
  "taking a teaspoon of salt": "salt shaker",
  "using toilet": "toilet seat",
  "sitting and watching tv": "television",
  "opening and closing the freezer": "icebox",
  "diapers (or baby-pants)": "diaper",
  "moped/motorcycle": "moped",
  "cleaning after toilet": "toilet paper",
  "dishwasher": "dishwasher",
  "opening and closing the refrigerator": "refrigerator",
  "answering the phone": "mobile phone",
  "alarm clock": "analog clock",
  "wheel barrow": "wheelbarrow",
  "listening to the radio": "radio",
  "dinner guests": "eating place"
}
```