San Luis Potosi. Photo: Abraham Pérez

ARTIFICIAL INTELLIGENCE FOR CLIMATE AND DEVELOPMENT

With Recommendations for USAID/Mexico

March 8, 2024

Author: Christopher Yeh

Final Submission Date: March 11, 2024

Contract Number: 7200AA20D00017/72052322F00003

COR: Karely Alvarado

Submitted by:

Rich Fromer, Chief of Party

LINC

1201 Connecticut Ave NW, Ste 200J Washington, DC 20036

Tel: +1 (202) 640 5462 | www.linclocal.org

Email: rfromer@mexicomesa.org

TABLE OF CONTENTS

ACRONYMS

AC	Air Conditioning
AI	Artificial Intelligence
AI2	Allen AI Institute
AV	Autonomous Vehicle
CAIC	Centre for AI and Climate
CCAI	Climate Change AI
CI	Conservation International
CO_2	Carbon Dioxide
DO2	Development Objective 2
EV	Electric Vehicle
GPAI	Global Partnership on AI
GPU	Graphics Processing Units
GHG	Greenhouse Gas
IPCC	Intergovernmental Panel on Climate Change
MESA	USAID/Mexico Monitoring and Evaluation Support for Adaptation Activity
ML	Machine Learning
NASA	National Aeronautics and Space Administration
OSD	USAID/Mexico Office of Sustainable Development
SURGES	USAID/Mexico Southern Mexico Generating Employment and Sustainability Project
TB	Terabyte
TNC	The Nature Conservancy
USAID	United States Agency for International Development
USFS	United States Forest Service
V2G	Vehicle-to-Grid

Introduction

USAID/Mexico's Development Objective 2 (DO2) is to advance transparent and climate-resilient economic growth. Efforts under this objective are centered on capitalizing on sustainable market opportunities to tackle greenhouse gas emissions, advancing climate-smart energy initiatives, and expanding the use of nature-based climate solutions for agriculture, forestry, and land-use. The Mission's Office of Sustainable Development (OSD) currently supports the implementation of six activities directly contributing to DO2: Co2munitario, Sustainable Landscapes Ventures, Sustainable Prosperous Communities, Sustainable Landscapes and Forest Transparency, SURGES,[1] USAID-NREL Partnership, and Partnership for Net Zero Cities.

Having noted the accelerating worldwide use of artificial intelligence (AI) in support of climate change objectives, OSD commissioned the USAID/Mexico Monitoring and Evaluation Support for Adaptation (MESA) Activity to conduct research to identify ways in which AI is already being applied in support of climate priorities, identify opportunities for USAID/Mexico and its activities to leverage AI for better results in current and future OSD programming, and provide additional resources on AI for climate change for the Mission staff's future reference. The report is organized along these lines of inquiry:

- **Section 1** presents a brief primer on AI methods (with a more detailed introduction to AI in Annex A) to help the reader with less experience in applications of AI understand the example uses discussed in Sections 2 & 3.
- **Section 2** details current ways AI models are being applied to the energy, transportation, forestry, and sustainable development sectors globally. At the end of each thematic area, a summary table provides potential considerations for the Mission's programming.
- **Section 3** provides concrete recommendations for how these models may be applicable to current, and to a lesser extent, future OSD programming.
- **Section 4** presents additional resources for OSD to review and engage as appropriate, including a list of relevant upcoming events, notable relevant organizations, and an annotated bibliography.

Section 1: Brief Primer on Artificial Intelligence

Before discussing how artificial intelligence (AI) can be used for climate and development initiatives, this brief primer provides some initial background and terms for the reader. For those interested or who feel that this section does not go into sufficient detail to answer background questions, Annex A presents a more complete AI primer.

Artificial intelligence (AI) is not a single, universal technology, rather it refers to a large collection of technologies with different capabilities, strengths, and weaknesses. The recent growth of AI comes from machine learning (ML), a statistical method that has largely improved upon and replaced the traditional approaches. Today, "AI" and "ML" are almost entirely synonymous, and this report will use these two phrases interchangeably.

Fundamentally, all ML technologies are trying to solve the same problem of fitting a function to data. That is, ML algorithms try to learn a mapping from inputs *x* to outputs *y*. For example, using hourly satellite images from the past 24 hours (input) to predict the hourly weather during the next 24 hours

[1] Southern Mexico Generating Employment and Sustainability Project

(output); having a chatbot respond to a user question or request (input) with any one of many possible answers (output); or having a smart thermostat use forecasts of temperature and energy (inputs) to determine weather to run the heat, air conditioning, or turn off (output).

In essence, practically all AI algorithms boil down to approximating arbitrary mathematical functions $y = f(x)$ that map those inputs to outputs. Given a large enough dataset of input-output examples (x,y), AI algorithms can usually learn a reasonable approximation of this input-output mapping.

The differences between AI technologies can be roughly categorized along four areas: goal of the model, format of data accepted by the AI algorithm, the way the data set is collected and/or designed, and the mathematical methods used.

GOAL OF THE MODEL

There are three common AI model goals:

- **Predictive AI** seeks to learn the single most likely output y given an input x. This is best when there is a correct answer that can be predicted using past examples of correct answers. The weather forecasting example above falls into this category.
- **Generative AI** seeks to generate a likely output y (or possibly several likely outputs) from the distribution of possible outputs given an input x. This goal is appropriate when we want to sample many possible outcomes y for each input x, and there might not be a single "correct answer" but rather a spectrum of possible answers. The chatbot example above falls into this category.
- **Reinforcement learning** seeks to take an optimal action y given an input x. This goal is appropriate when y corresponds to an action, instead of a prediction, and when we do not know the optimal action in advance, but there may be some metric by which to measure the optimality of an action taken. The smart thermostat controller example falls into this category. The AI agent guesses what it thinks is a good action ("turn off system," "turn on AC", or "turn on heater") which results in a change in the environment (change in the room temperature) as well as reward or penalty for the agent (deviation from the desired temperature range, amount of carbon emissions associated with the energy used). This reward signal is used to update the AI agent's decision-making process for similar inputs.

FORMAT OF DATA ACCEPTED BY THE AI ALGORITHM

Currently, there is not a "universal" AI algorithm that can handle all types of data inputs. Instead, individual AI algorithms are usually designed to handle a particular set of data formats or modalities. The different modern subfields of AI have generally developed around the needs of handling different data types, though it is important to note that these domains are not mutually exclusive.

- **Computer vision** refers to the subfield of AI that processes 2-D and 3-D images, videos, and other visual inputs. The weather forecast example uses computer vision with satellite imagery.
- **Natural language processing** refers to the subfield of AI that deals with text, language, and speech data. The chatbot example uses natural language processing to understand text input.
- **Network analysis** and **graph learning** refer to the subfield of AI that deals with data that is best represented via a network (*i.e.*, a "graph" in mathematical parlance). Examples of network data include the electrical grid, social networks, and transportation road networks.
- **Recommender systems** refer to the subfield of AI that processes human preferences. Examples include search engines (*e.g.*, Google) and product recommendations (*e.g.*, Amazon or Netflix).

- **Reinforcement learning** refers to the subfield of AI that processes data from interacting with an environment. The smart thermostat controller example uses reinforcement learning to improve its results by interacting with its environment.

THE WAY THE DATASET IS COLLECTED AND/OR DESIGNED

AI algorithms need to be given large amounts of training data, but perhaps even more important, is the structure of the input-output pairs.

- The most straightforward dataset design is (x,y) input-output pairs, like the weather forecasting example above, which could train on millions of pairs of historical weather observations 24 hours apart to predict future weather within the next 24 hours.
- A chatbot requires a complex dataset design, because large enough datasets for what "good" or "correct" outputs from a chatbot should look like do not exist. Instead, modern chatbots such as ChatGPT rely on three separate training datasets, for training the chatbot to learn different tasks.
    1. First is a large collection of example (*first sentence*, *next sentence*) pairs, from websites such as Wikipedia and The New York Times.
    2. Second is a collection of (*prompt*, *response*) pairs that demonstrate how the chatbot should respond to prompts. This dataset is generated by real humans.
    3. Finally, the chatbot is fine-tuned on a third dataset of (*prompt*, *better response*, *worse response*) triplets to improve its responses. The prompts are generated by real humans, the two responses are generated by the chatbot itself, and the better-vs-worse distinction is decided by human annotators.
- Reinforcement learning, like the smart thermostat controller, uses a different dataset design than the other examples. Humans are not suited to picking the best action y ("turn off system", "turn on AC", or "turn on heater") that simultaneously maximizes thermal comfort while reducing energy usage and associated emissions. Because this objective is complex and difficult to optimize directly, a dataset consists of ($x,\hat{y},r$) triplets, where $\hat{y}$ is a given action taken for the input x, and r is the corresponding measurement of success of that action (also called a reward). Based on this dataset, reinforcement learning algorithms "learn-by-doing" to create a map from an input x to an optimal action y, without ever being told the optimal action for each x.[2]

THE SPECIFIC MATHEMATICAL METHODS USED TO APPROXIMATE THE *X*-TO-*Y* MAPPING

Many different mathematical functions can be used to approximate an x-to-y mapping. One of the simplest mathematical models is linear regression, which is only useful if the x-to-y mapping has a linear structure. For images and text data, though, this linear structure no longer holds, and therefore more advanced mathematical models are necessary. The term "training" refers to adjusting the parameters (also called "weights") of the model to better fit the data. Generally speaking, the more parameters a model has, the better it is able to fit more complicated x-to-y mappings. However, tuning more parameters also has a higher computational cost and takes more time.

TRAINING VS. INFERENCE

The term "training" refers to the act of updating an AI's mathematical model parameters to fit the training data. The term "inference" refers to *using* a trained AI model, that is, running a trained AI model

[2] The dataset variable $\hat{y}$ indicates a potentially suboptimal action taken by the AI agent while learning. The final mapping variable y indicates the optimal action that the AI agent nears over time through reinforcement learning.

on new inputs. Inference is only possible after an AI model has been trained. Training tends to be much more computationally expensive, with most modern AI models requiring graphics processing units (GPUs) to speed up computation. In contrast, inference tends to be much less computationally expensive. Typically, training an AI model on a dataset is usually at least 100-1000x slower than running AI inference on the same inputs.

Section 2: AI for Sustainability and Climate Change

Addressing climate change involves two facets, mitigation and adaptation, both of which AI can help with. Mitigation refers to reducing greenhouse gas (GHG) emissions in the energy sector, transportation, buildings, industry, land use, and agriculture. Adaptation requires better climate modeling (including extreme weather prediction), risk modeling, and planning for resilience and disaster management.

AI FOR ENERGY

Machine learning techniques have long been applied to issues of energy. As early as the mid-1960s, the field of geostatistics developed “kriging” as a method for predicting the locations and sizes of ore and oil deposits based on limited observations (Matheron, 1963); the same method is now known as Gaussian processes regression in the machine learning community and is widely used in areas ranging from climate modeling to materials design.

Currently, the main use of AI for climate impact in the energy sector is for enabling a low- or zero-carbon electricity grid. The main challenge with modern renewable electricity is its variability—solar generation is only available during daytime on sunny days, wind turbines only generate energy when the winds blow, and hydropower can only function with sufficient rainfall collected in dams. Other, less variable renewable energy sources exist—e.g., tidal power and geothermal energy—but these energy sources have yet to be deployed at scale (<1% of U.S. energy capacity as of 2020).[3] In the era of rapidly increasing variable energy deployment, AI can serve three critical roles:

1. Forecasting electricity demand and renewable energy generation, to reduce the reliance on carbon-intensive generators;
2. Controlling smart energy systems to reduce emissions, given accurate electricity demand and generation forecasts;
3. Accelerating the research and development of new materials to increase renewable generation and/or energy storage density.

Demand and renewable generation forecasting

Grid operators have a long history of forecasting electricity demand at many scales, ranging from years-ahead forecasts that dictate capacity markets and weeks- and months-ahead forecasts for planning reserves, to short-term day-ahead and 15-minute forecasts for the day-ahead and real-time markets. The timescale where AI is most useful is in short-term (weeks-ahead to real-time) forecasting because it is where variability is highest and where data is most prevalent. To see why historical data is most readily available for short-term forecasting, note that a single year produces only 12 data points at monthly intervals vs. 35,040 data points at 15-minute intervals (365 days per year * 24 hours per day * 4 fifteen-minute intervals per hour = 35,040 fifteen-minute intervals per year).

---

[3] https://www.energy.gov/eere/geothermal/articles/now-available-iea-2020-us-geothermal-report

Until the rise of variable renewable energy, forecasting efforts were primarily focused on electricity demand; dispatchable generators (mainly fossil-fuel generators and to a large extent, hydropower) had pre-determined power capacities, and it was up to the grid operator/electricity market to dictate which generators would turn on or off. In contrast, solar and wind are non-dispatchable generators—a grid operator cannot dictate that a wind farm produces electricity when the wind is not blowing. Thus, the rise in renewables has been accompanied by a commensurate need for electricity generation forecasting. Nowadays, the role of the grid operator is to use renewable generation forecasts to decide which generators to dispatch to meet the *net* electricity demand, which is the total electricity demand minus the renewable energy generation.

Carbon intensity and carbon emissions forecasting

The abilities to accurately forecast electricity demand and renewable energy generation (and therefore which generators will need to run at each time step in the future) have also led to short-term forecasts of many other grid-related quantities:

- carbon intensity, commonly measured in units of carbon emissions per unit of electrical energy consumed (kg CO_2 / kWh)
- wholesale electricity prices
- curtailment (which is when a renewable power plant produces more electricity than is demanded, leading to effectively zero carbon intensity)

Predicting these quantities is particularly useful for coordinating flexible electrical loads—energy intensive tasks that have some flexibility in when they should be run. Both Google and Microsoft have set net-zero carbon emissions goals, and they use forecasts of carbon intensities and curtailment to decide when to run their data centers harder (often for training large AI models or indexing the internet).[4] For example, if Google forecasts a curtailment of wind energy tonight, it may choose to pause its internet indexing during the day and resume indexing at nighttime to take advantage of zero carbon intensity electricity.

Some other examples of flexible loads that can be intelligently scheduled around carbon intensity forecasts include air conditioning systems and thermostats (*e.g.*, Google Nest), refrigerators, EV charging, and grid-scale battery systems. These are often called "carbon-aware" energy systems, and they commonly employ both predictive AI for forecasting future carbon intensity as well as reinforcement learning algorithms for making optimal scheduling decisions.[5]

A number of entities have already adopted AI-powered technologies for forecasting short-term energy demand, renewable generation, and carbon emissions. Many electric grid system operators in the U.S. (*e.g.*, CAISO[6] and PJM[7]) either have plans to adopt or have already begun adopting AI-generated forecasts. Two notably private (non-governmental) organizations provide real-time carbon intensity data as well as AI-power forecasts for many regions of the world, including Mexico: WattTime[8] and

[4] Google: https://blog.google/inside-google/infrastructure/data-centers-work-harder-sun-shines-wind-blows/. Microsoft: https://devblogs.microsoft.com/sustainable-software/carbon-aware-kubernetes/
[5] https://chrisyeh96.github.io/sustaingym/
[6] https://www.caiso.com/Documents/2022-2026-Strategic-Plan.pdf
[7] https://learn.pjm.com/-/media/about-pjm/newsroom/fact-sheets/load-forecasting-supports-reliability-and-efficiency.ashx
[8] https://watttime.org/docs-dev/coverage-map/

ElectricityMap[9]. In regions where electricity is traded through commodities markets, as in the U.S. and parts of Europe, quantitative finance firms have long-relied on AI forecasts of energy prices to profit from arbitrage. Grid-scale battery systems, such as Tesla's Megapacks,[10] rely on AI technologies for deciding when to charge and discharge their batteries to both help stabilize the grid and profit from price arbitrage.

Battery materials development

Besides adding value to the energy grid, AI technologies have also already facilitated improvements in both "solar fuels" and battery technology. Historically, the process of discovering new materials is slow and tedious, as the fundamental quantum physics underlying chemical reactions are difficult for humans to reason about and extremely computationally expensive and time-consuming to simulate based on first principles. Instead, AI techniques can be applied throughout the materials development pipeline. Some AI models aim to speed up the chemical reaction simulations (*e.g.*, Caltech's OrbNet), whereas other AI models directly aim to predict the properties of chemicals (*e.g.*, Google DeepMind's AlphaFold and IBM Research's MolFormer). Recently, new AI generative models for chemical molecules have demonstrated promising abilities to directly predict chemical compounds that are likely to possess desired chemical properties.

One of the leading efforts in the AI for energy-related materials design space is the Open Catalyst Project by Facebook AI Research and Carnegie Mellon University. The Open Catalyst Project is specifically aimed at finding low-cost chemical catalysts to improve energy storage density and efficiency.

Additional applications of AI in energy

Below are some additional applications of AI in the energy sector for addressing climate change:

- AI-based system for controlling a nuclear fusion reactor[11]
- AI-powered design and placement of wind turbine blades using reinforcement learning or Bayesian optimization[12]
- AI-based suggestions for predictive maintenance in energy distribution pipelines[13] and electricity power lines[14]

[9] https://app.electricitymaps.com/map
[10] https://www.tesla.com/megapack
[11] https://www.wired.com/story/deepmind-ai-nuclear-fusion/
[12] Windscape AI (https://windscape.ai/) is an example of a startup company using AI for wind turbine blade optimization. Bayesian optimization is an AI technique for maximizing an arbitrary "black-box" function, meaning that only the inputs and outputs of the function can be observed, but the function itself is unknown. In the case of wind turbine blade optimization, the "black-box" function is the mapping between the turbine blade angle and the turbine's electricity output. While there may be a highly complex physics-based equation to describe this mapping, such an equation is often too complex to derive. Instead, this mapping is typically treated as a black-box function where input-output pairs can only be observed from historical data. Note that Bayesian optimization, which seeks to find the input that maximizes an unknown function, is different from traditional machine learning, which seeks to find the parameters of a mathematical model to best approximate the unknown function.
[13] https://www.plantengineering.com/articles/how-machine-learning-contributes-to-smarter-pipeline-maintenance/
[14] https://www.technologyrecord.com/article/eon-uses-drones-and-ai-to-inspect-power-lines-virtually

AI for Energy, Considerations for USAID/Mexico:

- USAID is not currently doing any work on grid-level energy generation or management, but if future activities begin to support grid-level efforts, several applications of AI could be relevant:
    - Support AI-based demand and renewable generation forecasting to improve management of energy resources
    - Support AI-based suggestions for predictive maintenance in energy distribution systems
    - Encourage electric utilities and/or energy generators to publicize AI-based forecasts of carbon emissions data to enable carbon-aware demand response from large electricity consumers (*i.e.*, shifting energy usage to low carbon intensity time periods)
- USAID's Net Zero Cities (specifically Objective 1) and NREL activities could potentially encourage the adoption of smart thermostats and other air conditioning equipment that shift energy usage to low carbon intensity time periods
    - Support the creation of "digital twins" of campuses and large facilities to enable AI-based optimization of electricity demand (see recommendation #5)

AI FOR TRANSPORTATION

The transportation sector accounts for about a quarter of global energy-related CO_2 emissions, with nearly two-thirds of transportation-related emissions coming from road travel. To reduce transportation-related emissions, AI methods can help with both increasing the efficiency of existing transportation systems as well as shifting people towards activities or transportation modes that have lower emissions.

AI algorithms help solve two fundamental issues in transportation systems: real-time monitoring and traffic prediction. AI solutions for solving these two issues alone do not directly help with reducing emissions, but they enable a host of downstream applications with stronger climate impacts.

Real-time transportation system monitoring and understanding

Fundamental to understanding any vehicular transportation system is real-time monitoring. Traditionally, vehicle traffic is counted via ground-based inductive loop detectors, but such detectors tend to be costly to install and maintain, and inductive loop sensors also struggle to identify cyclists and pedestrians. Instead, cameras mounted at intersections[15] or on buildings along roads[16] can use AI-based computer vision to detect and track individual vehicles, cyclists, and pedestrians with very high accuracy. However, a barrier to installing cameras is privacy concerns.

In addition to empowering camera-based sensors to count vehicles, cyclists, and pedestrians at specific locations, AI algorithms can use the raw traffic count data along with information about the road network to infer the most likely origin-destination pairs that induced the real-time observed traffic counts. This real-time "dynamic" origin-destination information may be useful, for example, for determining dynamic congestion pricing or for transit agencies and taxi companies to prioritize their operations to serve routes with highest demand.

[15] https://www.transportation.gov/sites/dot.gov/files/2021-04/UTC%20Newsletter_Quarterly_April_Q2_3-18.pdf
[16] https://telraam.net/en/what-is-telraam

Predicting traffic patterns

Perhaps the most prominent use case of AI in transportation is for modeling and predicting traffic patterns. Because roads are connected via a network, network analysis and graph processing AI technologies are well-suited for this task. Google Maps, for example, uses a Graph Neural Network to predict near-term traffic conditions.[17]

Real-time monitoring of transportation systems and the ability to predict traffic patterns in turn enable a number of additional AI-driven transportation applications that are relevant to climate action.

Efficient vehicle route planning

Whether moving passengers or freight, route planning is an essential logistical component in transportation. A prime example is Google Maps' directions, which rely on a combination of real-time traffic and forecasted traffic data to calculate the fastest routes.[18] On-demand ride services such as Uber[19] and Lyft[20] likewise use AI to optimize driver routes and improve arrival time calculations. Delivery companies such as Amazon also use AI technologies (specifically reinforcement learning algorithms) to plan their delivery routes,[21] and AI tools can provide additional value in optimizing the complex interaction between shipment sizes, transportation modes, and origin-destination pairs.

Improved public transit experience

A barrier to increasing passenger adoption of public transit is its unreliability, or at least perceived unreliability.[22] Utilizing AI-powered traffic prediction as well as historical transit data, various mobile apps such as Google Maps[23] and Transit App[24] provide users with real-time predictions of bus arrival times that are more accurate than the official bus schedules, and Transit App can also detect unannounced bus detours.[25]

Shared mobility and transit planning

AI tools that identify long-term origin-destination trends as well as real-time dynamic origin-destination information provide both policymakers and companies appropriate information to improve shared mobility and transit routes to serve the highest demand transportation routes, which may differ at

---

[17] https://blog.google/products/maps/google-maps-101-how-ai-helps-predict-traffic-and-determine-routes/
[18] https://blog.google/products/maps/google-maps-101-how-ai-helps-predict-traffic-and-determine-routes/
[19] https://www.uber.com/blog/deepeta-how-uber-predicts-arrival-times/
[20] https://www.lyft.com/blog/posts/how-lyft-uses-ai-to-get-you-where-you-want-to-go-faster
[21] https://aws.amazon.com/blogs/supply-chain/aws-last-mile-solution-for-faster-delivery-lower-costs-and-a-better-customer-experience/ and https://finance.yahoo.com/news/amazon-using-ai-deliver-packages-154018094.html
[22] See, for example, https://www.metro.net/about/metros-customer-experience-survey-identifies-reliability-frequency-safety-cleanliness-and-homelessness-as-top-improvement-areas-the-annual-survey-informs-cx-action-plan-to-address-custom/
[23] https://blog.research.google/2019/06/predicting-bus-delays-with-machine.html
[24] https://blog.transitapp.com/can-we-make-montreals-buses-more-predictable-no-but-machines-can-e42f28a1a0ba/
[25] https://blog.transitapp.com/transit-can-now-detect-where-your-bus-has-been-detoured/

different times of the day, week, and year. In theory, optimally matching shared mobility services with these routes would increase vehicle sharing and therefore reduce emissions per capita.

Traffic signal optimization

Traditional traffic signal timings typically use either fixed duration cycles (which may follow a time-varying schedule over the course of a day) or simple rules based on car counts from inductive loop sensors embedded into the roadway near intersections. Traffic signals on the same road often operate independently, and even when they are coordinated, the rules are again often simple and sub-optimal.

AI technologies (especially using reinforcement learning) for optimizing traffic signal timings have long been a topic of academic research and have shown significant promise for improving traffic flow. A notable example is Google's "Green Light" initiative which is now deployed in 12 cities around the world and claims up to 30% reduction in traffic stops and 10% reduction in greenhouse gas emissions.[26] Notably, Google's Green Light initiative only uses historical data from Google Maps to suggest changes to existing traffic light timings. A comprehensive AI system that jointly optimizes multiple traffic lights, such as SUMO-RL,[27] may have the potential for more substantial traffic and emissions improvements.

Improved traffic enforcement, enabling better public transit

In addition to informing traffic models, AI-powered cameras and other roadway sensors can also help facilitate faster and more reliable public transit. For example, New York's Metropolitan Transportation Authority (the city's public transit agency) has installed AI-powered cameras onto its buses. These cameras automatically identify illegal parking violations (*e.g.*, cars illegally driving or parked in dedicated bike or bus lanes or at bus stops) and send video footage evidence along with vehicle license plate information to the relevant law enforcement agencies to issue citations.[28] Such citations aim to deter drivers from illegally parking cars in bus lanes and bus stops, thereby allowing buses to travel faster. While the technology is very new, a pilot program in Manhattan, New York, suggested an improvement of 5% in bus speed travel times.[29]

Beyond vehicular traffic, a number of other areas of transportation can be readily improved with AI technologies.

Intelligent electric vehicle (EV) charging and vehicle-to-grid (V2G) operations

As briefly mentioned in the earlier section on AI for Energy, EV charging systems can be designed to operate in a "carbon aware" manner in order to reduce the carbon emissions associated with the electricity needed to charge EVs. An intelligent EV charging system (also called "managed" EV charging) does not always charge the EV at the charger's maximum charging rate. Instead, if it predicts that the grid's carbon intensity will decrease in the near future, it may choose to charge slowly when the grid's carbon intensity is high and wait for the carbon intensity to decrease before charging more quickly. By either directly asking the EV driver for a desired length of the charging session, or by predicting the

---

[26] https://sites.research.google/greenlight/
[27] https://lucasalegre.github.io/sumo-rl/
[28] https://www.hayden.ai/applications/transit-buses
[29] https://www.hayden.ai/press/hayden-ai-completes-first-phase-of-new-york-mta-automated-bus-lane-enforcement-system-expansion

charging session length based on previous interactions with the user, an intelligent EV charging station can thus reduce EV charging carbon emissions while still ensuring that the EV is charged to full capacity by the end of the charging session. Both the forecasting of grid carbon intensity and the optimization of time-varying charging rates can be done with AI techniques. Besides reducing carbon emissions, intelligent EV charging systems can also reduce charging costs by incorporating forecasts of future electricity prices. PowerFlex chargers pioneered intelligent EV charging, with over 10,000 such intelligent chargers installed across the U.S., mostly at workspace locations where vehicles are often parked for multiple hours at a time.[30] Note, though, that the benefits of intelligent EV charging are only realizable if the grid's carbon intensity and/or electricity prices are variable; if they are always constant, then intelligent EV charging provides no environmental benefits.

A potential future iteration of intelligent EV chargers may further include V2G capabilities, which allow an EV that is plugged into a charger to *discharge* energy into the grid from its battery. Instead of charging slowly during periods of high carbon intensity, and assuming that the EV was charged using lower-carbon intensity electricity, a V2G charger could draw energy from the EV battery to supply the grid with cleaner electricity. Sufficient adoption of EVs and intelligent V2G chargers may help lower the cost of a transition to a low carbon grid.[31]

Reducing heat-trapping airplane contrails

While the AI-powered transportation technologies described above focus on ground transportation systems, AI may also help reduce carbon emissions associated with the hard-to-decarbonize aviation industry, which currently accounts for about 2% of global emissions.[32] According to the Intergovernmental Panel on Climate Change (IPCC), a significant fraction (with some estimates up to 57%) of aviation's global warming impact comes from contrails created by airplanes, the thin clouds created in the wake of an airplane's flight path,[33] because contrails formed in the evening prevent heat from radiating into space.

Google Research has developed AI-powered contrail forecast maps that identify humid regions of the atmosphere that are most conducive to contrails. Routing flight paths around these contrail-forming regions has been demonstrated to be a cost effective (<\$25/ton CO_2e) method for reducing global warming.[34]

Finally, AI technologies may enable a number of "moonshot" projects with the potential for high-impact changes in transportation.

Autonomous vehicles (AVs) for platooning and traffic smoothing

Autonomous vehicles enabled by AI technology may help reduce climate impacts of transportation through multiple mechanisms. First, AVs may be programmed to drive more efficiently than humans. Second, a fleet of on-demand AVs could reduce personal car ownership, therefore reducing the carbon emissions associated with car manufacturing. Additionally, AVs can reduce energy consumption by platooning (driving very close together, as if they were connected train cars, to reduce air resistance) by

[30] https://www.powerflex.com/
[31] https://pubs.rsc.org/en/content/articlehtml/2022/ya/d2ya00204c
[32] https://www.iea.org/energy-system/transport/aviation
[33] https://www.ipcc.ch/report/ar6/wg3/downloads/report/IPCC_AR6_WGIII_Chapter10.pdf, p.1087
[34] https://blog.google/technology/ai/ai-airlines-contrails-climate-change/

leveraging communication technologies that allow vehicles to brake and accelerate simultaneously. Platooning is particularly effective for reducing the energy consumption of large trucks, although such technology has only been demonstrated in small scale pilots.[35]

Finally, even a smaller number of AVs on the road can help reduce emissions through traffic smoothing—that is, driving in a manner that aims to reduce the occurrence of stop-and-go traffic jams. The CIRCLES Consortium, in particular, has demonstrated in a large-scale test that a small fleet of AVs can collaboratively help reduce traffic congestion.[36]

Autonomous drone delivery

Last mile delivery of goods is estimated to contribute up to 50% of the carbon emissions associated with e-commerce delivery.[37] While intelligent vehicle routing can help reduce emissions, an alternative solution is to use AI-powered autonomous drones for delivery. Amazon, Google's Wing division, and Zipline operate drone delivery businesses, with programs in the U.S., Australia, Africa, and Japan.[38] If the drones are electric powered and charged with renewable electricity, they may significantly reduce last-mile delivery emissions.

As the examples above suggest, AI technologies can help decarbonize the transportation industry, often by increasing the energy efficiency of transportation. However, there is a risk of Jevons paradox—*i.e.*, increasing energy efficiency of transportation might lead to more demand for transportation and therefore more vehicle miles traveled, erasing the total carbon emissions benefits.

AI for Transportation, Considerations for USAID/Mexico:

- USAID's Net Zero Cities (specifically Objective 2) and NREL activities could potentially encourage use of AI to accelerate the decarbonization of transportation (see recommendation #4)
- USAID is not currently supporting value chain producers and intermediaries with distribution related improvements, but if future activities or interventions reach to transportation of goods, AI applications could be further investigated:
  - Efficient delivery vehicle route planning
  - Optimizing the complex relationships among shipment size, delivery methods, traffic patterns, and emissions

AI FOR FORESTRY

In the field of forestry, the various applications of AI are driven by the data modalities that are available. Remote sensing imagery from satellites, planes, and drones may include multispectral daytime and nighttime imagery, LIDAR, and radar measurements. On-the-ground sensors may include camera traps, acoustic sensors, tree inventories, and environmental DNA sampling.

---

[35] https://highways.dot.gov/research/laboratories/saxton-transportation-operations-laboratory/Truck-Platooning

[36] https://news.berkeley.edu/2022/11/22/massive-traffic-experiment-pits-machine-learning-against-phantom-jams/

[37] https://clean-mobility.org/wp-content/uploads/2022/07/Secret-Emissions-of-E-Commerce.pdf

[38] Amazon: https://www.amazon.com/gp/help/customer/display.html?nodeId=T3jxhuvPfQ629BOIL4, Google Wing: https://wing.com/, Zipline: https://www.flyzipline.com/

Satellite imagery and remote sensing

In particular, satellite-based remote sensing stands out for its ubiquity. NASA's Landsat satellites take a complete image of Earth approximately once every 16 days at a resolution of 30 meters per pixel (30m/px),[39] and its geostationary GOES-16/18 satellites image all of North America once every 10 minutes at a resolution of up to 2,000m/px. These satellite images are free for the public to access. Higher resolution satellite imagery can also be purchased from private companies; for example, Planet provides 0.5m/px imagery with a 5-day revisit period.[40] Furthermore, many of these satellites take measurements not only in the visible light spectrum, but also in near-infrared and infrared (thermal) bands that are particularly useful for cloud and moisture detection.

While the ubiquity of satellite imagery promises to enable near real-time global ecosystem monitoring, the large volume of remote sensing data poses significant computational and storage challenges. For example, NASA's Landsat satellites alone generate around 1 terabyte (TB) of imagery data per day.[41] Processing the large amounts of remote sensing data is generally infeasible on personal computers. Instead, remote sensing computations are increasingly cloud-based. Two notable cloud-based tools, Google's Earth Engine[42] (launched in 2010) and Microsoft's Planetary Computer[43] (launched in 2022), provide both data access and computing resources on a unified platform and have enabled a number of diverse applications.

Deforestation tracking

Identifying deforestation in satellite imagery using AI was the original use case that motivated Google's development of Earth Engine. Specifically, Google developed both the AI software along with the necessary computational hardware for detecting illegal logging in the Amazon rainforest in near real-time, the results of which could be reported to relevant law enforcement authorities.[44] Notably, whereas legal logging tends to follow more structured patterns, illegal logging tends to produce uneven or jagged cut patterns that are relatively easy for AI to recognize in satellite images.

Alternatively, organizations such as Rainforest Connection rely on ground-based acoustic sensors to detect potential illegal logging activities.[45] AI-powered microphones connected to solar panels can detect the sounds of chainsaws and logging vehicles in real-time. When deployed in protected forests, these systems have been effective at thwarting illegal logging activities.

Carbon stock estimation and verifying carbon credits

Estimating the amount of carbon stock in a forest is useful for both assessing the amount being sequestered as well as pricing the carbon, should it be sold as part of a carbon offset program. Some AI approaches directly estimate carbon stock within each parcel of land from satellite imagery, whereas other approaches first try to estimate tree canopy height and canopy cover, and then approximate forest

---

[39] https://www.usgs.gov/landsat-missions/landsat-8
[40] https://www.planet.com/
[41] https://landsat.gsfc.nasa.gov/article/imaging-the-past/
[42] https://earthengine.google.com/
[43] https://planetarycomputer.microsoft.com/
[44] https://googleblog.blogspot.com/2009/12/seeing-forest-through-cloud.html
[45] https://rfcx.org/

carbon stock as a function of tree canopy height. Recent work has shown how satellite imagery can be useful for measuring compliance in a payment-for-forest-protection program.[46] A notable example is GainForest, a Swiss non-profit that has pioneered an AI-based "proof-of-care" system that uses remote sensing imagery as well as field data to verify the achievement of conservation milestones before distributing carbon/conservation credits to communities.[47]

Wildfire monitoring and forest fire management

AI-powered climate models, especially for drought forecasting, are particularly useful for identifying regions with the highest fire risk. AI models can also predict the hypothetical spatial progression of a wildfire across different terrains.

Whereas satellite image-based systems are useful for *tracking* larger fires, ground-based video camera systems are often preferred for initial fire *detection*. Because video cameras effectively capture multiple images every second, various computer vision models for ground-mounted video cameras have demonstrated fire detection capabilities within the first minute of fire ignition. A number of companies have now commercialized such AI tools, with software products targeted for government fire management organizations as well as utility companies to improve wildfire response.[48]

For active wildfires, computer vision tools can also accurately track fire boundaries from satellite imagery. Geostationary satellites (*e.g.*, NASA's GOES-16/18) that provide frequent imagery updates every few minutes allow near real-time fire tracking. Google has built such automated wildfire detection and tracking into Google Maps in order to provide users with near real-time information on fire extent.[49] Because visible light does not easily penetrate clouds and smoke, these methods often rely on the infrared (thermal) bands in the satellite imagery to identify the fire location.

Biodiversity monitoring

AI methods have proven to be highly scalable and effective at monitoring wildlife via camera traps, microphones, and citizen science efforts.[50] Camera traps are cameras that are placed in the wild, often mounted to trees, to take photos at regular intervals (*e.g.*, every 15 seconds) or when triggered by motion. Traditional wildlife monitoring relies on humans to manually review thousands of camera trap images taken each month to detect and track wildlife populations. Alternatively, computer vision tools can automatically recognize animals in images and determine their species with very high accuracy, typically reducing the amount of time and cost for processing images by a factor between 2-10x, compared to manual human annotators.[51]

Just as microphones can be used to detect logging sounds, microphones can also detect different animal species within a forest. However, microphones are only useful for detecting species that make sounds, such as birds and mammals.

---

[46] https://www.science.org/doi/10.1126/science.aan0568
[47] https://gainforest.earth/
[48] Examples of commercially available AI-powered real-time fire detection from ground-based camera systems include FireScout AI (https://firescout.ai/), exci (https://www.exci.ai/), and DigitalPath (https://www.nytimes.com/2023/08/24/us/wildfires-ai-detection-california.html)
[49] https://blog.research.google/2023/02/real-time-tracking-of-wildfire.html
[50] https://dl.acm.org/doi/full/10.1145/3466857
[51] https://www.nature.com/articles/s41467-022-27980-y

Identify and intercept illegal poaching

Large wildlife parks around the world often may not have sufficient rangers to consistently monitor every region of the park for illegal poaching activity. To alleviate this issue, Harvard and Carnegie Mellon University researchers have trained reinforcement learning algorithms to predict illegal poaching activity and recommend proactive patrols to deter poachers.[52] These reinforcement learning algorithms recommend patrol routes that optimally balance the objectives of exploration and exploitation—that is, "exploration" to search for poachers across all regions of the park, and "exploitation" to focus on regions of the park that have previously been identified as poaching hot spots.

AI for Forestry, Considerations for USAID/Mexico:

- USFS is actively using AI models to support its deforestation prevention efforts and carbon stock estimation:
    - Computer vision AI to process satellite images to generate deforestation warnings
    - Predictive algorithm to identify potential wildfire hotspots using satellite images and atmospheric variables
    - Predictive AI model to estimate stored carbon using data from the National Forestry and Soil Inventory (INFyS) of the National Forestry Commission (CONAFOR) and remote sensing data
- Other USAID/Mexico activities that work to reduce and prevent deforestation (e.g., SURGES, CI) can also use AI-based tools (see recommendation #3):
    - Remote sensing for tracking deforestation and reforestation over time
    - Satellite and ground-based sensors to identify likely illegal logging to support conservation enforcement
- USAID/Mexico activities like Pronatura that also support communities to access carbon credits can use AI-based remote sensing to track carbon stocks over time (see recommendation #2)
- USAID/Mexico activities that support biodiversity (*e.g.*, the USFS activity's BIOCOMUNI biodiversity management system) can use AI to support biodiversity monitoring:
    - Remote sensing to monitor wildlife populations and locations
    - Predictive algorithms to identify human threats to biodiversity such as illegal poaching or logging

AI for Sustainable Development

AI tools, particularly for processing satellite imagery, hold significant promise for monitoring and informing sustainable development outcomes in the face of climate change. These outcomes, whether in economics, agriculture, or health, tend to rely on manual data collection which is both time-consuming and costly. AI methods provide a complementary means of understanding these outcomes on a large scale, filling in data gaps from manually collected sources.

AI for mapping economic livelihood

As the effects of climate change disproportionately affect poor and under-resourced communities, accurate data on the economic livelihoods of these communities is important for policymakers to

[52] https://seas.harvard.edu/news/2020/06/preventing-poaching

provide appropriate humanitarian aid and support, especially after major natural disasters or other adverse conditions. Since 2016, a rapidly growing body of research has demonstrated the ability of computer vision methods that use a combination of daytime and nighttime satellite images to create high resolution "poverty maps" particularly for regions of the world where survey data is scarce.[53] The AI methods can also identify changes in economic livelihoods over time, thereby providing a low-cost and scalable means of evaluating the success of economic policies.[54] Such AI-generated poverty maps helped inform the deployment of cash transfer programs in urban Africa, for example, in the midst of the COVID-19 pandemic.[55] Another recent study demonstrated a combination of AI and causal inference methods to quantitatively measure the livelihood impact of rural electrification policies.[56]

AI in support of agricultural production

Computer vision algorithms have also been used for forecasting crop yields, even for smallholder agriculture where field sizes may be 0.1 hectares or less.[57] The multispectral bands from satellite imagery, particularly in the near-infrared and visible red spectrum, provide a strong signal about soil health, vegetation density, and plant health.[58] As these forecasts can be made early in each growing season, they can inform policies around food security and humanitarian aid when the forecasts suggest abnormally low agricultural productivity.

For larger-scale farms, AI-powered robots promise to enable "precision agriculture." A major success has been in autonomous weeding robots that roam through crop fields to detect weeds using computer vision and then kill the weeds, whether by firing high-powered laser beams[59] or using a metallic robotic arm to pluck weeds from the soil.[60] These robots help reduce the amount of herbicides needed, improving the health of crop plants, the soil, and other fauna in the vicinity of farms. Other agricultural robots use computer vision to provide plant health analytics and phenotyping on the level of individual plants, which may accelerate crop trials and improve agricultural science for more robust crop yields.[61]

AI for pollution and methane emissions detection

Pairing AI with remote sensing has also enabled highly scalable methods for detecting pollution. A recent study produced daily wildfire smoke PM2.5[62] estimates for the entire continental U.S. at spatial resolution of 10km/px, providing the first large-scale spatial and temporal analysis of wildfire smoke, which is increasingly common under climate change.[63] Furthermore, in March 2024, a new satellite from the Environmental Defense Fund, MethaneSAT, was launched to specifically measure methane

---

[53] https://www.science.org/doi/10.1126/science.aaf7894
[54] https://www.nature.com/articles/s41467-020-16185-w
[55] https://openknowledge.worldbank.org/handle/10986/35003
[56] https://www.nature.com/articles/s41586-022-05322-8
[57] https://sustainlab-group.github.io/sustainbench/
[58] https://www.usgs.gov/landsat-missions/landsat-normalized-difference-vegetation-index
[59] https://carbonrobotics.com/
[60] https://www.ekobot.se/products/ekobot-weai/
[61] https://mineral.ai/
[62] PM2.5 refers to fine particulate matter that is 2.5 microns in diameter or less. PM2.5 is a standard metric for quantifying unhealthful air pollution. See, *e.g.*, https://pubmed.ncbi.nlm.nih.gov/35742761/.
[63] https://pubs.acs.org/doi/10.1021/acs.est.2c02934

emissions, with Google providing compute resources to enable AI-based analyses.[64] As methane is a more potent greenhouse gas than carbon dioxide, identifying methane emission sources and especially unexpected methane leaks is critical to combating global warming.

AI for Sustainable Development, Considerations for USAID/Mexico:

- USAID/Mexico is not currently focusing efforts on pollution or emissions detections, but if future activities target pollution or have the opportunity to address unexpected emissions, AI tools could be applied:
    - Remote sensing to monitor pollution levels from fires, vehicle emissions, etc. could bring new understanding to associated risks and inform the public
    - Predictive algorithms could complement remote sensing to mitigate polluting vehicles or industry before pollution levels reach highs
    - Remote sensing could allow officials to identify emissions of methane and address the sources quickly
- Across its Mission-wide portfolio, USAID seeks to design activities to support vulnerable communities throughout Mexico. AI tools could support identification of those communities through remote sensing and predictive algorithms.
- USAID Activities that are working to improve livelihoods, specifically for small-holder farmers, can use AI-based tools, including remote sensing for tracking soil health and predicting crop yields

Section 3: Recommendations and Priority Areas for USAID/Mexico

1. Identify and recognize the challenges of access to computational resources unique to each AI application.

There are two major criteria that dictate computational resources necessary for each AI application: data location, and training vs. inference.

Data location: cloud vs. local vs. edge

Where an AI model is run—either in the cloud, on local computers, or on edge devices (*e.g.*, smartphones, drones, cameras)—is entirely dependent on where the data is stored. Many factors affect where data can be stored:

- Dataset size: How large is the dataset? How quickly is new data being acquired, if at all?
    - Cloud storage is especially preferred when new data is regularly acquired, because cloud storage automatically scales to the size of your dataset. So-called "Automated ML" (AutoML) tools are also provided by all of the large cloud services, allowing for non-experts to take advantage of the latest AI models, with little or even no code necessary.
    - Local storage is typically cheaper than cloud storage, but requires more expertise to set up. If the dataset is rarely updated (*e.g.*, at most once per year), then processing the

64 https://blog.google/outreach-initiatives/sustainability/how-satellites-algorithms-and-ai-can-help-map-and-trace-methane-sources/

    - data using AI on a local computer may be sufficient, assuming that the dataset is small enough (typically ≤1TB) to fit on a local computer.
    - Edge devices are generally unable to store large datasets, but they may be used for data gathering, and they may transmit the data to a local computer or the cloud for further processing.
- Privacy or security: Are there security considerations about transmitting the data to the cloud? Are there privacy restrictions about who may access the data?
    - Both cloud and local storage, when properly set up, is very secure. However, setting up cloud infrastructure with the necessary security and privacy settings may be particularly challenging for non-experts.
    - Edge devices may be preferred if there are security or privacy considerations regarding data access. In particular, if data is not allowed to leave edge devices due to privacy or security concerns, yet AI models should be trained on such data, then federated learning techniques become necessary.[65]
- Collaboration needs: Does the data and/or AI model need to be accessible by multiple parties?
    - Cloud storage is the easiest for data sharing, whereas local and edge devices are less suitable for large collaborations.
- Internet access: In areas where internet access is limited, running AI models directly on edge devices may be the only feasible option. However, edge devices come with their own limitations, primarily limited data storage and compute capabilities. These are described in the next section.

AI training vs. inference

Another key consideration for compute is whether an AI model needs to be continuously updated, an AI model needs to only be trained once and then deployed, or if an existing pre-trained AI model is already sufficient for the task at hand. As explained in the primer section of this report, training an AI model on new data is significantly much more expensive in terms of both monetary cost and time needed compared to “inference,” which refers to running an already trained AI model. Typically, edge devices such as smartphones do not have the requisite computational resources to train AI models, nor large enough storage capacity to store data sets needed for training AI models. However, edge devices may be used for data gathering, and they may transmit the data to a local computer or the cloud for training AI models.

2. Leverage AI-based remote sensing tools for monitoring forest carbon stocks and verifying carbon credits.

USAID/Mexico activities like Pronatura México support the generation and selling of carbon credits. Whereas Pronatura México is concerned with following the Climate Action Reserve Protocol to generate carbon credits for forest conservation, other activities may be more concerned with carbon credits generated by smallholder farmers and producers for land-use conservation. Under the Climate Action

[65] Federated learning is a technique for training a combined model by learning from multiple decentralized, distributed edge clients. When combined with other privacy-preserving AI techniques such as differential privacy, federated learning enables AI training with minimal security and privacy concerns. See, *e.g.*, https://www.mdpi.com/2076-3417/12/18/9124.

Reserve Protocol, both forest conservation[66] and land-use conservation[67] may, in-part, be verified through high-resolution satellite imagery. (While the Climate Action Reserve does not currently have a land-use protocol for Mexico, its land-use conservation protocols for U.S. and Canada grasslands both allow for satellite imagery-based verification.) Because manually estimating carbon stocks and land-use change from high resolution satellite imagery may be very time-consuming, AI tools may help automate this process.

As described in the AI for Forestry section above, AI-based remote sensing tools can provide real-time insights into forest canopy cover, land-use, and carbon stocks. USFS currently uses an ensemble of over 40 AI models to more accurately predict carbon stocks using data from the Inventario Nacional forestal y de Suelos (INFyS) from the Comisión Nacional Forestal (CONAFOR) and remote sensing data of climatic, vegetation, and topographic variables. As both the satellite imagery (typically from NASA Landsat or ESA Sentinel-2 satellites) and computation resources (from Google's Earth Engine or Microsoft Planetary Computer) are generally free of charge for nonprofit organizations and most government agencies, adopting AI tools for remote sensing applications could be a low-cost and scalable solution for other USAID/Mexico activities working in forestry and agricultural monitoring. Note, though, that some carbon credit protocols require the use of very high resolution (*e.g.*, ≤3m/px) satellite imagery, which is generally not freely available and is only accessible through commercial satellite providers such as Planet Labs or DigitalGlobe. However, such imagery can be imported into Google Earth Engine for processing.[68] More integrated AI-based remote-sensing carbon credit verification tools such as GainForest[69] are under development, though it is unclear whether these tools are directly suitable for the needs of USAID/Mexico activities.

3. Use AI tools for deforestation detection.

Many USAID/Mexico projects (*e.g.*, USFS, Pronatura México, TNC, CI, and SURGES) are concerned with deforestation. As described in the AI for Forestry section above, AI tools can help with deforestation detection. There are generally two types of AI-based tools (satellite imagery or auditory sensors) for deforestation detection, with different strengths and weaknesses.

Satellite imagery tools can detect deforestation in near real-time, with latency dependent on imagery availability. Free, publicly available imagery is generally lower resolution with a slower revisit period (*e.g.*, Sentinel-2 has 10m/px resolution at a 5-day revisit), whereas Planet Labs provides 3m/px resolution at a daily revisit rate. A key benefit of remote-sensing tools is their low startup cost; they do not require the deployment of specialized sensors or equipment. This is particularly useful for monitoring areas without internet or electricity access. Furthermore, as mentioned in the previous recommendation, the satellite imagery data and computational resources may be freely available. USFS is currently using this type of AI-based satellite imagery processing in support of its deforestation warning system.

Ground sensors, especially microphones, can detect deforestation in real-time, provided that they have internet connectivity. Electricity access is not necessary, as these microphones can be powered by solar

[66] https://www.climateactionreserve.org/wp-content/uploads/2023/12/Mexico-Forest-Protocol-V3.0_ENG_Errata-and-Clarifications.pdf
[67] https://www.climateactionreserve.org/wp-content/uploads/2020/02/Grassland_Protocol_V2.1.pdf
[68] https://developers.planet.com/docs/integrations/gee/
[69] https://gainforest.earth/

panels. The two key benefits for ground-sensors is that they can truly detect deforestation sounds in real-time, and consequently, they provide an exact timestamp for the deforestation activity, which may be useful for law enforcement to act on this information. However, ground sensors may take significantly more effort to install and calibrate, especially in more remote forest regions.

4. Use AI tools to accelerate the decarbonization of transportation.

USAID's Net Zero Cities (specifically Objective 2) and NREL activities could potentially encourage use of AI to accelerate the decarbonization of transportation. Net Zero Cities Intermediate Result 2.1 specifically aims to implement mobility and planning strategies to encourage the use of public transit and cycling over personal vehicle ownership. AI tools may assist with the planning process by accurately identifying the most important origin-destination pairs for public transit to serve. Additionally, AI-powered cameras installed on buses and roadways may help law enforcement issue citations to deter vehicles from illegally parking or otherwise obstructing public transit and cycling infrastructure. Third, AI can help accurately predict public transit arrival and departure times based on traffic patterns to encourage public trust in public transit. NREL's collaboration with the state of Yucatán to introduce an electric bus corridor may also benefit from these same AI technologies.

Net Zero Cities Intermediate Result 2.2 aims to reduce emissions from vehicle operations, including the acceleration of electric vehicle fleet adoption. For this task, intelligent AI-powered electric vehicle charging stations may help reduce carbon emissions (as described in the AI for Transportation section), especially if the electric grid's carbon intensity is variable.

Finally, while the Net Zero Cities and NREL activity reports do not explicitly mention "traffic," USAID could consider supporting efforts by municipalities, traffic officials, or public transit authorities to use AI for traffic light optimization to reduce congestion as well as traffic system monitoring to support congestion pricing for car use.

5. Prioritize digital twins and emissions forecasting to enable actionable signals towards the net-zero building goal.

Among the objectives that the USAID/Mexico Net Zero Cities and NREL activities are concerned with is reducing emissions from buildings. One particularly effective technique is to first accurately forecast the carbon intensity of the grid, and then control energy systems (*e.g.*, air conditioning systems, EV charging, water heaters, batteries, computers, heat pumps, etc.) to shift usage towards periods of lower intensity.

The first step, forecasting grid carbon intensity, is mostly a solved problem. As mentioned in the AI for Energy section, real-time grid carbon intensity data and AI-powered forecasts are already available for Mexico through two providers: WattTime[70] (an American nonprofit) and Electricity Maps[71] (a European startup). Grid utilities may already provide such information and forecasts as well.

The second step, taking advantage of the forecasted carbon intensity signal, is more challenging. While individual smart devices, such as Nest smart thermostats[72], can already automatically leverage carbon intensity forecasts, they generally operate individually without coordinating with other energy systems.

[70] https://watttime.org/
[71] https://www.electricitymaps.com/
[72] https://nestrenew.google.com/welcome/

Especially for larger buildings and campuses with many different energy systems, it may be worth creating a “digital twin” of the building or campus. A “digital twin” of a physical system refers to an accurate digital simulation of the system. It is a detailed digital model that mimics how the real system behaves. Increasingly common in the energy and transportation domains, digital twins are used to simulate complex systems like power grids, traffic networks, or air conditioning units in buildings. A digital twin can be thought of as a virtual playground or testbed for trying out new controllers or seeing how a system might respond to unusual inputs, without any risk to the actual physical system itself. For example, a digital twin of a power grid could be used to simulate how the grid will perform under different conditions, like extreme weather or high demand. This helps the grid operator plan ahead and make necessary adjustments to harden the grid against future disruptions.

Creating accurate digital twins of physical infrastructure can be a time-consuming and costly process. Therefore, they are best suited for larger projects (*e.g.*, large building complexes and campuses) where there is more room for emissions reductions and cost-savings from the AI optimizations enabled by a digital twin. Several technology firms (*e.g.*, Bosch[73] and Siemens[74]) already provide digital services for creating digital twins.

For AI applications, a key benefit of digital twins is the ability to synthetically generate large datasets without relying on gathering field data. (Of course, a small amount of field data is often necessary for validating the faithfulness of a digital twin compared to the physical system.) Specifically, digital twins help generate counterfactual data—*i.e.*, “what if” scenarios that differ from actual observations.

Combining digital twin simulations with accurate AI-powered grid emissions forecasting enables “carbon-aware” energy and transportation applications, as described in the AI for Energy section above. Optimizing existing systems to shift workloads and energy consumption from periods of high carbon intensity on the grid to periods of lower carbon intensity provides a low-cost method for reducing carbon emissions, especially during this energy transition period where fossil fuels remain the dominant source of energy on the grid.

Section 4: Additional Resources

RELEVANT EVENTS

Tackling Climate Change with Machine Learning Workshop Series

The “Tackling Climate Change with Machine Learning” workshop series is hosted 2-3 times each year by Climate Change AI (CCAI) at the largest annual international AI research conferences. Each workshop features invited guest speakers, panel discussions, spotlight talks for the best accepted research paper submissions, as well as poster sessions for the remaining accepted paper submissions. Attendees usually tend to be members of the AI and climate change research community, although participants from industry and government are also welcome. Typical attendance is in the 200-500 person range. Links to full recordings of previous workshops can be found on the event website for each of the previous events: https://www.climatechange.ai/events.

---

[73] https://www.boschbuildingsolutions.com/xc/en/news-and-stories/digital-building-twins/

[74] https://www.siemens.com/global/en/products/buildings/digital-building-lifecycle/building-twin.html

The next workshop in this series takes place at the annual International Conference on Learning Representations.

- Date: May 11, 2024
- Location: Vienna, Austria, and virtual (online)
- Website: https://www.climatechange.ai/events/iclr2024

Google Geo for Good Summits

Google Earth Outreach hosts an annual Geo for Good Summit in person at the Google Bay View Event Center in Mountain View, California. The summit teaches practical tools, especially the use of Google Earth Engine and Google Cloud, for collecting, hosting, analyzing, visualizing, and publishing map data. Historically, the speakers and sessions have had a strong focus on leveraging AI for climate and environmental sustainability applications. Recordings of sessions from previous summits can be found at https://earthoutreachonair.withgoogle.com/events/geoforgood23.

The next Google Geo for Good Summit has not yet been announced. Historically, applications to attend open in May, while the summit itself is held in October.

ACM COMPASS

The ACM SIGCAS/SIGCHI Conference on Computing and Sustainable Societies ("ACM COMPASS" for short) is an annual conference hosted by the Association for Computing Machinery (ACM) Special Interest Group on Computing and Society (SIGCAS) and the ACM Special Interest Group in Computer Human Interaction (SIGCHI) since 2018. The ACM is the world's largest and preeminent society for scientific and educational computing, and COMPASS is its annual forum focused on the intersection of computing and sustainability. COMPASS is a full academic research conference with research talks, poster sessions, and workshops. Research presented at COMPASS is not limited to AI-specific applications, but a significant proportion of accepted papers each year leverage AI in at least some capacity.

The 2024 ACM COMPASS conference details are as follows:

- Date: July 8-11, 2024
- Location: Indraprastha Institute of Information Technology Delhi (IIIT-Delhi), Delhi, India
- Website: https://compass.acm.org/

NOTABLE RELEVANT ORGANIZATIONS

Climate Change AI (CCAI)

Climate Change AI (https://www.climatechange.ai/) is a non-profit focused on tackling climate change using machine learning tools. Based in the U.S. but with a global presence, CCAI is primarily run by volunteers across academia and industry. One of the primary CCAI activities is its "Tackling Climate Change with Machine Learning" workshop series described above. CCAI also issues a monthly newsletter highlighting the latest research on AI for climate applications, relevant events, and job postings.

Centre for AI and Climate (CAIC)

The Centre for AI and Climate (https://www.c-ai-c.org/) is a non-profit based in London, U.K., focused on building tools for applying machine learning to accelerate climate action. Notably, CAIC helped lead the development of the report "Climate Change and AI: Recommendations for Government Action" for the

Global Partnership on AI (GPAI), which serves as a guidebook for how governments can facilitate the responsible application of AI to address the climate crisis.

Allen Institute for AI (AI2)

The Allen Institute for AI (https://allenai.org/), founded by Microsoft co-founder and philanthropist Paul G. Allen, is a non-profit research institute focused on high-impact AI research and engineering. A key initiative at AI2 is "AI for the Environment" which has included AI projects ranging from climate modeling to wildlife monitoring and conservation.

ANNOTATED BIBLIOGRAPHY

Burke, Marshall, et al. "Using Satellite Imagery to Understand and Promote Sustainable Development." *Science*, vol. 371, no. 6535, Mar. 2021. https://doi.org/10.1126/science.abe8628.

Professor Burke's paper is a comprehensive review of the promises and challenges of pairing AI techniques with satellite imagery to understand and promote sustainable development. The paper has four key messages. First, AI models using satellite imagery inputs are strongly performant at predicting key sustainable development outcomes. Second, the key bottleneck to better AI models is labeled training data. Third, satellite-based approaches complement, but should not fully replace, ground-based data collection. Finally, adoption of satellite-based approaches in public-sector decision-making has so-far been limited but should be increased. This report specifically highlights applications in agriculture, fisheries, health, and economics.

Clutton-Brock, Peter, et al. *Climate Change and AI: Recommendations for Government Action*. Global Partnership on AI, Nov. 2021, https://gpai.ai/projects/responsible-ai/environment/climate-change-and-ai.pdf.

Formed in 2020, the Global Partnership on AI (GPAI) is an international organization hosted by the OECD that brings together subject matter experts to guide the responsible development in AI while seeking to address the UN Sustainable Development Goals. GPAI's 2021 report on Climate Change and AI provides "actionable recommendations as to how governments can support the responsible use of AI in the context of climate change." These recommendations encompass three main categories: using AI for climate change mitigation and adaptation; reducing the possible negative impacts of AI on the climate; and enhancing implementation and governance capabilities.

The recommendations in the GPAI report are primarily targeted at government policymakers. It includes suggestions on data access standards, research funding for AI technologies, carbon pricing policies that account for emissions from AI computations, and criterion for assessing the climate impact of AI technologies. Finally, the report also includes 13 specific case studies on AI for climate impact, ranging from using AI to estimate global carbon emissions from satellite imagery and ground sensors, to reducing energy usage in Google's data centers.

Dannouni, Amane, et al. *Accelerating Climate Action with AI*. Boston Consulting Group, https://www.gstatic.com/gumdrop/sustainability/accelerating-climate-action-ai.pdf.

This report, commissioned by Google and prepared by the Boston Consulting Group in time for COP28, focuses on three topics. First, it provides a number of case studies where AI can help with climate action, such as using AI to measure CO_2 emissions or generating heat maps of flood risk. It also identifies climate risks associated with AI technologies, especially the large amounts of electricity and water demanded by data centers. Finally, it provides a set of policy recommendations aimed at encouraging more uses of AI for climate action while reducing the climate risks of deploying AI technology. A very short summary of this report is available on the Google Blog: https://blog.google/outreach-initiatives/sustainability/report-ai-sustainability-google-cop28/.

Gomes, Carla, et al. "Computational Sustainability: Computing for a Better World and a Sustainable Future." *Communications of the ACM*, vol. 62, no. 9, Aug. 2019, pp. 56–65. *ACM Digital Library*, https://doi.org/10.1145/3339399.

Professor Gomes, who coined the phrase "computational sustainability," defines it to mean the research field to "identify, formalize, and provide solutions to computational problems concerning the balancing of environmental, economic, and societal needs for a sustainable future." The article is not limited to AI technologies nor climate change applications, but instead considers computational advancements at large as well as the broader set of U.N. Sustainable Development Goals, of which Climate Action is one such goal. This report is useful for providing a more general outlook on the intersection of computational technology and sustainability, especially how many computational technologies were specifically developed for, or at least partially inspired by, sustainability challenges over the past two to three decades.

Rolnick, David, et al. "Tackling Climate Change with Machine Learning." ACM Computing Surveys, vol. 55, no. 2, Feb. 2022, p. 42:1-42:96. *ACM Digital Library*, https://doi.org/10.1145/3485128.

This survey article, first published as a preprint in 2019, was one of the first papers to systematically categorize the numerous applications of AI technologies for climate change adaptation and mitigation. Many examples in this report come from Professor Rolnick's paper. In particular, Table 1 in the article provides a comprehensive cross-map between climate change solution domains and selected areas of AI research that are relevant to each domain, and an interactive version is available online at https://www.climatechange.ai/summaries. Although the article uses a significant amount of terminology specific to machine learning, its intended audience also includes corporate leaders and governments.

ANNEX A. FULL PRIMER ON ARTIFICIAL INTELLIGENCE

Before discussing how artificial intelligence (AI) can be used for climate and development initiatives, it is important to understand that AI is not a single, universal technology. AI methods can loosely be categorized as either symbolic (using predefined rules and logic to deduce results) or statistical (using induction from data). The recent growth of AI comes from machine learning (ML), a statistical method that has largely improved upon and replaced the traditional symbolic approaches. Today, "AI" and "ML" are almost entirely synonymous, and this report will use these two phrases interchangeably.

ML refers to a large collection of technologies, each with different capabilities, strengths, and weaknesses. Fundamentally, all of these technologies are trying to solve the same problem of fitting a function to data. That is, ML algorithms try to learn a mapping from inputs x to outputs y.

This report will use three running examples when discussing AI technologies:

- **Weather forecasting:** to predict hourly weather in Mexico City for the next 24 hours. An AI model might use the following inputs and outputs:
    a. Input *x* (24x256x256x3 + 3x24 = 4,718,664 numbers)
        i. hourly satellite imagery (images of dimension typically 256x256x3, corresponding to width x height x channels) from the previous 24 hours
        ii. hourly weather data (temperature, precipitation, humidity) from previous 24 hours
    b. Output *y*: 48 numbers, corresponding to the forecasted hourly precipitation and temperature values for the next 24 hours

- **Chatbot:** to create a chatbot like ChatGPT. The AI model's inputs and outputs are:
    c. Input *x*: sequence of words provided by user
    d. Output *y*: sequence of words, corresponding to the most likely set of words that follow the user input

- **Smart thermostat controller:** an AI model that controls a thermostat to keep the room temperature within a given range, while also trying to minimize the carbon emissions associated with the air conditioning / heater energy usage. The AI model's inputs and outputs could be:
    e. Input *x* (10 numbers):
        i. current temperature
        ii. current time of day
        iii. today's sunset and sunrise time
        iv. forecasted hourly external temperature for the next 3 hours
        v. forecasted electricity carbon intensity (a measure of how clean the energy production is at a given time) for the next 3 hours
    f. Output *y*: an action chosen from "turn off system", "turn on AC", "turn on heater"

In essence, practically all AI algorithms boil down to approximating arbitrary mathematical functions $y = f(x)$. Given a large enough dataset of input-output examples (x,y), AI algorithms can usually learn a reasonable approximation of this input-output mapping. The differences between AI technologies can be roughly categorized along 4 areas:

1. Goal of the model

There are 3 common AI model goals, which can roughly be categorized as predictive AI, generative AI, and reinforcement learning.

In **predictive AI**, the goal is to learn the single most likely output *y* given an input *x*. This goal is most appropriate when there is a single "correct answer" for each input, and we have sufficient historical data consisting of (*x*,*y*) pairs that are accessible that we could use to train a model. The weather forecasting example falls into this category. If the forecasting model predicts that tomorrow at 9am the temperature will be 25C, we can wait until tomorrow 9am and check whether the temperature is actually 25C.

In **generative AI**, the goal is to sample a likely output *y* (or possibly several likely outputs) from the distribution of possible outputs given an input *x*. This goal is appropriate when we want to sample many possible outcomes *y* for each input *x*. Here, there might not be a single "correct answer" but rather a spectrum of possible answers. The chatbot example falls into this category. Even if asked the same prompt (e.g., "write a short love story"), the user may want the chatbot to generate different answers each time. The different answers may reflect different writing styles or creative expression, for example.

In **reinforcement learning**, the goal is to take an optimal action *y* given an input *x*. This goal is appropriate when *y* corresponds to an action, instead of a prediction. Often, it may be hard to collect a dataset of input-optimal action (*x*,*y*) pairs, especially when we do not know the optimal action in advance, but there may be some metric by which to measure the optimality of an action taken. The smart thermostat controller example falls into this category. The AI agent takes an action ("turn off system," "turn on AC", or "turn on heater") which results in a change in the environment (change in the room temperature) as well as reward or penalty for the agent (any deviation from the desired temperature range, as well as the amount of carbon emissions associated with the energy used). This reward signal is then used to update the AI agent's decision-making process.

2. Format of data accepted by the AI algorithm

Currently, there is not a "universal" AI algorithm that can handle all types of data inputs. Instead, AI algorithms are usually designed to handle a particular set of data formats or modalities. The different modern subfields of AI have generally developed around the needs of handling different data types.

Computer vision refers to the subfield of AI that processes 2-D and 3-D images, videos, and point clouds (3-D plots, *e.g.*, from LIDAR cameras). The weather forecasting example would use a computer vision model in order to process satellite imagery inputs. Other examples include the visual processing units in self-driving cars, facial recognition software, automated X-ray medical image diagnoses, and image/video generation tools.

Natural language processing refers to the subfield of AI that deals with text, language, and speech data. The chatbot example uses natural language processing to understand the human text input. Other examples include voice recognition systems (*e.g.*, Siri and Alexa), text-to-speech systems, and translation systems (*e.g.*, Google Translate).

Network analysis and **graph learning** refer to the subfield of AI that deals with data that is best represented via a network (*i.e.*, a "graph" in mathematical parlance). Examples of network data include the electrical grid, social networks, and transportation road networks.

Recommender systems refers to the subfield of AI that processes human preference data. Examples include search engines (*e.g.*, Google Search) and product recommendation services (*e.g.*, Amazon and Netflix).

Reinforcement learning refers to the subfield of AI that processes data from interacting with an environment. Reinforcement learning is the field of AI most related to robotics. The smart thermostat controller example uses reinforcement learning to learn how to best control a thermostat in the context of the environment.

These domains are not mutually exclusive. For example, a text-to-image generator such as OpenAI's DALL-E or StabilityAI's StableDiffusion combines natural language processing with computer vision algorithms. Many robotic systems combine reinforcement learning with computer vision algorithms to enable robotic systems to understand their surroundings and learn to achieve a goal.

3. The way the dataset is collected and/or designed

AI algorithms need to be given large amounts of training data, but perhaps even more important, is the structure of the input-output pairs. It may be instructive to walk through the dataset design for the three running examples.

The **weather forecasting** example has the most straightforward dataset design. An ideal training dataset would consist of millions of (x,y) pairs corresponding to historical weather observations. This is known as **supervised training**.

The **chatbot example** has a complex dataset design, because there do not exist large enough datasets for what "good" or "correct" outputs from a chatbot should look like. Instead, modern chatbots such as OpenAI's ChatGPT rely on three separate training datasets, each for training the chatbot to learn different tasks.

4. First, a chatbot is trained on a large collection of (*first sentence*, *next sentence*) pairs, scraped from websites such as Wikipedia and The New York Times. This collection has billions (possible trillions) of such pairs. By training on this next-sentence-prediction dataset, the chatbot becomes proficient at understanding the meaning of words, as well as proper grammar and punctuation. However, this dataset is insufficient for the chatbot to understand how to actually behave like a chatbot.
5. Second, the chatbot is trained on a collection of (*prompt*, *response*) pairs that demonstrate how the chatbot should respond to prompts. This dataset usually contains about a hundred thousand examples, and these prompts and responses are generated by real humans specifically for the purpose of training the chatbot to behave more chatbot-like.
6. Finally, the chatbot is fine-tuned on a third dataset of (*prompt*, *better response*, *worse response*) triplets. The prompts are generated by real humans, but the two responses are generated by the chatbot itself, and the better-vs-worse distinction is decided by human annotators. The chatbot is trained to be more likely to output the better response and less likely to output the worse response. It is through this final dataset that the chatbot learns to be more "aligned" with human values, such as outputting more factually accurate statements and avoiding racist responses.

The **smart thermostat controller** example has a different dataset design than the other examples. If we already knew what the best action y ("turn off system", "turn on AC", or "turn on heater") was for each input x, then we could construct such a dataset of (x,y) pairs and use supervised training, as in the weather forecasting example. However, we humans are probably not the best at picking the best action that simultaneously maximizes thermal comfort while reducing energy usage and associated carbon emissions. Because this objective is complex and difficult to optimize directly, we instead collect a dataset of $(x,\hat{y},r)$, where $\hat{y}$ is a possibly suboptimal action for the input x, and r is the corresponding objective value (also called a reward). Based on this dataset, reinforcement learning algorithms learn a mapping from an input x to an optimal action y, without being told the optimal action for each x.

4. The specific mathematical methods used to approximate the x-to-y mapping

Many different mathematical functions can be used to approximate an x-to-y mapping. One of the simplest mathematical models is linear regression, which is only useful if the x-to-y mapping has a linear structure. For images and text data, though, this linear structure no longer holds, and therefore more advanced mathematical models are necessary. It turns out that polynomials, neural networks, and Fourier functions are all "universal function approximators," meaning that when trained on enough data, they can almost exactly approximate *any* x-to-y mapping. (Using neural networks as the approximation function is known as *deep learning*.) The reason AI engineers may choose one "universal function approximator" model over another usually boils down to the model's computational costs, interpretability, and ability to generalize to unseen data. This is more the realm of AI research and engineering, and arguably less important for businesses and managers seeking to deploy AI.

TRAINING VS. INFERENCE

The term "training" refers to the act of updating an AI's mathematical model parameters to fit the training data. The term "inference" refers to *using* a trained AI model, that is running a trained AI model on new inputs. Inference is only possible after an AI model has been trained. Training tends to be much more computationally expensive, with most modern AI models requiring graphics processing units (GPUs) to speed up computation. In contrast, inference tends to be much less computationally expensive. Typically, training an AI model on a dataset is usually at least 100-1000x slower than running AI inference on the same inputs.