

DISSERTATION

STUDYING CLIMATE INTERVENTION SCENARIOS WITH DATA SCIENCE METHODS

Submitted by

Charlotte J. Connolly

Department of Atmospheric Science

In partial fulfillment of the requirements

For the Degree of Doctor of Philosophy

Colorado State University

Fort Collins, Colorado

Fall 2025

Doctoral Committee:

Advisor: Elizabeth A. Barnes

James W. Hurrell

David Randall

Chuck Anderson

Copyright by Charlotte J. Connolly 2025

All Rights Reserved

ABSTRACT

STUDYING CLIMATE INTERVENTION SCENARIOS WITH DATA SCIENCE METHODS

Stratospheric aerosol injection (SAI) is a proposed climate intervention strategy that may alleviate the most severe climate change impacts. The ways in which a possible SAI deployment may occur are vast and require social, environmental, and political consideration. To help scientists decide which of the many possible futures to study, scenario design methods are often used. SAI scenarios generally fall into one of two categories. The first, coordinated global deployment, is characterized by global cooperation among all or most of the world's population in deciding how to deploy SAI. This work explores a common tool used alongside coordinated global deployment climate model simulations to quantify the tool's sensitivity to internal variability. The second category, single actor deployment, is characterized by a single party or a small group independently deploying SAI to achieve their own climate goals. Under a range of these scenarios, atmospheric responses may motivate additional counter SAI deployments or retaliation from other parties. This work quantifies atmospheric responses shortly after a wide range of single actor deployments.

Finally, this work concludes by introducing a novel machine learning emulator designed to explore SAI deployment scenarios for use alongside earth system models to inform future simulation design. This machine learning emulator is comprised of multiple interpretable machine learning models that are trained on simulations from the Community Earth System Model version 2 including a novel simulation titled, Climate Response After Stratospheric Sulfur dioxide injection (CRASSULA). Here the emulator design is described along with how this emulator can be used for scenario design. The work in this dissertation emphasizes the need to question how SAI scenarios are simulated, what assumptions are made, and what possible scenarios are being overlooked.

ACKNOWLEDGEMENTS

The work in this dissertation was partially funded by Defense Advanced Research Projects Agency (DARPA) grant HR00112290071 and NSF CAREER AGS-1749261 under the Climate and Large-Scale Dynamics program.

This dissertation exists in its current form solely because of the many people who supported me throughout this process. My committee, Elizabeth A. Barnes, James W. Hurrell, David Randall, and Chuck Anderson always challenged me and helped me grow as a scientist. The Barnes group was always there to answer my questions. My best friends Daniel and Drew helped me get through the largest challenges by being wonderful listeners and providing emotional support. Lastly, all of the my friends and mentors who inspired me and helped me discover my passion for science, teaching, and mentoring.

DEDICATION

To:

My Mama Linda Connolly who taught me how to be selfless.

My Papa Roger Connolly who taught me how to be tough.

My Grandma Nellie Davis who taught me how to be creative.

My Grandpa Roger Davis who taught me how to be resourceful.

TABLE OF CONTENTS

ABSTRACT	ii
ACKNOWLEDGEMENTS	iii
DEDICATION	iv
LIST OF TABLES	vii
LIST OF FIGURES	viii
Chapter 1 Introduction	1
Chapter 2 Quantifying the Impact of Internal Variability on the CESM2 Control Algo- rithm for Stratospheric Aerosol Injection	5
2.1 Introduction	5
2.2 Methods	6
2.2.1 Controller Inputs	6
2.2.2 Changes to the Controller	8
2.3 Results	9
2.4 Discussion	15
2.5 Conclusion	16
Chapter 3 Atmospheric Response Shortly After a Single Actor Deployment	18
3.1 Introduction	18
3.2 Climate Response After Stratospheric Sulfur dioxide Injection (CRASSULA)	19
3.2.1 Designing CRASSULA	19
3.2.2 Challenges while Simulating CRASSULA	21
3.3 Composites and Results	22
3.3.1 AOD and Net Flux at the Surface	23
3.3.2 Temperature	24
3.3.3 Precipitation	26
3.3.4 Ozone	30
3.4 Conclusion	31
Chapter 4 Emulating the Temperature Response to Stratospheric Aerosol Injection Sce- narios	34
4.1 Introduction	34
4.2 Data	35
4.3 Machine Learning Training and Validation	37
4.3.1 Internal Variability Machine Learning Model	39
4.3.2 CO ₂ Machine Learning Model	42
4.3.3 AOD Machine Learning Model	43
4.3.4 SAI Scenario Learning Model	45

4.4	The Multi-Model Approach for Scenario Exploration	47
4.5	Conclusion	51
Chapter 5	Closing Remarks	53
5.1	Future of SAI Scenarios	53
5.2	Future of SAI Machine Learning Emulators	55
Bibliography	58
Appendix A	Supporting Information for Chapter 2	72
Appendix B	Supporting Information for Chapter 4	78
B.1	Details about the internal variability machine learning model	78
B.2	Details about the CO ₂ machine learning model	79
B.3	Details about the AOD machine learning model	81
B.4	Details about the SO ₂ machine learning model	82

LIST OF TABLES

4.1	Data used to train machine learning models	36
-----	--	----

LIST OF FIGURES

2.1	Schematic showing patterns that make up two different controller inputs. The base injection is the amount injected given only the base state while the new injection is the injection amount when all components are summed. Percent change shows how much internal variability changes the total injection as a function of the base state.	10
2.2	Percent change in total SO ₂ injection as a function of (a) ENSO, (b) NAO, and (c) SAM events. Solid lines use data from ARISE-SAI-CTRL and dashed lines use data from CESM2-LE. Green lines use year 2035 base state and orange lines use year 2045 base state. Black dashed line marks zero percent change.	11
2.3	Percent change in total SO ₂ injection as a function of two internal variability indices using composites from the CESM2-LE. Top row uses the year 2035 base state and bottom row uses the year 2045 base state. Black line in each panel separates positive percent change (red shading) from negative percent change (blue shading).	12
2.4	Mt. Pinatubo’s impact on the total injection where (a) are the temperature anomalies associated with the Mt. Pinatubo eruption (volcano component of controller input). The new injection is the total SO ₂ injected given the base state and the volcano component. Percent change shows as a function of the base state, how much the Mt. Pinatubo eruption changes the total injection. Panels (b), (c), and (d) are similar to Figure 2.2 but also include the volcano component in the controller input.	15
3.1	Shows the SAI locations in CRASSULA along with how many times each location is sampled for 1Tg/yr (a), 4Tg/yr (b), 8Tg/yr (c), and 16Tg/yr (d).	22
3.2	Average AOD anomaly (a-e) and average net shortwave flux at the surface (f-j) for SO ₂ injections at 30°N. The injection scenarios are split into four injection rates 1Tg/yr (a, f), 4Tg/yr (b, g), 8Tg/yr (c, h), and 16Tg/yr (d, i) and stippling represents a p-value of less than 0.05. Panels (e) and (j) are global average anomalies for year one (purple triangle) and year two (green circle). Dark purple and dark green indicate ensemble means and light green and light purple show global mean anomalies for each ensemble member. Black horizontal line denotes zero.	25
3.3	Average AOD anomaly (a-e) and average net shortwave flux at the surface (f-j) for SO ₂ injections at 30°S. The injection scenarios are split into four injection rates 1Tg/yr (a, f), 4Tg/yr (b, g), 8Tg/yr (c, h), and 16Tg/yr (d, i) and stippling represents a p-value of less than 0.05. Panels (e) and (j) are global average anomalies for year one (purple triangle) and year two (green circle). Dark purple and dark green indicate ensemble means and light green and light purple show global mean anomalies for each ensemble member. Black horizontal line denotes zero.	26

3.4	Average near surface temperature anomaly for SO ₂ injections at 30°N (a-e) and 30°S (f-j). The injection scenarios are split into four injection rates 1Tg/yr (a, f), 4Tg/yr (b, g), 8Tg/yr (c, h), and 16Tg/yr (d, i) and stippling represents a p-value of less than 0.05. Panels (e) and (j) are global average anomalies for year one (purple triangle) and year two (green circle). Dark purple and dark green indicate ensemble means and light green and light purple show global mean anomalies for each ensemble member. Black horizontal line denotes zero. Note the different y-axis on panels (e) and (j).	27
3.5	Average precipitation anomaly for SO ₂ injections at 30°N (a-e) and 30°S (f-j). The injection scenarios are split into four injection rates 1Tg/yr (a, f), 4Tg/yr (b, g), 8Tg/yr (c, h), and 16Tg/yr (d, i) and stippling represents a p-value of less than 0.05. Panels (e) and (j) are global average anomalies for year one (purple triangle) and year two (green circle). Dark purple and dark green indicate ensemble means and light green and light purple show global mean anomalies for each ensemble member. Black horizontal line denotes zero. Note the different y-axis on panels (e) and (j).	28
3.6	Average precipitation anomaly for SO ₂ injections at the equator. The injection scenarios are split into four injection rates 1Tg/yr (a), 4Tg/yr (b), 8Tg/yr (c), and 16Tg/yr (d) and stippling represents a p-value of less than 0.05.. Panel (e) is global average anomalies for year one (purple triangle) and year two (green circle). Dark purple and dark green indicate ensemble means and light green and light purple show global mean anomalies for each ensemble member. Black horizontal line denotes zero.	29
3.7	Average present change in ozone for SO ₂ injections at 30°S (a,b), equator (c,d), and 30°N (e,f). Stippling represents a p-value of less than 0.05. Panels b, d and, f are global average anomalies for year one (triangle) and year two (circle). Pink and purple are averages with tropospheric ozone. Blue and orange are averages with stratospheric ozone. Dark purple, pink, blue, and orange indicate ensemble means and light purple, pink, blue, and orange show global mean anomalies for each ensemble member. Black horizontal line denotes zero.	33
4.1	Schematic of the four machine learning models that comprise the machine learning emulator (black boxes). Inputs are indicated by purple squares and blue squares indicate inputs that are predictions from other machine learning models. Arrows represent the flow of information within this multi-model architecture.	38
4.2	Mean absolute error the on piControl-WACCM data using the internal variability machine learning model (a), on SSP2-4.5 using the CO ₂ machine learning model (b), on CRASSULA using the AOD machine learning model(c), and on CRASSULA using the SAI machine learning model (d) for respective testing data set.	40
4.3	Global mean near surface land temperature from an ARISE-SAI-1.5 ensemble member (black line) compared to data generated with the machine learning models. The internal variability machine learning model predicts temperature anomaly relative to 281.5K (green line), the CO ₂ machine learning model adds heating associated with the amount of CO ₂ globally (red line), and the SAI machine learning model adds cooling in response to SAI scenario (blue line)	41

4.4	Globally average error from the AOD machine learning model using 10 sets of 30-years from CRASSULA to identify the distribution of errors. The purple line is the average error for using the AOD machine learning model to predict a single year with the shading representing the min and max error. The orange line shows the globally averaged error when using the AOD machine learning model autoregressively with the orange shading representing minimum and maximum error.	44
4.5	(c) The global mean change in AOD in response to a single continuous injection (pink line) compared to a scenario where a second party deploys SAI at a new location at year 4 (purple line). Panels (a) and (b) shows the AOD in year 10 for scenario 1 and scenario 2.	46
4.6	Global mean land temperature for SSP2-4.5 (red shading), ARISE-SAI-1.5 (green shading), data generated by the SAI emulator to replicate SSP2-4.5 (purple lines), and data generated by the SAI emulator to replicate AIRSE-SAI-1.5 (green lines). Shaded areas represent the range between the minimum and maximum values in SSP2-4.5 and ARISE-SAI-1.5.	48
4.7	Probability ensemble members experience perceived failure in the 10 year period following an SAI deployment. Perceived failure is categorized as a positive trend in near surface temperature across the 10 years post deployment. (a) Probability of perceived failure of the 10 member ARISE-SAI-1.5 simulation adapted from Keys et al., (2022). (b) Probability of perceived failure from the 10 member AIRSE-SAI-1.5 as predicted by the SAI emulator (c) Probability of perceived failure sampling internal variability from the piControl-WACCM simulation. (c) Probability of perceived failure sampling internal variability from the SSP2-4.5 simulation.	50
A.1	Black line shows the global average temperatures. Red line is the linear trend fit to the global averaged temperatures 10 years prior to Mt. Pinatubo. The blue line is extrapolating the red line for an additional 2 years and is used to calculate the temperature anomalies associated with the Mt. Pinatubo eruption.	73
A.2	Difference between the 2045 base state and 2035 base state.	73
A.3	Table showing which base state and internal variability patterns are used to make each of the figures in the manuscript.	74
A.4	The portion of the SO ₂ injection (Tg/year) in response to (a) ENSO, (b) NAO, and (c) SAM using the base state years 2035 (green line) and 2045 (orange line). Solid lines use data from ARISE-SAI-CTRL and dashed lines use data from CESM2-LE.	75
A.5	Percent change in total SO ₂ injection as a function of two internal variability indices but using composites from ARISE-SAI-CTRL. Top row uses the year 2035 base state and bottom row uses the year 2045 base state. Black line in each panel separates positive percent change (red shading) from negative percent change (blue shading).	76
A.6	Deviations in T0, T1, and T2 as a function of ENSO and NAO using composites from CESM-LE. Top row uses the year 2035 base state and bottom row uses the year 2045 base state. Black line in each panel separates positive error (red shading) from error (blue shading).	77
A.7	Total SO ₂ injection as a function of keeping either T0, T1, or T2 constant and varying the other two errors.	77

B.1	Figure showing what simulations were used to train each individual machine learning models	78
B.2	Panel (a) shows 50 years of global mean temperature from piControl-WACCM (black line) compared to predicted global mean temperature (red line). Panel (b) shows average errors of the machine learning model prediction (red line) compared to the error of a persistence prediction (purple line). Blue vertical lines show when a persistence prediction out performs the machine learning model prediction.	79
B.3	Panel (a) shows 50 years of global mean temperature from piControl-WACCM (black line) compared to predicted global mean temperature (red line). Panel (b) shows average errors of the machine learning model prediction (red line) compared to the error of a persistence prediction (purple line). Blue vertical lines show when a persistence prediction out performs the machine learning model prediction.	80
B.4	Predicted global average AOD (black line) compared to the true global average AOD (black line). Green line represents predictions when the SO ₂ input is kept constant and only the initial AOD input is changed. The blue line represents predictions when the initial AOD input is kept constant and only the SO ₂ input is changed.	81
B.5	Schematic of the SAI machine learning model architecture. Each layer is labeled with the input dimensions. Layers are colored to represent the type of layer. Dark purple in a convolutional layer. Pink is an average pooling layer. Purple is a dense layer. Light blue is a upsampling layer. Arrows indicate skip connection and Teal squares represent the inputs and outputs.	83
B.6	Global mean land temperature for SSP2-4.5 (red shading), ARISE-SAI-1.5 (green shading), data generated by the SAI emulator to replicate SSP2-4.5 (purple lines), and data generated by the SAI emulator to replicate AIRSE-SAI-1.5 (green lines). Shaded areas represent the range between the minimum and maximum values in SSP2-4.5 and ARISE-SAI-1.5.	84
B.7	Schematic of the internal variability machine learning model architecture. Each layer is labeled with the input dimensions. Layers are colored to represent the type of layer. Dark purple in a convolutional layer. Pink is an average pooling layer. Purple is a dense layer. Light blue is a upsampling layer. Arrows indicate skip connection and Teal squares represent the inputs and outputs.	85
B.8	Schematic of the CO ₂ machine learning model architecture. Each layer is labeled with the input dimensions. Layers are colored to represent the type of layer. Dark purple in a convolutional layer. Pink is an average pooling layer. Purple is a dense layer. Light blue is an upsampling layer. Arrows indicate skip connection and Teal squares represent the inputs and outputs.	86
B.9	Schematic of the AOD machine learning model architecture. Each layer is labeled with the input dimensions. Layers are colored to represent the type of layer. Dark purple in a convolutional layer. Pink is an average pooling layer. Purple is a dense layer. Light blue is a upsampling layer. Arrows indicate skip connection and Teal squares represent the inputs and outputs.	87

Chapter 1

Introduction

Reducing carbon dioxide (CO₂) and other greenhouse gas emissions remains the top priority for averting climate change impacts. However, current pledges to reduce global emissions are likely insufficient for keeping global mean temperature below the 2°C global mean temperature target set by the Paris Agreement [1, 2]. As the threat of rising global temperatures increases, there is a call for research into methods of deliberate interventions in the Earth system, known as climate intervention strategies [3–5]. While climate intervention strategies are not without risk, these strategies could be used as a last resort to temporarily mitigate the worst climate change impacts and provide more time for global emissions to be reduced [6–8].

Stratospheric aerosol injection (SAI) is one such proposed method of climate intervention where aerosols or their precursors are continuously injected into the stratosphere to reflect some incoming shortwave radiation and temporarily alleviate some of the worst consequences of climate change. Importantly, SAI does not revert the Earth system to a previous state as it does not perfectly counter the warming response to greenhouse gases or remove CO₂ from the atmosphere [9, 10]. For example, SAI can slow or reverse melting sea ice in the Arctic [11] but cannot stop surface ocean acidification caused by increased concentrations of carbon dioxide in the atmosphere [12, 13]. SAI also has its own risks and uncertainties such as possible slow down to stratospheric ozone recovery through impacts to stratospheric chemistry and stratospheric circulation [14, 15].

Some uncertainties and risks can be managed with careful development of the SAI deployment strategy, which describes how an SAI deployment may occur. This includes when the injection starts, where the injections occur, how much is injected, and for how long the injection continues [16–18]. Deployment strategies generally fall into two categories. The first is a *coordinated global deployment* which is characterized by global cooperation among all or most of the world's population in deciding how to deploy SAI. Global cooperation allows for flexibility in the deployment strategy so that possible negative consequences of SAI may be minimized [19]. Through

cooperation and regulation required to make a coordinated global deployment successful may be challenging given complex geopolitics [20]. While this deployment strategy often attempts to minimize negative impacts, there can be some trade-offs with different temperature targets [21] and temperature targets that cause rapid cooling may be just as harmful as rapid warming from climate change [22].

The second deployment scenario is a *single actor deployment* which is characterized by a single party or a small group independently deploying SAI to achieve their own climate goals [3, 23, 24]. When considering single actor scenarios, regional responses, community opinions about climate change and climate intervention, and the geopolitical landscape may all influence decisions about an SAI deployment or additional responses by other actors [25, 26]. For example, extreme events following a deployment may be falsely or correctly attributed as an SAI response and motivate subsequent decisions including retaliation or additional climate intervention [26, 27]. Consequently, the decision of whether to deploy SAI or not is one of complex social, environmental, and political considerations [26, 28].

This dissertation works to better understand possible future scenarios that involve SAI and how they are simulated in Earth system models. To do this, the implications about how coordinated global deployments are often simulated and possible atmospheric responses shortly after a single actor deployment are quantified and discussed. Chapter 2 investigates how sensitive a control theory method that is commonly used in model simulations of coordinated global deployment scenarios is to internal variability. This control theory method is referred to as a control algorithm and keeps a system at a predetermined target by applying different forcings [29]. In coordinated global deployment model simulations, the control algorithm maintains global mean temperature and temperature gradients by determining how much sulfur dioxide (SO₂) to inject. To quantify its sensitivity to internal variability, an offline version of the control algorithm is passed patterns of internal variability to quantify its sensitivity to each. Understanding how sensitive the control algorithm is to internal variability motivates a discussion about whether the magnitude at which the control algorithm is sensitive to internal variability is acceptable.

The next two chapters introduce and use a novel simulation, titled Climate Response After Stratospheric Sulfur dioxide Injection (CRASSULA), which serves two complementary purposes. The first is to provide further information about atmospheric responses shortly after a range of single actor deployments and the second is to serve as training data for a machine learning emulator that allows quick and cheap exploration of further scenarios.

Chapter 3 quantifies how common atmospheric variables respond shortly after a single actor deployment. Many scenarios of single point injections have unfavorable atmospheric impacts such as regional changes to precipitation [30, 31]. While unfavorable, technology advancements, decreasing cost of implementing climate intervention, and increasing climate change impacts, all increase the likelihood of a single actor deployment. CRASSULA is used to investigate responses in near-surface temperature, precipitation, and ozone one and two years after a single actor deployment. Quantifying impacts shortly after a single actor deployment will highlight regions that are at risk for experiencing impacts. By showing possible climate states shortly after a single actor deployment, this chapter emphasizes the need for understanding these impacts for communication and for planning any necessary efforts to mitigate SAI impacts.

These chapters highlight the need to consider how SAI scenarios are simulated and scenarios of single actor deployments that may have impacts on communities. Often, the limited computing time or limited computing resources restrict what scenarios can be simulated in an Earth system model. Chapter 4 introduces an initial version of a machine learning emulator that predicts the near surface land temperature response to SAI and can be used to explore SAI scenarios quickly and cheaply. This emulator is not designed to replace Earth system models for SAI impact research. It is designed to be used to explore various injection scenarios and provide information about possible sensitivities or dependencies on internal variability. The emulator can also be used to consider impacts alongside possible political and social responses which may motivate additional SAI deployments. This information can then be used to design scenarios that are simulated with an Earth system model. This chapter outlines the steps used to train the machine learning emulator along with how it is used to investigate SAI scenarios.

The content in Chapter 2 has been peer-reviewed and published in *Earth's Future* [32]. As such, this paper has been adopted for use in this dissertation:

Connolly, C., Prewett, E., Barnes, E. A., & Hurrell, J. W. (2024). Quantifying the impact of internal variability on the CESM2 control algorithm for stratospheric aerosol injection. *Earth's Future*, 12(6), e2023EF004300.

Chapter 2

Quantifying the Impact of Internal Variability on the CESM2 Control Algorithm for Stratospheric Aerosol Injection

2.1 Introduction

Several modeling projects have been conducted to understand how the climate system may respond to additional SO_2 in the stratosphere [33–37]. Some of these simulations implement a feedback algorithm, called a controller, which determines how much aerosol to inject into the stratosphere to maintain the climate system at pre-established temperature targets [29, 36, 37].

The Assessing Responses and Impacts of Solar climate intervention on the Earth system with Stratospheric Aerosol Injection (ARISE-SAI-1.5) simulations, performed with version two of the Community Earth System Model (CESM2; [38]), implement a controller to keep global mean surface temperature near 1.5°C while also maintaining temperature gradients so that atmospheric circulations are minimally impacted [37]. The controller accomplishes this by comparing the global temperature (T_0), the north-south temperature gradient (T_1) and the Equator-to-pole temperature gradient (T_2) to predetermined targets of 288.64, 0.8767, and -5.89 K, respectively [37, 39]. Deviations between the T_0 , T_1 , and T_2 values calculated from model output and the individual predetermined targets are used by the controller to determine how much SO_2 to inject at four different locations (30°N , 15°N , 15°S , 30°S).

In the ARISE-SAI-1.5 simulations, the controller impacts the climate system by determining how much SO_2 is needed to maintain the climate system at the pre-determined targets. Additionally, since the controller determines injection amounts based on deviations of T_0 , T_1 , and T_2 from their respective targets, global and regional temperature patterns driven by internal climate variability can impact injection amounts. A handful of studies have begun to explore how the

controller and the simulated climate system impact one another. For example, MacMartin et al. (2014) show that the way in which the controller is tuned and the lag between the controller input and the response of the system can impact the internal variability of the climate system. Diao et al. (2023) use data from the ARISE-SAI-1.5 simulations to show that ENSO accounts for 70% of the year-to-year variability in injection anomalies determined by the controller. In this work, we further explore how internal variability and volcanic eruptions impact SO_2 injection determined by the controller by passing global temperature maps with different internal variability patterns into an offline version of the ARISE-SAI-1.5 controller.

2.2 Methods

The ARISE-SAI-1.5 controller sensitivity to internal variability is quantified by creating controller inputs, for which the warming pattern and the patterns of internal variability are known, and passing them to the controller. The way in which the warming patterns and patterns of internal variability are calculated is provided in section 2.2.2.1. An offline version of the ARISE-SAI-1.5 controller is used to explore a large range of climate states without having to run additional simulations. Details about the changes made to the ARISE-SAI-1.5 controller are in section 2.2.2.

2.2.1 Controller Inputs

The 10 member ARISE-SAI-1.5 control simulation (ARISE-SAI-CTRL) is used to create the controller inputs [37]. The ARISE-SAI-CTRL comes from the same model configuration the ARISE-SAI-1.5 controller was tuned for but does not contain any SAI. Using the ARISE-SAI-CTRL means that the controller inputs will not contain any SAI driven cooling.

Every controller input map contains one forced component which describes the climate warming trend. The forced component is defined as the smoothed annual-mean ensemble mean near surface temperature using years 2035–2070 from the ARISE-SAI-CTRL. However, since 10 members are not enough to remove all internal variability [40], the ensemble mean is smoothed by fitting a third order polynomial to the timeseries at each grid point. These smoothed data are used as the

base states, and we focus on the years 2035 and 2045 in this study. The year 2035 defines when SAI begins in ARISE-SAI-1.5. Year 2045 defines when SAI begins in a different set of simulations which are designed to inform about the atmospheric responses after a delayed deployment using the same AIRSE-SAI-1.5 temperature targets [19].

Unforced components, or *internal variability patterns*, are defined as monthly temperature anomalies composited based on internal variability events. Any number of internal variability patterns of interest can be added onto a base state to quantify their impacts on total injection amounts. This work focuses on variability associated with the El-Niño Southern Oscillation (ENSO; [41]), the Southern Annular Mode (SAM; [42]), the North Atlantic Oscillation (NAO; [43]), and the eruption of Mt. Pinatubo [44]. These modes of variability are selected because each produces strong temperature anomalies in different regions of the globe. ENSO influences temperature predominantly at low latitudes, the NAO influences temperature at the high latitudes of the Northern Hemisphere, the SAM influences temperature at the high latitudes of the Southern Hemisphere, and a Pinatuno-like volcanic eruption influences temperatures globally.

The climate indices used to composite temperature anomalies associated with ENSO, NAO, and SAM events are calculated using sea surface temperature and sea level pressure from the ARISE-SAI-CTRL and the CESM2-LE. Methods used to calculate each climate index are as follows:

- The ENSO index is defined by the Nino3.4 index [41] based on the five month average sea surface temperature within the 5°N-5°S, 120-170°W region. A positive ENSO event is characterized by warmer than average temperatures in central and eastern Tropical Pacific and a negative ENSO has cooler than average temperatures [41].
- The NAO index is defined by the principal component timeseries of the leading empirical orthogonal function of surface pressure anomalies within 20-80°N, 90°W-40°E [43]. When the NAO is positive, sea level pressures are anomalously low over the subpolar North Atlantic. The opposite occurs during the negative phase of the NAO. The sea level pressure anomalies drive circulation responses that impact regional temperature patterns in a variety

of ways [43, 45, 46]. For example, when the NAO is in its negative state it drives warmer than average temperatures across Northern Europe [47].

- The SAM index is calculated as the principal component of the leading empirical orthogonal function of sea level pressure over the region 20-90°S [42]. When the SAM is positive, sea level pressures over the Southern Hemisphere polar region are anomalously low and temperatures are cooler than average. However, when the SAM is negative, sea level pressures are anomalously high and temperatures are warmer than average [48, 49].

The temperature anomaly pattern associated with the Mt. Pinatubo eruption is defined as the average temperature anomaly for two years following the eruption (June 1991–June 1993). The 100 member CESM2-LE warming trend is estimated by fitting a line at every grid point to the ensemble mean surface temperature anomalies timeseries 10 years prior to the eruption (May 1981–May 1991). This line is extrapolated to June 1993, two years following the eruption, and then subtracted from the ensemble mean. Assuming the internal variability is removed by calculating the ensemble mean of 100 members and that the linear fit represents a short term continued warming trend, subtracting the linear fit from the ensemble mean estimates the temperature anomalies associated with the eruption of Mt. Pinatubo. Figure A.1 in the Appendix A demonstrates an example of fitting a line to calculate temperature anomalies associated with Mt. Pinatubo.

2.2.2 Changes to the Controller

The ARISE-SAI-1.5 controller is a proportional-integral control algorithm, or PI controller [50]. With a PI controller, the proportional term accounts for the current error between model output and the predetermined targets and the integral term accounts for any persistent errors in time. Constants, called gains, are tuned to determine how much of each component is needed to maintain the system at the user-specified targets [29, 50, 51]. The active controller in the ARISE-SAI-1.5 simulations has a ramp up time of 5 years, which reduces shock to the system, and considers errors from previous years in the calculation via the integral portion of the controller. More details about the complete ARISE-SAI-1.5 simulations and its active controller can be found in Richter et al.

(2022) and Kravitz et al. (2017) and the sources within. This work utilizes an offline version of the ARISE-SAI-1.5 controller where the gain values are kept the same (i.e., no additional tuning) and the controller is not connected to an active simulation.

Some additional changes are made to the offline ARISE-SAI-1.5 controller for this work. First, since the offline controller is not connected to a simulation and cannot shock the simulated climate system by suddenly injecting large amounts of SO_2 , the ramp-up period is reduced from 5 years to 1 year. Second, the offline controller only receives one input at a time; therefore the controller does not have errors from previous years to use when calculating an injection amount for the current input. These changes ensure that when a temperature pattern is fed through the controller, the injection amount is determined by a single temperature pattern and not an evolving state

2.3 Results

The total injection when only the base states are passed into the controller quantifies the total injection in response to the climate warming signal. For the base states of 2035 and 2045, the injections are 0.43 Tg/year and 1.44 Tg/year, respectively (Figure 2.1). Next, different combinations of internal variability patterns are added onto these base states to create new controller inputs that, when passed into the controller, quantify the impact of internal variability on the total injection amounts.

Figure 2.1 shows how the controller responds to the same patterns of internal variability occurring under different background warming. The difference between the 2035 base state and 2045 base state is found in Appendix A, Figure A.2. Controller input (a) in Figure 2.1 shows the base state from 2035, the temperature anomaly pattern associated with an ENSO index between 1.0 and 1.2, and the temperature anomaly pattern associated with NAO index between -1.2 and -1.0. When these three patterns are added together and then passed into the controller, the controller injects 0.71 Tg/year of SO_2 into the stratosphere. Adding the same internal variability patterns onto the base state 2045 (controller input (b)), the total injection increases to 1.56 Tg/year. The two patterns of internal variability shown in Figure 1 are responsible for increasing the total injection by 0.28

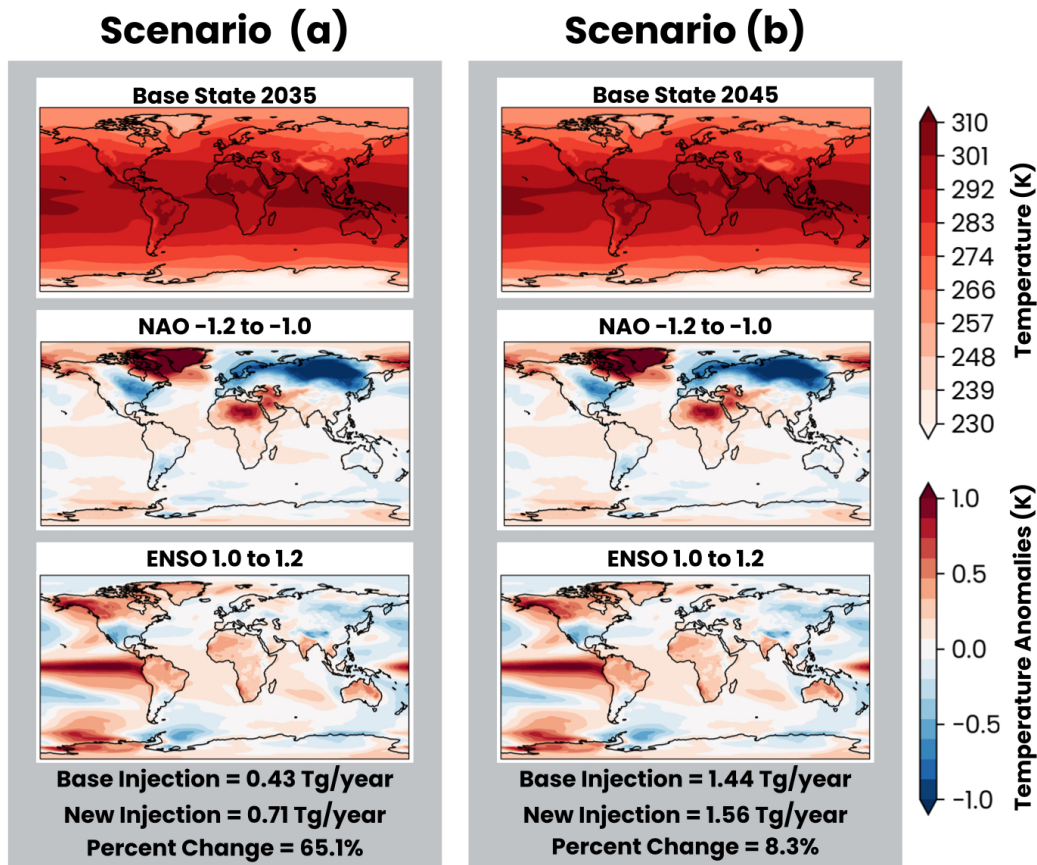


Figure 2.1: Schematic showing patterns that make up two different controller inputs. The base injection is the amount injected given only the base state while the new injection is the injection amount when all components are summed. Percent change shows how much internal variability changes the total injection as a function of the base state.

Tg/year in 2035 and by 0.12 Tg/year in 2045. These increases are similar in magnitude, but in relation to the base injection (i.e., percent change), the patterns of internal variability have a greater impact in 2035 than 2045: 65.1% increase compared to a 8.3% increase. This shows the amount of SO_2 injected in response to internal variability in 2035 is not equal to the amount of SO_2 injected in response to the same internal variability in 2045. In the following analysis, we add different combinations of internal variability patterns onto each base state (see A.3 in the Appendix A for what climate indices are used in each figure).

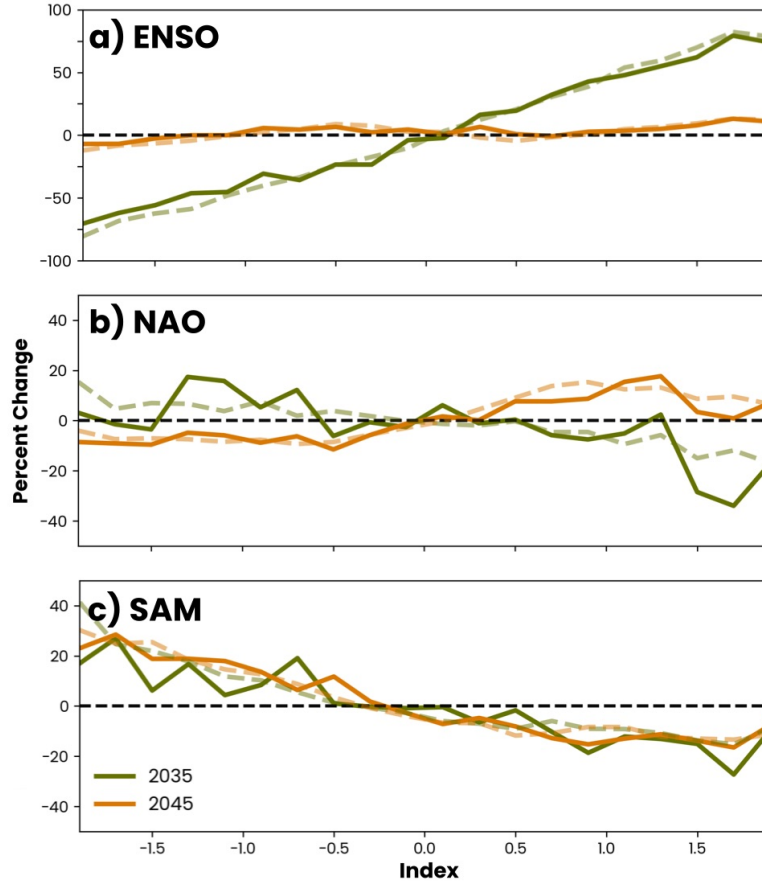


Figure 2.2: Percent change in total SO₂ injection as a function of (a) ENSO, (b) NAO, and (c) SAM events. Solid lines use data from ARISE-SAI-CTRL and dashed lines use data from CESM2-LE. Green lines use year 2035 base state and orange lines use year 2045 base state. Black dashed line marks zero percent change.

Since the impacts from internal variability on the controller-determined total injection depends on the base state, the ENSO, NAO, and SAM impacts on the total injection amounts are quantified as percent changes in Figure 2.2 (see A.4 in the Appendix A for total change). The percent change is calculated as the change in injection amount divided by the base injection amount (Figure 2.1). In Figure 2.2, positive ENSO events are shown to increase the amount of SO₂ injected and negative ENSO events decrease the amount SO₂ injected (Figure 2.2a), since positive ENSO events are shown to increase the global average temperature and negative events do the opposite [52]. The stronger the ENSO event, the greater the impact on the total injection, although, the impact of

ENSO anomalies on the controller decreases substantially from year 2035 to year 2045. This is because as the climate warming signal increases, the ENSO internal variability pattern is a smaller percentage of the input and thus a smaller role in the total injection amount.

The NAO has a smaller impact on the total injection in 2035 when compared to ENSO and its impact switches signs from 2035 to 2045. The SAM also has a smaller impact on the total injection than ENSO but its impact does not change from 2035 to 2045. Similar SAM and NAO impacts exist in both the ARISE-SAI-CTRL and CESM2-LE data and are therefore not a result of noise in the composites but a response to the internal variability patterns. In Figure 2.2, the base state pattern is the only difference between the green and orange lines in each panel, further demonstrating how the same internal variability pattern can have a different impact on the total SO₂ injection depending on the background state.

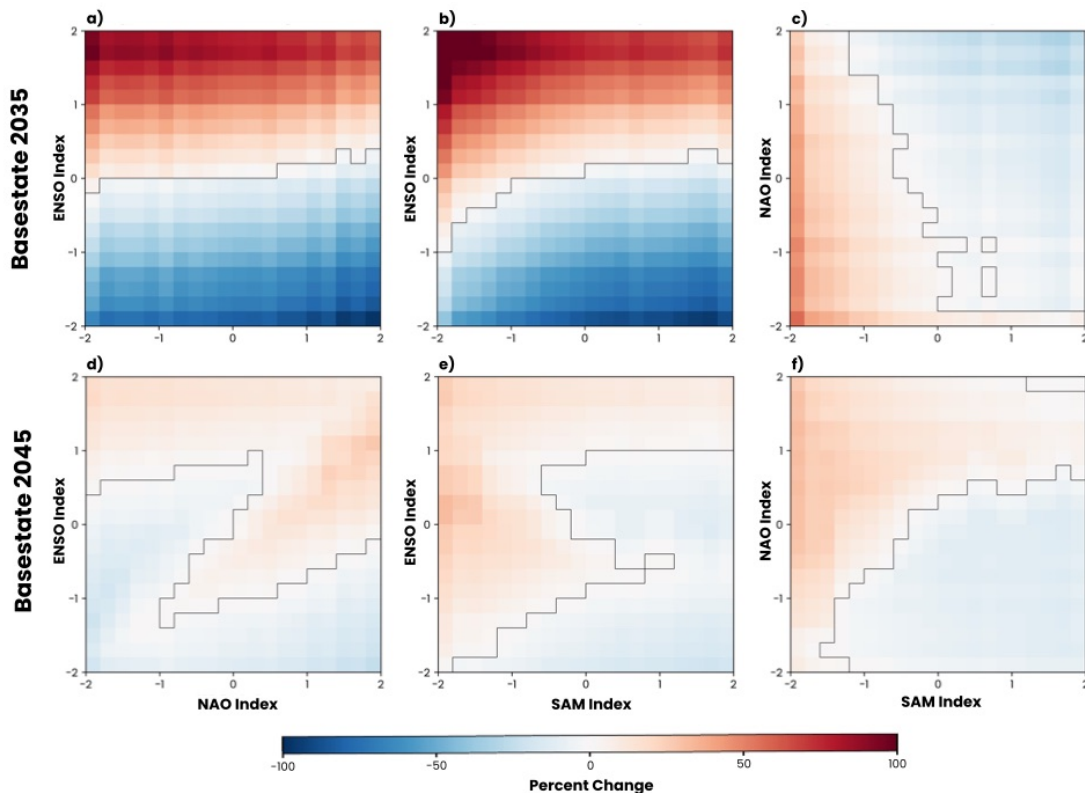


Figure 2.3: Percent change in total SO₂ injection as a function of two internal variability indices using composites from the CESM2-LE. Top row uses the year 2035 base state and bottom row uses the year 2045 base state. Black line in each panel separates positive percent change (red shading) from negative percent change (blue shading).

Next, in Figure 2.3, we show how the total SO₂ injection changes as a function of the combination of the two climate indices, with the top row using the 2035 base state and the bottom row using the 2045 base state. Since the controller responds similarly whether anomalies are calculated from ARISE-SAI-CTRL or CESM2-LE data as shown in Figures 2.2, Figure 2.3 shows results only using CESM2-LE anomalies. Results using ARISE-SAI-CTRL are in Supporting Information S5. Responses quantified in Figure 2.2 are also seen in Figure 2.3, such as the increase in total injection in response to a positive ENSO and a negative SAM. Figure 2.3 further shows that when these events occur together in 2035, the total injection increases more than when the patterns occur individually. A similar but opposite response is seen when a negative ENSO and a positive SAM occur simultaneously and drive a larger decrease in the total injection.

The impact on the total SO₂ injection from multiple internal variability patterns can change based on the base state (Figure 2.3). When using the 2035 base state, the largest impacts typically occur when the internal variability events are the strongest, as shown by the largest magnitudes of percent change found in the corners of the top row panels in Figure 2.3. For a base state year of 2045 (bottom row), we find that the largest magnitude changes no longer necessarily occur when the internal variability events are strongest. For instance, when the NAO is positive, the strongest impact to the total injection occurs when the ENSO index is near one rather than two (Figure 2.3d). When looking at the T0, T1, and T2 errors for the individual temperature patterns in Figure 3 (see A.6 in the Appendix A), the sign of the T1 error relative to the T1 target (0.8767) changes sign from negative in 2035 to positive in 2045 while the sign of T0 and T2 errors stay the same. The T1 value describes the north-south temperature gradient where a positive T1 value means the Northern Hemisphere is warmer than the Southern Hemisphere. Therefore, the sign change in T1 errors is likely in response to the uneven hemispheric warming that occurs in response to climate change. For more details about how deviations in T0, T1, and T2 change the total SO₂ injection see A.7 in the Appendix A.

We now explore the controller sensitivity to a volcanic eruption represented by the temperature anomaly pattern associated with the 1991 Mt. Pinatubo eruption (Figure 2.4a). Introducing the

volcanic eruption temperature pattern to the 2035 and 2045 base states decreases the amount of SO₂ the ARISE-SAI controller injects. When the volcanic pattern is added to the 2035 base state alone, the controller injects nothing and when added to the 2045 base state, the injection decreases by 0.58 Tg/year. The Mt. Pinatubo eruption injected approximately 10 Tg of SO₂ into the stratosphere [53,54] and was estimated to cool the Earth's surface by 0.5°C [55]. Therefore, a volcanic eruption the size of the Mt Pinatubo eruption would reduce the errors in T0 and thus decrease the total injection determined by the controller. In 2035, the global cooling in response to a Pinatubo-like eruption is enough to negate all experienced global-mean warming (at least from the controller's perspective), removing the need to inject any additional SO₂. The amount of SO₂ naturally injected by Mt Pinatubo is not enough to combat the amount of warming experienced in 2045.

Including an internal variability pattern in addition to the Mt. Pinatubo eruption pattern allows for the quantification of how much a Pinatubo-like eruption in combination with internal variability impacts the controller-determined SO₂ injection (Figures 2.4b, c, and d). In 2035, when a Pinatubo-like eruption removes the need to inject SO₂, only an ENSO event stronger than 0.5 forces the controller to inject. Warming associated with a positive ENSO greater than 0.5 is enough to cause the ARISE-SAI-1.5 controller to inject despite the volcanic eruption.

In 2045, when the controller input also contains a single internal variability pattern, a Pinatubo-like eruption decreases the total injection by about 40% as shown by the orange line in Figure 2.4 centered around -40] % rather than around 0% as it is in Figure 2.2. Additionally, the slopes of the orange lines in Figure 2.4 are smaller than the slopes of the orange lines in Figure 2.2 indicating that internal variability has a smaller impact on the total injection when occurring alongside a Pinatubo-like eruption. In summary, a volcanic eruption the size of Pinatubo in 2035 can reduce the total SO₂ injection by 100% except when the ENSO index is greater than 0.5. In 2045, a volcanic eruption of the same size can reduce the total SO₂ injection by about 40% and lessen the impacts of ENSO, NAO, and SAM on total injection amounts.

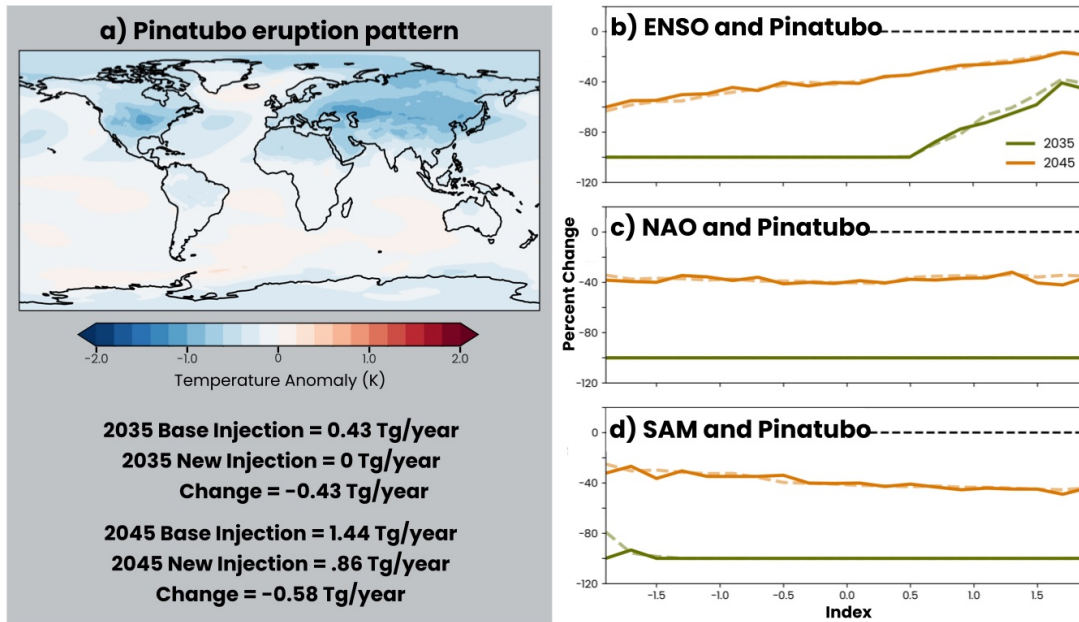


Figure 2.4: Mt. Pinatubo’s impact on the total injection where (a) are the temperature anomalies associated with the Mt. Pinatubo eruption (volcano component of controller input). The new injection is the total SO_2 injected given the base state and the volcano component. Percent change shows as a function of the base state, how much the Mt. Pinatubo eruption changes the total injection. Panels (b), (c), and (d) are similar to Figure 2.2 but also include the volcano component in the controller input.

2.4 Discussion

By design, controllers respond to variability of a system and therefore work well in systems with uncertainty [56]. However, a controller’s ability to respond and impact internal variability can result in complicated feedbacks where the controller can amplify or attenuate the frequency of internal variability, a feature explored thoroughly in MacMartin et al. (2014). These features of a controller are considered and balanced during the tuning phase of a controller. While this may present a challenge toward implementing a control algorithm in reality, Kravitz et al. (2014) showed that a control algorithm designed in one model could be used to meet the targets in a different model, demonstrating the controller’s ability to generalize to different systems. The results in this work show a way to quantify a controller’s sensitivities to a variety of temperature patterns post tuning, including to those outside of the system used to tune the control algorithm. While the method produces some climate states that may have statistically low chances of occurring or that may never occur, it allows for quick and cheap quantification of internal variability’s impact on the

total injection determined by the controller. Results in this work are confined to the 2035 and 2045 base states calculated from the ARISE-SAI-1.5 control simulations (i.e., temperature patterns are from the system the controller was tuned for). Given that this work shows that the internal variability's impact on the total injection depends on the background warming, using a different emissions scenario or model for the base state may result in different quantified sensitivities.

Once sensitivities are quantified, one can consider whether the magnitude in which different internal variability patterns impact the total injection is acceptable. For example, consider the ARISE-SAI-1.5 controller's response to a Pinatubo-like eruption. The controller injects less when there are naturally occurring aerosols cooling the planet. However, in regards to patterns of internal variability, is it acceptable that more SO_2 is injected when the atmospheric-ocean system is in an El Niño phase rather than a La Niña phase? Or should there be focus on ways to ensure that the majority of the SO_2 injection is in response to climate warming signal alone? Doing so would require the ability to separate the forced and unforced response in our current atmosphere or to predict them with considerable accuracy. Given that knowing or predicting the forced or unforced response with high accuracy is an ongoing area of research [57–59], implementing current methods to determine these responses would introduce further uncertainty into the feedback system.

2.5 Conclusion

This work quantifies the ARISE-SAI-1.5 controller sensitivity to internal variability and demonstrates a method that allows for a quick and effective quantification of controller sensitivity post tuning. The ARISE-SAI-1.5 controller's response to patterns of internal variability associated with ENSO, NAO and SAM as well as a Pinatubo-like eruption are quantified as these patterns cover Northern Hemisphere, Southern Hemisphere, and global temperature impacts. Focus is placed on quantifying these patterns of internal variability in relation to years 2035 and 2045, which correspond to the deployment year in ARISE-SAI-1.5 and the deployment year in delayed start, respectively (MacMartin et al., 2022). Using these two base state years, we show that internal variability's impact on the total injection is dependent on the background warming it is occurring

under. Using this method to explore and quantify sensitivities of a tuned controller provides the opportunity to explore controller responses to a system it is not tuned for, facilitates sensitivity comparisons between scenarios and earth system models, and may promote discussion about the extent to which an SAI-controller responds to variability internal to the climate system.

Chapter 3

Atmospheric Response Shortly After a Single Actor Deployment

3.1 Introduction

As technology advances and the cost of climate interventions are projected to decrease [60], it becomes more likely a single actor may deploy SAI to achieve their own climate goals. Single point injections are shown through model simulations to have global and often undesirable impacts [30, 31, 61]. These impacts include single hemisphere injections causing the Intertropical Convergence Zone to shift in response to unequal cooling of the hemispheres, impacting precipitation amounts for regions at and near the equator [31], and injections at the equator reducing equatorial precipitation via overcooling the tropical region more than the high latitudes [30]. An SAI deployment from a single actor may successfully drive regional cooling or mitigate local climate change impacts, but may also drive additional remote and possibly unfavorable responses. Understanding atmospheric responses on this short timeframe will help understand possible climate states post a single actor deployment, communicate possible impacts to affected communities, and strategize any necessary mitigation efforts.

Much of the research that investigates the atmospheric response to SAI focuses on coordinated deployment scenarios [36, 37, 62, 63], or timeframes many years after deployment and averaged across multiple years so that the atmospheric response from the deployment can be more easily identified [14, 64, 65].

To provide further information about atmospheric responses shortly after a single actor deployment and to train a machine learning learning emulator for quicker and cheaper exploration of deployment scenarios, this work introduces a novel simulation titled Climate Response After Stratospheric Sulfur dioxide Injection (CRASSULA). This chapter introduces the dataset and fo-

cuses on the atmospheric responses to SAI deployment strategy shortly after deployment. Chapter 4 will discuss the training and use of the machine learning emulator for scenario design.

3.2 Climate Response After Stratospheric Sulfur dioxide Injection (CRASSULA)

CRASSULA is a novel dataset simulated to train the machine learning model introduced in Chapter 4. CRASSULA is simulated with CESM2, run with a 0.9 latitude by 1.25 longitude grid, 70 vertical level resolution, and the middle-atmosphere version of the CESM2-WACCM6 atmospheric [66]. Sea surface temperatures are prescribed to seasonal values and the simulation is run with a perpetual 2000 emissions forcing so that responses in the atmosphere can be more easily attributed to the injections of SO₂ into the stratosphere.

SO₂ injections in CRASSULA occur for two years at a specified rate and location after which they abruptly stop and a new injection starts. Twenty-nine different injection locations are sampled in CRASSULA. The locations include seven latitudes (0, ±15, ±30, and ±50 latitude) and four longitudes (20, 125, 180, and 280 longitude) with an additional injection at the high latitudes of the Northern Hemisphere (70°N, 180°E). Injection height is predetermined as a function of latitude so that the injection is just above the tropopause. Four different injection rates are sampled (1, 4, 8, and 16 Tg/year) at each injection location. In total, 116 unique injection scenarios are simulated. Most injection scenarios are sampled eight times, however, due to challenges during simulating CRASSULA, some injection scenarios are sampled less frequently. See Figure 3.1 for the injection locations and how many times each is sampled. Section 3.2.2 provides more details about the challenges faced while simulating CRASSULA. CRASSULA also has a 30-year control simulation with the same configuration but with no injections of SO₂ into the stratosphere.

3.2.1 Designing CRASSULA

The original motivation for simulating CRASSULA was to train the machine learning emulator outlined in Chapter 4. However, given the large environmental footprint of running Earth

system models, it was important to ensure that CRASSULA could be used beyond training a machine learning emulator [67, 68]. Physical constraints such as limited core hours, limited access to a super computer, and limited storage, restricted the possible scope of CRASSULA. Designing CRASSULA became a balance of ensuring the data would be useful for training the machine learning emulator, that it fell within the physical computing constraints, and that it had applicability to research beyond training the machine learning emulator.

To get as many training samples as possible and to allow the machine learning emulator to learn possible nonlinearities in the response to different consecutive injections, each injection in CRASSULA occurs continuously for two years before it abruptly stops and a new one starts. The two year duration was motivated by the time in which atmospheric responses to volcanic impacts are often studied [69]. Early tests of the CRASSULA model configuration (not shown) and some of the results shown here indicate that year one contains some influence from the previous injection while year two is influenced by the current injection. This means that year one can be used to look at atmospheric responses to different injection scenarios with careful consideration about the composition of CRASSULA (continuous injection at different locations every two years) and year two may provide unique details about atmospheric responses to SAI scenarios shortly after deployment. For this reason, to identify regional atmospheric response shortly after a single actor deployment, this work focuses on the second year after deployment and uses year one only for global averages.

Another decision about the design of CRASSULA was to identify which variables and on what temporal resolution to save. With the 1,784 years simulated in CRASSULA (originally 1,856 intended years) and limited storage space, there had to be extra consideration about what variables would make CRASSULA most useful to answer future questions about single actor deployments. Since there is no way to foresee all ways in which a dataset such as CRASSULA will be useful, the variables saved from CRASSULA followed the recommendation of Richter et al., (2022), which contains a table of the minimum recommended model output for an ARISE-SAI simulations (Table A1 in Appendix A).

The model configuration used to simulate CRASSULA was also an important design choice, as it affects how this data can be used, how it can be compared to existing data, and how many resources were necessary to simulate CRASSULA. Some decisions were straightforward such as using the middle atmosphere component of CESM2-WACCM6 as this component is cheaper than the full CESM2-WACCM6 atmosphere component but still resolves important stratospheric chemistry necessary for SAI research [70]. A less straightforward decision was the choice to simulate CRASSULA with a data ocean rather than a dynamic ocean. There were three reasons this decision was made. First, the data ocean is marginally cheaper than a dynamic ocean, reducing the necessary resources to simulate CRASSULA. Second, not sampling the ocean states meant fewer samples needed to train the machine learning emulator. Lastly, the ocean response to external forcing is slow due to larger specific heat content and so the ocean response to injections sampled in CRASSULA would be minimal [71]. While some studies show that large volcanic eruptions can have atmospheric impacts decades after eruptions due to the slow response of the ocean [72], the purpose of CRASSULA is to study the atmospheric responses shortly after deployment, not the long lived impacts of aerosols in the stratosphere.

The final choice for the model configuration is deciding what historical forcing will be used in CRASSULA. The historical forcing impacts model details such as background emissions and land use. It was important that background emissions in CRASSULA stay constant so that atmospheric responses can be more easily attributed to the SAI scenario rather than the changing background state. The perpetual year 2000 is selected to mimic deploying SAI alongside greenhouse gas emissions.

3.2.2 Challenges while Simulating CRASSULA

CRASSULA was run on the Cheyenne supercomputer leading up to its decommission at the end of 2023. After seven years of service, Cheyenne was operating unreliably causing active jobs to fail randomly. This made finishing the simulations by the end of 2023 challenging. The failing jobs extended the timeline causing some data to unknowingly hit the maximum file life and be

deleted. These problems resulted in 72 years of the 1,856 intended years of CRASSULA to be lost. This equals a loss of 3.8% of the data and a total of 36 injection scenarios. Figure 3.1 shows the missing injection locations for the four injection rates colored by how many times each is sampled in CRASSULA. Note that despite some missing data, all injection scenarios are sampled six or more times with most scenarios sampled more than seven times. These gaps can be supplemented by additional runs, but there are currently no plans to do so.

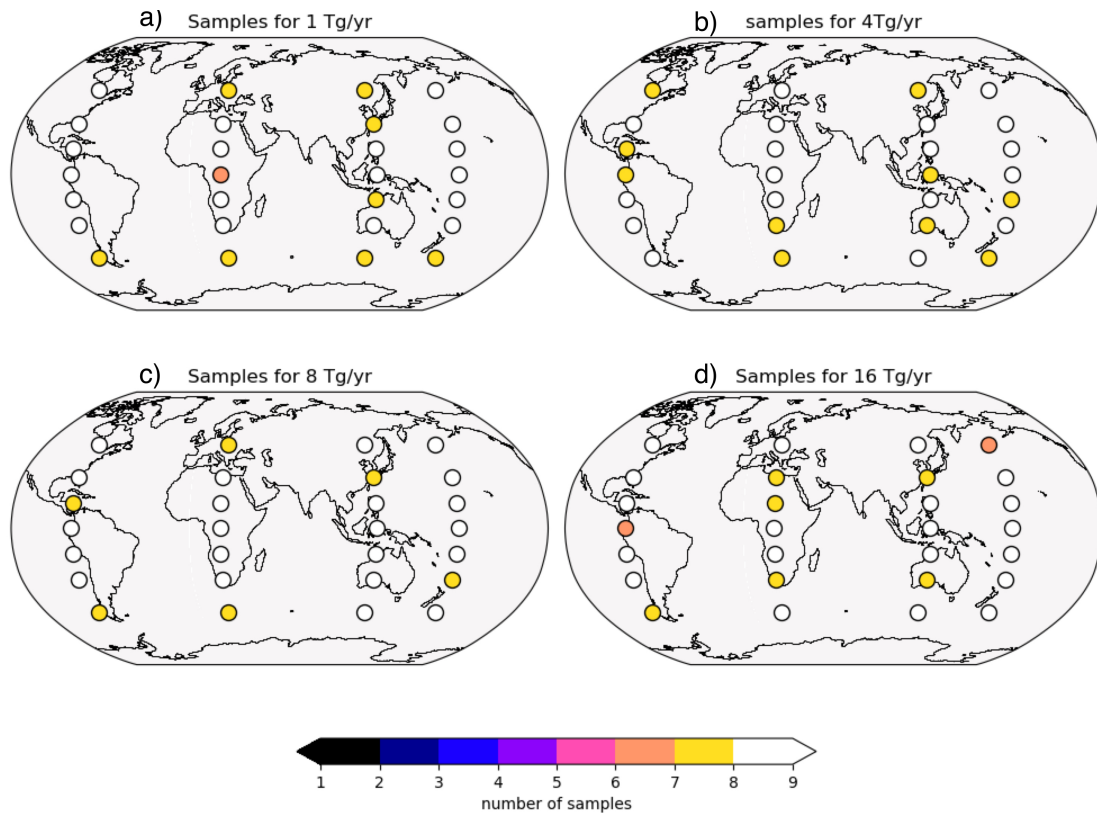


Figure 3.1: Shows the SAI locations in CRASSULA along with how many times each location is sampled for 1Tg/yr (a), 4Tg/yr (b), 8Tg/yr (c), and 16Tg/yr (d).

3.3 Composites and Results

With the unique composition of CRASSULA established, sections 3.3.1 through 3.3.4 review the atmospheric responses shortly after an SAI deployment. This study focuses on injections at 30°N, 30°S, and at the equator to represent single injections in each hemisphere and at the equa-

tor. The four longitudes sampled in CRASSULA are all used in the composites to increase the number of samples and reduce possible impacts from internal variability. Grouping scenarios with injections at different longitudes but the same latitude is reasonable, as aerosol are evenly mixed zonally by the second year after deployment, therefore changes in longitude have very little impact on atmospheric responses on this timeframe [17, 73].

3.3.1 AOD and Net Flux at the Surface

SAI cools the planet by reflecting incoming shortwave radiation. This is represented in Figure 3.2 and Figure 3.3 as increased aerosol optical depth in the hemisphere with the SO_2 injection and a reduction in net shortwave radiation at the surface. The AOD and net surface shortwave radiation anomalies also change linearly with changes to the SO_2 injection rate.

It should be emphasized that, while most of the response in AOD is located in the hemisphere of the injection, the student t-test used to stipple the panels indicate significant values in the opposite hemisphere. While significant values (stippling) do not confirm that a value is or is not a response of the current injection, they do offer evidence that the anomaly was unlikely to have occurred by chance alone. In this case, the stippling is likely caused by the previous injection. As previously mentioned, CRASSULA injects at the same location for two years when it then instantly switches to a new location and rate, therefore, there is always an injection occurring. This is important for training the machine learning emulator (Chapter 4), but may appear as lingering impacts in this study. The stippling in the hemisphere opposite the hemisphere with the injection in Figure 3.2a-d and Figure 3.3a-d are likely a result of impacts from previous injections as, on this timeframe, aerosols typically stay confined to the hemisphere where the injection occurred [17, 74]. Additional evidence of this is that the net shortwave flux anomaly at the surface shown in Figure 3.2f-i and Figure 3.3f-i indicate that most of the radiative impact is contained to the hemisphere of the injection.

Figure 3.2e, j and Figure 3.3e, j show the globally averaged anomalies for AOD and net shortwave flux at the surface. The mean across all data is shown in the dark purple (year one) and

dark green (year two) while globally averaged anomalies for each individual ensemble member is shown in the light purple (year one) and light green (year two). The ensemble mean shows the average global mean response while the global means of individual ensemble members represent a range of possible climate states following a specific injection scenario. The global mean values shown in Figure 3.2e and Figure 3.3j show a reduction in the spread of possible mean climate states from year one to year two. This is likely again due to the continuous injection. Since the injections sampled in CRASSULA are randomized, each sample is coming from a different climate state. By year two, the response to the current injection dominates in both AOD and net shortwave flux at the surface.

Figure 3.2 and Figure 3.3 provide context for how the novel CRASSULA simulation behaves to the continuous randomized injections. Now, we use the same steps to make composites of temperature, precipitation, and ozone.

3.3.2 Temperature

The predominant goal of SAI is to cool or stabilize the planet to counteract warming from additional greenhouse gases in the atmosphere and studies have shown that the cooling response depends on the deployment strategy [63, 65, 75]. The large sample of injection scenarios in CRASSULA provides a unique opportunity to study the near surface temperature response to SAI shortly after deployment. Figure 3.4 shows the near surface temperature anomaly in year two for injections at 30°N and 30°S. In these figures, the near surface temperature shows little change over the oceans. Since there is a measurable response in the net shortwave flux at the surface (Figure 3.2f-i and Figure 3.3f-i) and a temperature response over land, the data ocean is likely modulating the signal over the oceans. Even with a dynamic ocean, the surface ocean is slower to respond to external forcing and so there would likely be little to no response in the ocean on these timescales [76, 77]. The smaller response in near-surface temperature over the ocean is apparent in the global mean values in Figure 3.4e and j where the response in the Northern Hemisphere is larger (note the different y-axis).

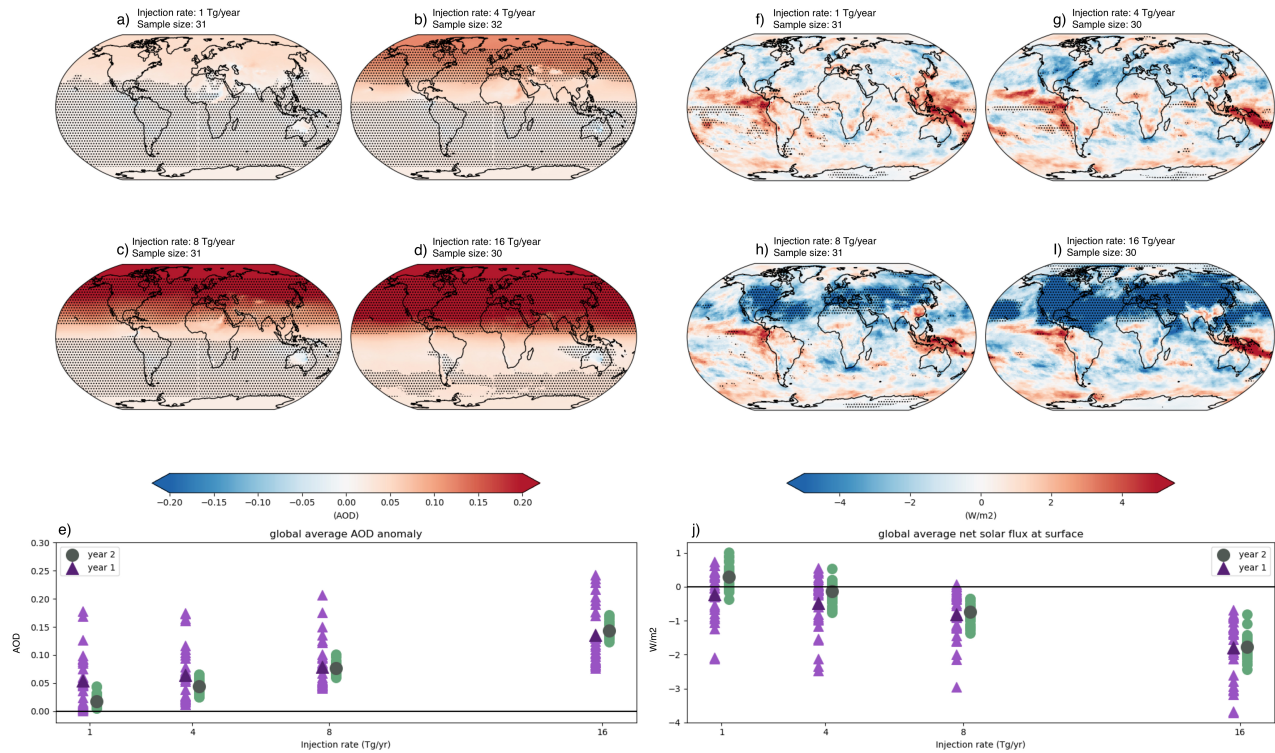


Figure 3.2: Average AOD anomaly (a-e) and average net shortwave flux at the surface (f-j) for SO₂ injections at 30°N. The injection scenarios are split into four injection rates 1Tg/yr (a, f), 4Tg/yr (b, g), 8Tg/yr (c, h), and 16Tg/yr (d, i) and stippling represents a p-value of less than 0.05. Panels (e) and (j) are global average anomalies for year one (purple triangle) and year two (green circle). Dark purple and dark green indicate ensemble means and light green and light purple show global mean anomalies for each ensemble member. Black horizontal line denotes zero.

The globally averaged values in Figure 3.4e and j show larger variability in possible global average temperatures in the second year after deployment. Injections in the Northern Hemisphere increase the likelihood of cooler mean climate states. Injections in the Southern Hemisphere increase the likelihood of both warmer and cooler mean climate states. The warmer mean climate states present during Southern Hemisphere injections occur due to increased temperatures over Europe and Asia (Figure 3.4f-i). This is likely due to the SO₂ in the stratosphere of the Southern Hemisphere causing stratospheric heating and impacting the meridional temperature gradient from the equator to the polar region of the Northern Hemisphere. Previous studies have shown this change in the temperature gradient increases the momentum flux poleward, strengthening the eddy driven jet [64].

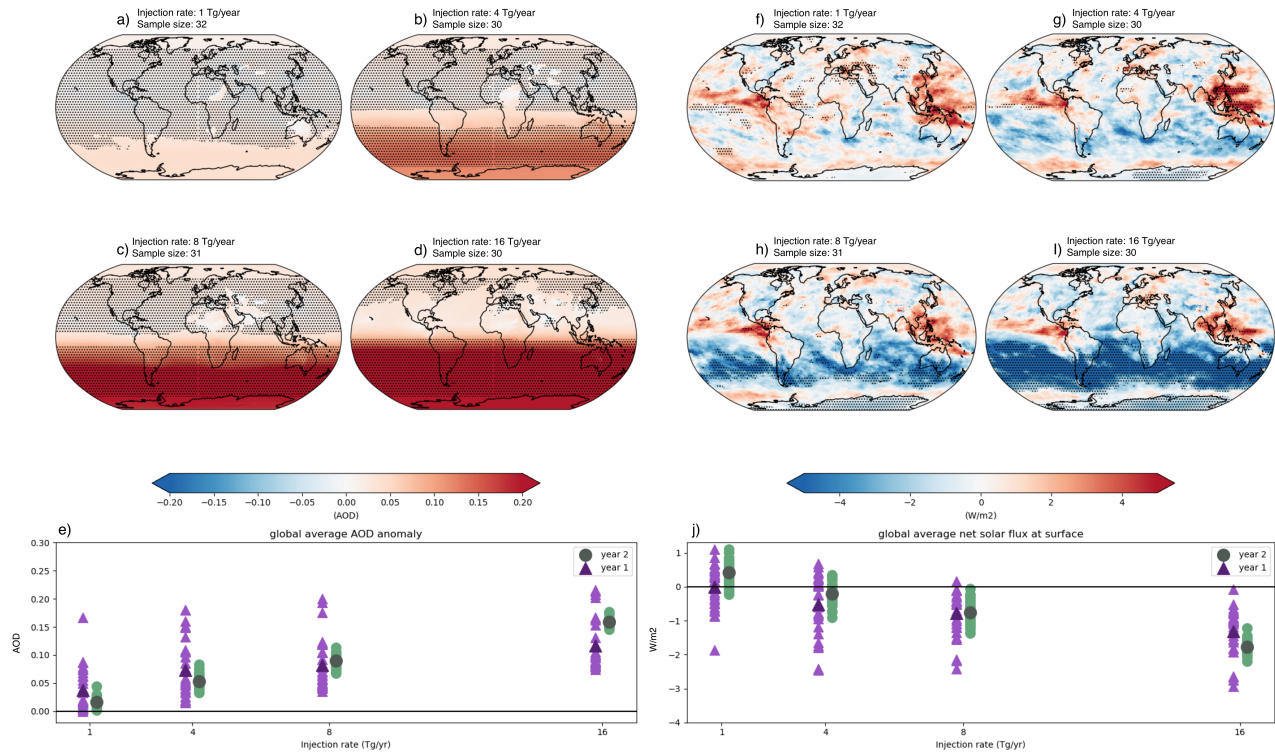


Figure 3.3: Average AOD anomaly (a-e) and average net shortwave flux at the surface (f-j) for SO₂ injections at 30°S. The injection scenarios are split into four injection rates 1Tg/yr (a, f), 4Tg/yr (b, g), 8Tg/yr (c, h), and 16Tg/yr (d, i) and stippling represents a p-value of less than 0.05. Panels (e) and (j) are global average anomalies for year one (purple triangle) and year two (green circle). Dark purple and dark green indicate ensemble means and light green and light purple show global mean anomalies for each ensemble member. Black horizontal line denotes zero.

Bednarz et al., (2022), found a similar response looking over longer timeframes. That work also identifies a similar response in the Southern Hemisphere to injection in the Northern Hemisphere, however, that signal is not present in the near surface temperature composites in Figure 3.4a-d. This indicates that SAI deployments in the Southern Hemisphere can impact circulation and near surface temperature in the Northern Hemisphere as soon as two years after deployment.

3.3.3 Precipitation

Moisture content in the atmosphere is expected to increase as a result of warming induced by additional greenhouse gases [78]. Thus, SAI is expected to reduce global mean precipitation along with reducing global mean temperatures. However, SAI does not perfectly reverse the impacts of greenhouse gases on precipitation [10], some injection scenarios have been documented to reduce

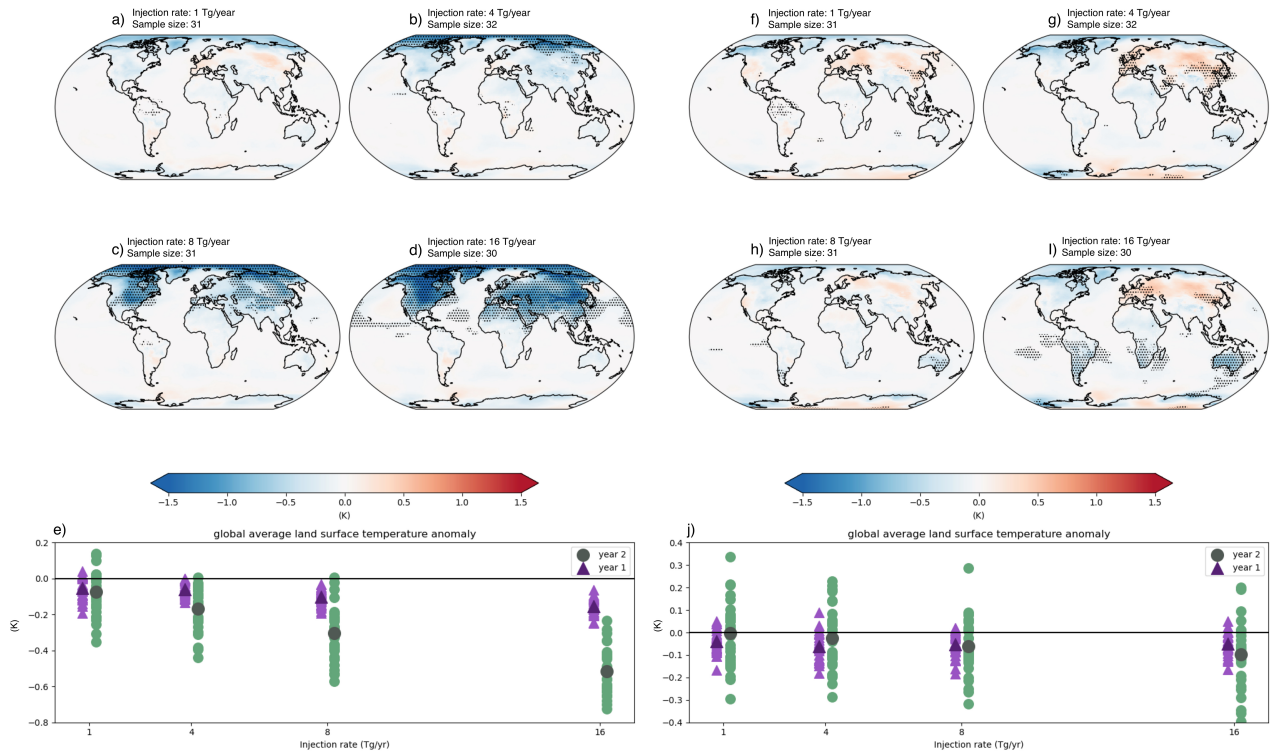


Figure 3.4: Average near surface temperature anomaly for SO₂ injections at 30°N (a-e) and 30°S (f-j). The injection scenarios are split into four injection rates 1Tg/yr (a, f), 4Tg/yr (b, g), 8Tg/yr (c, h), and 16Tg/yr (d, i) and stippling represents a p-value of less than 0.05. Panels (e) and (j) are global average anomalies for year one (purple triangle) and year two (green circle). Dark purple and dark green indicate ensemble means and light green and light purple show global mean anomalies for each ensemble member. Black horizontal line denotes zero. Note the different y-axis on panels (e) and (j).

global mean precipitation too much [31] or drive unique regional responses [79, 80]. Calculating average precipitation anomalies based on injection scenarios can quantify if precipitation responds shortly after deployment and by how much.

Figure 3.5 shows the precipitation anomaly for year two for injections at 30°N (a-e) and 30°S (f-g). Due to precipitation having large regional variability [81], there are fewer notable and significant anomalies for precipitation compared to near surface temperature. The global mean anomaly values in Figure 3.5e-j shows very little change across injection rates, however, the variability in possible global mean precipitation from year one to year two does greatly change. Similar to near surface temperature, the variability in year two is greater than the variability in year one.

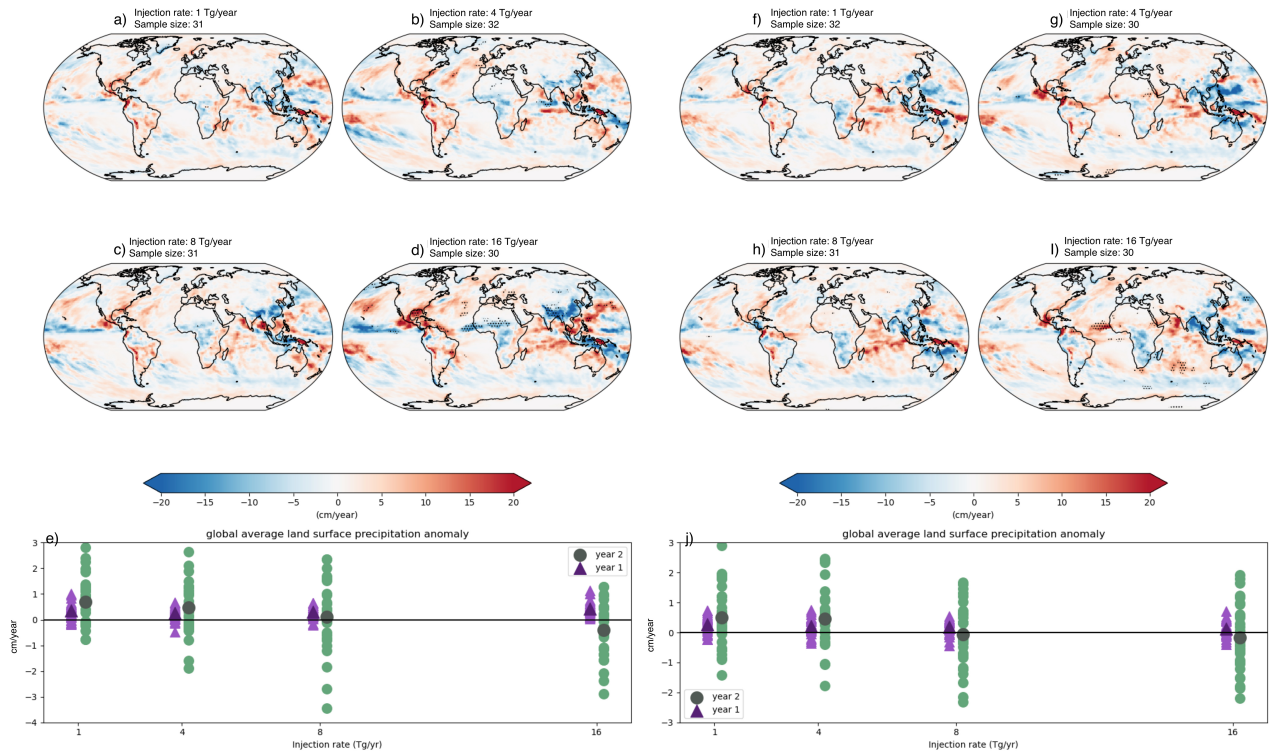


Figure 3.5: Average precipitation anomaly for SO₂ injections at 30°N (a-e) and 30°S (f-j). The injection scenarios are split into four injection rates 1Tg/yr (a, f), 4Tg/yr (b, g), 8Tg/yr (c, h), and 16Tg/yr (d, i) and stippling represents a p-value of less than 0.05. Panels (e) and (j) are global average anomalies for year one (purple triangle) and year two (green circle). Dark purple and dark green indicate ensemble means and light green and light purple show global mean anomalies for each ensemble member. Black horizontal line denotes zero. Note the different y-axis on panels (e) and (j).

While precipitation has little response to single hemispheric injections in this timeframe, precipitation does show measurable response to equatorial injections shortly after deployment. As previously mentioned, equatorial injections have been shown to reduce equatorial precipitation. Previous studies show that this is due to heating of the lower stratosphere [82, 83] which causes a stalling of the quasi-biennial oscillation [84].

Figure 3.6e highlights a lag in global precipitation response as almost no change in global mean precipitation occurs during year one for any injection rate. For year two there are linearly larger decreases in global mean precipitation in response to increases in injection rate. In global mean precipitation, an injection of 4Tg/yr or larger is needed to see a reduction in global mean precipitation in year two. However, the global mean precipitation values for individual samples

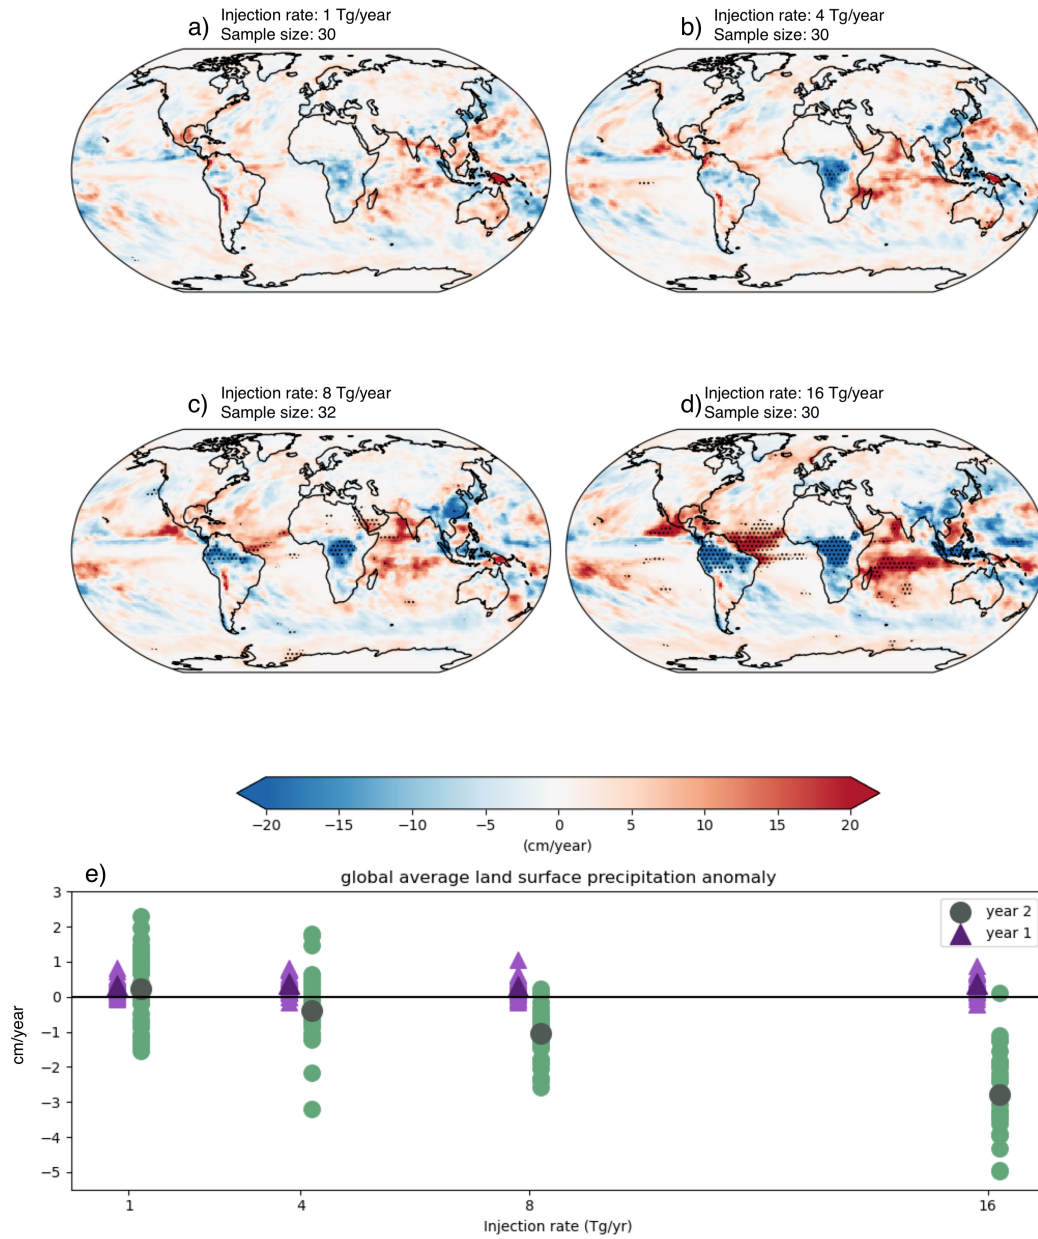


Figure 3.6: Average precipitation anomaly for SO₂ injections at the equator. The injection scenarios are split into four injection rates 1Tg/yr (a), 4Tg/yr (b), 8Tg/yr (c), and 16Tg/yr (d) and stippling represents a p-value of less than 0.05.. Panel (e) is global average anomalies for year one (purple triangle) and year two (green circle). Dark purple and dark green indicate ensemble means and light green and light purple show global mean anomalies for each ensemble member. Black horizontal line denotes zero.

show increased variability in possible mean climate states in year two with injections of only 1 Tg/yr. This includes increased occurrences of mean climate states with increases in precipitation and decreases in precipitation.

3.3.4 Ozone

Sulfate aerosols can impact ozone concentrations by catalyzing heterogeneous chemistry and impacting ozone transport [14, 36, 85, 86]. These impacts on ozone concentrations depend on the region and the injection scenario. For example, injection scenarios that strengthen the polar vortex will drive a decrease in ozone concentration at the high latitudes due to a reduction of poleward ozone transport [85]. Understanding how an injection scenario may impact ozone concentrations is important since some injection scenarios could slow down the recovery of the stratospheric ozone layer. Ozone responses shortly after a single actor deployment are uncertain and CRASSULA provides an opportunity to explore how quickly these changes occur.

When investigating ozone responses shortly after a single actor deployment, it is important to note that the impact SAI has on stratospheric ozone depends on the deployment year. The process of ozone loss happens via heterogeneous chemistry and depends on the chloroflourocarbon (CFC) load in the stratosphere. Since CRASSULA is simulated with continuous year 2000 emissions and the CFC load has continued to decrease each year in the real atmosphere [87], the magnitude of ozone impacts shown here will be larger than the ozone responses to future single actor deployments. CRASSULA can still provide information about the pattern of ozone response and how quickly ozone responds to single actor deployments. For this reason, only ozone response to injections of 16Tg/yr and global averages are shown in Figure 3.7 so that attention can be placed on the pattern of ozone response and changes from year one to year two.

The ozone responses shown in Figure 3.7 indicate that ozone responses can occur as quickly as two years after deployment. The patterns of ozone loss are similar, but weaker in magnitude, to those found in Bednarz et al. (2023) when looking at ozone responses over a 15-year period. These responses are summarized briefly here. Equatorial injections result in a layered ozone response in the stratosphere where ozone decreases in the lower stratosphere, increases in the middle stratosphere, and decreases in the upper stratosphere. This tiered response occurs due to changes in tropospheric circulation (equatorial upwelling and the Brewer-Dobson circulation) and changes

to heterogeneous chemistry [85]. When averaged across the entire stratosphere and troposphere, an equatorial injection drives an overall decrease in stratospheric ozone (Figure 3.7b, d, f).

In a single hemisphere injection there is a decrease in stratospheric and tropospheric ozone in response to ozone destruction in the stratosphere and less ozone available for transport at the high latitudes from the stratosphere into the troposphere. The opposite is true for the hemisphere opposite of the injection where stratospheric ozone increases. Previous research shows that as ozone is created in the tropical stratosphere, the stronger northern hemisphere component of the Brewer-Dobson Circulation carries the new ozone northward [85].

Given the long recovery time of stratospheric ozone [88] some features of 3.7 are likely lingering responses from previous injections. For example, all panels of zonal mean ozone show a negative ozone anomaly in the hemisphere opposite of the injection. Given that aerosols predominantly stay contained in the hemisphere of the injection on this timeframe [17,74], these reductions are likely not a response to the current injection. Furthermore, the feature does not reach statistical significance ($p > 0.05$). While significant values (stippling) do not confirm that a value is or is not a response of the current injection, they do offer evidence that the anomaly was unlikely to have occurred by chance alone. Overall, we find that ozone responds quickly to SAI. Single actor injections can impact stratospheric ozone, tropospheric ozone, and circulations as soon as two years after deployment.

3.4 Conclusion

Single hemisphere or equatorial injections come with a number of undesirable impacts [30,31,61]. However, this does not mean that single actor scenarios are not possible in the future. As technology improves, making climate intervention methods possible and cheaper, and climate impacts become more extreme, a single actor may be motivated to deploy SAI. Given the possibility of undesirable consequences, it is crucial that these scenarios receive attention so that responses shortly after a single actor deployment are understood, can be properly communicated, and any necessary mitigation strategies can be implemented. CRASSULA provides a unique opportunity

to investigate atmospheric responses to different injection scenarios shortly after deployment. The results highlighted here indicate that some atmospheric responses to SAI scenarios can occur as quickly as one to two years after deployment, depending on the rate of injection. Additionally, in smaller injection rates that have no response in the ensemble mean, this work indicates that there is increased variability in possible global mean precipitation and global mean near surface temperature states after a deployment. These atmospheric responses may be short lived but have major impacts on communities.

These results call attention to important nuances when considering future climate states after a single actor deployment. Even if a single actor injection were halted quickly after deployment, it does not mean that the atmospheric responses will be minimal. Additionally, the longer any single actor deployment continues, the larger probability of significant atmospheric responses. This work highlights why these scenarios, although not favorable, cannot be ignored.

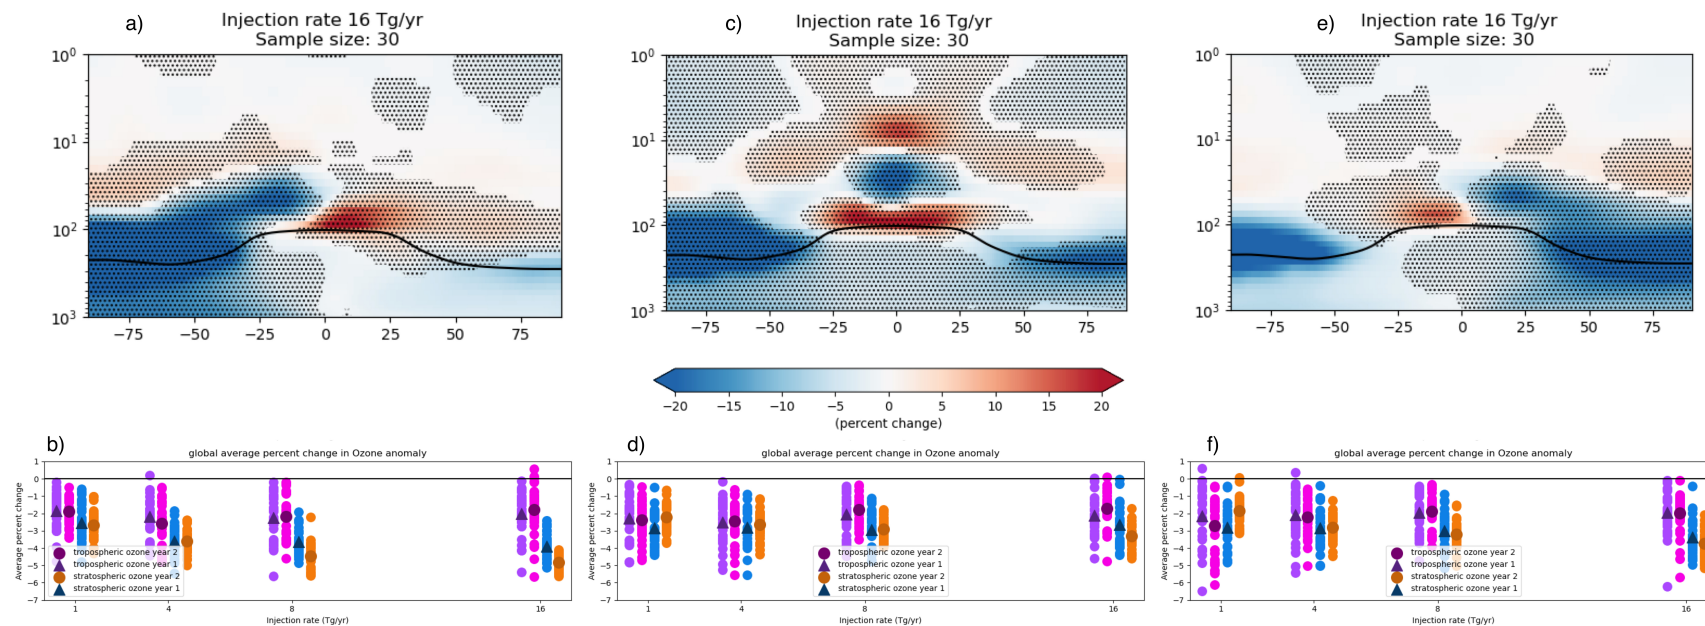


Figure 3.7: Average percent change in ozone for SO₂ injections at 30°S (a,b), equator (c,d), and 30°N (e,f). Stippling represents a p-value of less than 0.05. Panels b, d and, f are global average anomalies for year one (triangle) and year two (circle). Pink and purple are averages with tropospheric ozone. Blue and orange are averages with stratospheric ozone. Dark purple, pink, blue, and orange indicate ensemble means and light purple, pink, blue, and orange show global mean anomalies for each ensemble member. Black horizontal line denotes zero.

Chapter 4

Emulating the Temperature Response to Stratospheric Aerosol Injection Scenarios

4.1 Introduction

Earth system models are a powerful tool to simulate the response to hypothetical climate intervention strategies, but limited resources constrain which SAI scenarios can be simulated [19]. To help scientists decide which possible futures to simulate in an Earth system model, scenario design methods are often used. Scenarios constrain what climate futures are simulated by directing efforts towards plausible futures that can inform decision-making under uncertainty and are used to understand possible climate futures in response to decisions [89, 90]. The focus on past and future decisions in scenario design means that what is considered a plausible scenario changes based on the person designing the scenario and with changes in policies and society.

Scenario design can direct the planning process for creating future scenarios surrounding the deployment of SAI. Designing scenarios enable scientists to evaluate a range of potential climate futures, factoring in changes to emissions, culture, and political landscapes. Scenarios are especially beneficial when evaluating plausible climate futures that involve single actor scenarios, scenarios where an SAI deployment is in response to a previous SAI deployment, or for an SAI deployment that is in response to an extreme weather event [24, 91].

To inform future decisions about SAI deployment, it is important to have a sufficient understanding of the range of possible climate responses under SAI. Since running an Earth system model can be slow and expensive, emulators have been proposed as a tool that can cheaply and quickly explore a large range of climate responses to SAI scenarios. To date, there are emulators designed to emulate the impact of terminating an SAI deployment [92] and a linear emulator

to investigate how reducing the solar constant (reducing incoming shortwave radiation) impacts temperature and precipitation [93].

This work introduces an initial version of a machine learning emulator which will be updated for a future publication. The emulator is designed to explore SAI deployment scenarios by predicting the near surface temperature response to a multi-year SAI scenario. By generating data at one year increments, the generated data can be used for considering the motivation of subsequent actors by being able to sample different internal variability patterns and possible SAI impacts. After using the machine learning SAI emulator to design the deployment scenario, the parameters can then be simulated in an Earth system model.

This machine learning emulator is created by independently training multiple interpretable models that together create the complete SAI emulator. This novel method for creating a machine learning emulator allows for greater interpretability so that a user can better understand why a prediction was made. In total, this machine learning emulator is comprised of four machine learning models that are trained on simulations from CESM2 [38], including the novel dataset titled, CRASSULA. These four machine learning models independently predict internal variability, temperature responses to the total CO₂ in the atmosphere, aerosol optical depth (AOD) responses to SAI, and temperature responses to changes in AOD. Details about how each of these models are designed and trained is found in Section 4.3, with additional information in Appendix B.

4.2 Data

To train the machine learning emulator, annual mean temperature, annual mean AOD, and annual mean global CO₂ amount in parts per million (PPM) are used from five existing CESM2 simulations and the novel CESM2 simulation named CRASSULA (Table 4.1). Details about the simulations are found in the following sections.

Perpetual 1850 Control Simulations: Two perpetual 1850 pre-industrial control simulations are used. The first simulation is run with the Whole Atmosphere Community Climate Model version 6 (CESM2-WACCM6; [66, 94]) with a 0.9 latitude by 1.25 longitude grid (piControl-

Table 4.1: Data used to train machine learning models

Short Name	Reference
piControl-WACCM	[94]
piControl-CAM	[95]
SSP2.4-5	[96]
LE-HIST	[97]
ARISE-SAI-1.5	[37]
CRASSULA	documented in Chapter 3

WACCM). The second perpetual 1850 pre-industrial control simulation is run with Community Atmosphere Model version 6 (CAM6; [98]) with a 0.9 latitude by 1.25 longitude grid (piControl-CAM). Two different simulations were utilized rather than just one to provide additional samples of 1850 internal variability. While these simulations utilize different atmospheric components of the CESM2 model, both were selected for specific goals. The piControl-WACCM simulation uses the same atmospheric component as CRASSULA and the piControl-CAM is larger and provides additional training years.

Shared Socioeconomic Pathways Middle of the Road Rcenario: The Shared Socioeconomic Pathway 2-4.5 (SSP2-4.5; [37, 96]) simulation is a future emission scenario, often referred to as the “middle of the road” scenario [99]. This simulation uses CESM2-WACCM6 with 0.9 latitude by 1.25 longitude grid. SSP2-4.5 has 10 ensemble members, where five members simulate years 2015-2100 [96] and five members simulate years 2015-2070 [37].

Historical CESM2 Large Ensemble: The CESM2 Large Ensemble is a 100 member ensemble simulated with a 0.9 latitude by 1.25 longitude grid (CESM2-LE; [97]) and uses CAM6 as the atmosphere model. The ensemble members are forced with historical forcing from 1850 to 2014. To train and evaluate the machine learning emulator, only 10 members from CESM2-LE are used so as to not overwhelm the training data with data simulated with the CAM6 atmospheric component.

ARISE-SAI-1.5: The ARISE-SAI-1.5 ensemble is designed to simulate a plausible coordinated global deployment of SAI where SAI is deployed to maintain global mean temperature at

about 1.5°C above pre-industrial temperature while minimizing impact to global temperature gradients [37]. These goals are accomplished through the use of a control algorithm which determines how much sulfur dioxide to inject into the stratosphere across the four locations (15°S, 15°N, 30°S, 30°N and 180°E at height 21.5 km; [29, 37, 39]). ARISE-SAI-1.5 uses CESM2-WACCM6 with 0.9 latitude by 1.25 longitude grid. It has 10 ensemble members simulating the years 2035-2070. This simulation provides the emulator information about simultaneous injections. Refer to Richter et al., (2022) and Kravitz et al., (2017) for more details about the ARISE-SAI-1.5 ensemble or the control algorithm.

CRASSULA: CRASSULA was introduced in detail in section 3.2. For training the machine learning emulator, CRASSULA required a bias correction as the average land surface temperature calculated from the CRASSULA control simulations is about 1K cooler than year 2000 in the CESM2-LE or the historical simulations that accompany SSP2-4.5 simulation. CESM2-WACCM6 has been documented to have a cold bias in the global surface temperature compared to the CESM2-WACCM6 full atmosphere component [70].

To correct the bias, the average difference between the SSP2-4.5 historical simulation and the CRASSULA control is determined. Ten years (1995-2005) from the SSP historical simulations are used to reduce noise from internal variability. The bias correction is defined as the ensemble mean of the SSP historical ten year average minus the average of the near surface temperature from the CRASSULA control. To adjust the near surface temperature in CRASSULA, the bias correction map is added onto each annual mean near surface temperature map.

4.3 Machine Learning Training and Validation

Four individual U-net machine learning models make up the machine learning emulator [100]: termed the internal variability machine learning model (*IVM*), the CO₂ machine learning model (*CO2M*), the AOD machine learning model (*AODM*), and the SAI machine learning model (*SAIM*) (Figure 4.1). When put together, these four machine learning models makeup the complete SAI scenario emulator. The machine learning emulator requires four user defined inputs:

Machine Learning Emulator

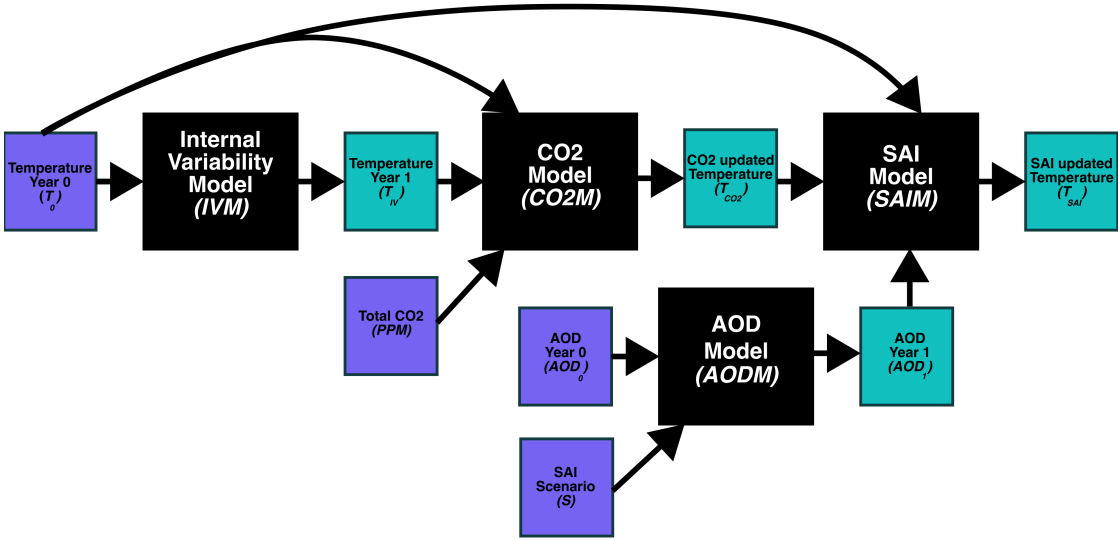


Figure 4.1: Schematic of the four machine learning models that comprise the machine learning emulator (black boxes). Inputs are indicated by purple squares and blue squares indicate inputs that are predictions from other machine learning models. Arrows represent the flow of information within this multi-model architecture.

an initial annual mean temperature map (T_0), an initial annual mean AOD map (AOD_0), a CO_2 amount in PPM (PPM), and the injection scenario (S). SAI scenarios are often identified by the location (latitude, longitude, and height) of injection along with the rate of injection. The task of the SAI scenario emulator is to predict the near surface land temperature for the following year. Future versions of the SAI scenario emulator can be trained to predict temperature over the oceans and SAI responses in other variables. Figure 4.1 shows how the four models (black boxes), the four user defined inputs (purple boxes), and how the information flows through the four U-nets to create the SAI scenario emulator (arrows).

Individual machine learning models are more interpretable as they allow for impacts (I) on the future climate state to easily be attributed to either internal variability, additional CO_2 (I_{CO_2}), or the SAI scenario (I_{SAI}). This is accomplished by calculating the difference in the predicted future state from each of the machine learning components as shown in equation 4.1 and equation 4.2.

$$I_{SAI} = T_{SO_2} - T_{CO_2} \tag{4.1}$$

$$I_{CO_2} = T_{CO_2} - T_{IV} \quad (4.2)$$

where T_{SO_2} is the annual mean temperature predicted by the SAI machine learning model, T_{CO_2} is the annual mean temperature predicted by the CO₂ model, and T_{IV} is the annual mean temperature predicted by internal variability model.

Model components are trained and validated individually. Average global mean error is reported in Figure 4.2 along with coefficient of determination (R_2); defined as the ratio of explained variance by the model to the total variance [101]. Additional validation metrics and information about the machine learning models, such as model architectures and loss functions, can be found in Appendix B.

4.3.1 Internal Variability Machine Learning Model

The atmospheric response to volcanic eruptions, the natural analog to SAI, is known to depend on internal variability [102]. Therefore, to use the SAI machine learning emulator for scenario design, it must have information about the current state of the atmosphere. For this reason, the foundation of this machine learning emulator is a model that predicts internal variability for the following year. Predicting annual mean temperature for the following year based on the annual mean temperature of the current year is a challenging problem [103]. The goal of this model is not to perfectly predict the future annual mean temperature map, but instead, provide a best guess of the following year's internal variability to build the rest of the machine learning emulator off of. While the R_2 score in Figure 4.2a is negative, the internal variability model does do better than a persistence prediction (Figure B.2 in Appendix B).

The internal variability machine learning model is trained on near surface temperature from both 1850 control simulations. Figure B.1 summarizes which simulations are used to train each machine learning model. This model takes in a centered and standardized annual temperature map as an input (T_0). The centering preprocessing removes the global mean temperature in each sample so that all inputs have the same mean. The prediction from the internal variability machine

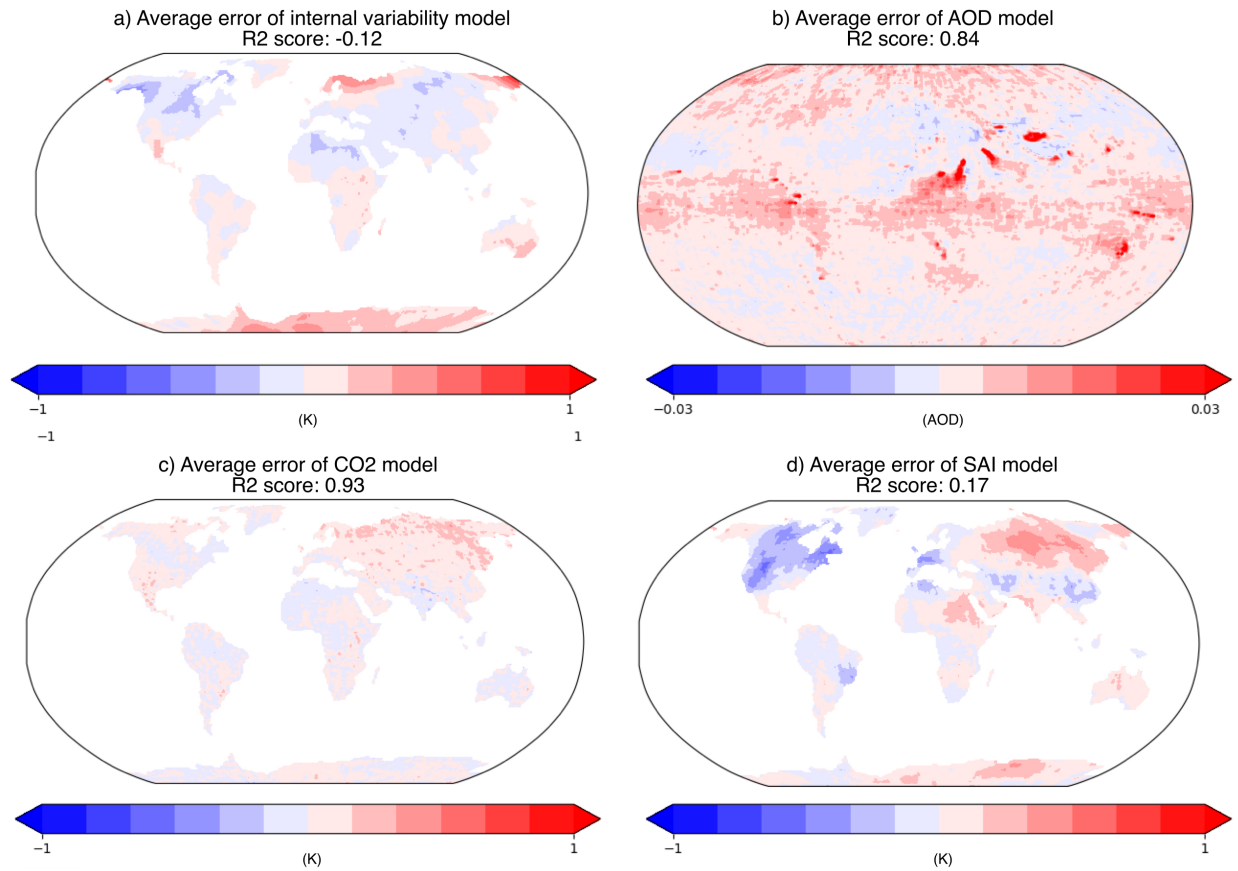


Figure 4.2: Mean absolute error the on piControl-WACCM data using the internal variability machine learning model (a), on SSP2-4.5 using the CO₂ machine learning model (b), on CRASSULA using the AOD machine learning model (c), and on CRASSULA using the SAI machine learning model (d) for respective testing data set.

Example of generated global mean land temperature for ARISE-SAI-1.5

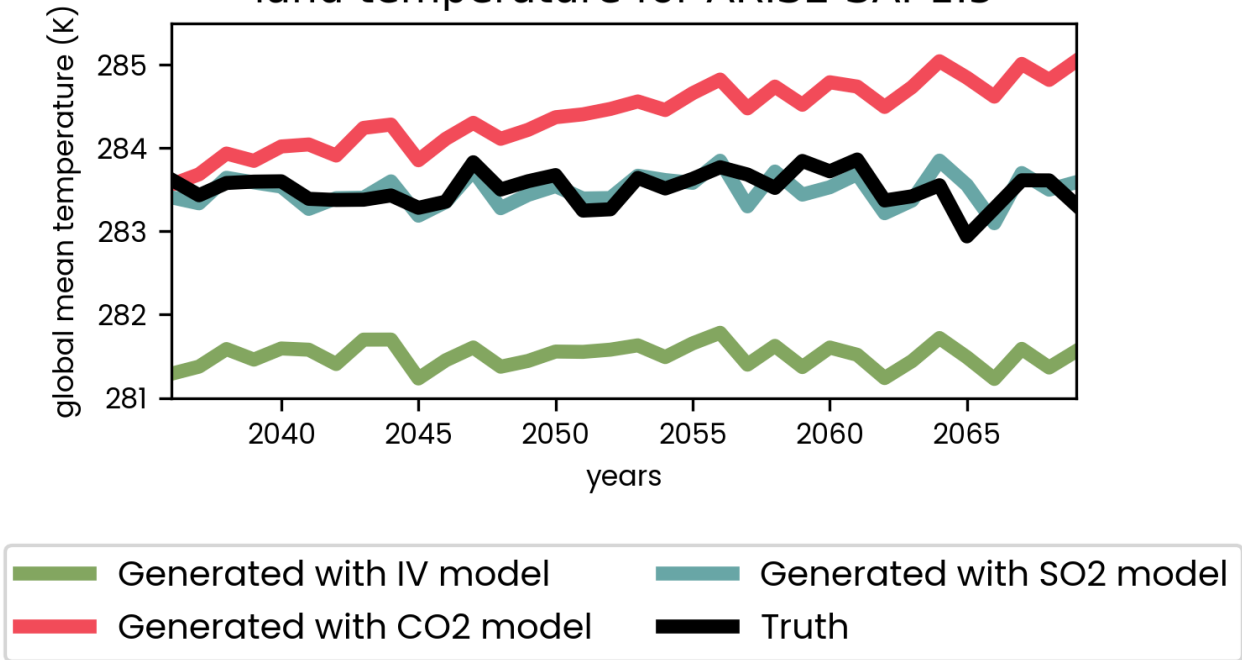


Figure 4.3: Global mean near surface land temperature from an ARISE-SAI-1.5 ensemble member (black line) compared to data generated with the machine learning models. The internal variability machine learning model predicts temperature anomaly relative to 281.5K (green line), the CO₂ machine learning model adds heating associated with the amount of CO₂ globally (red line), and the SAI machine learning model adds cooling in response to SAI scenario (blue line)

learning model predicts the following year’s temperature by including an additional layer as the final layer of the machine learning model. This forces the model to learn the change in near surface temperature from one year to the next and add it onto the input temperature.

Figure 4.3 is an example of generating ARISE-SAI-1.5 global mean near surface land temperature using the full SAI emulator. The green, red, and blue lines represent output from the internal variability, CO₂, and SAI machine learning models respectively. The prediction from the internal variability machine learning model varies around 281.5K, the average global mean near surface land temperature of the training data. This shows how the internal variability model creates a foundation for the other machine learning models to build off of.

4.3.2 CO₂ Machine Learning Model

The CO₂ model takes the initial annual temperature map (T_0), the prediction from the internal variability model, and a single value describing the global amount of CO₂ (PPM). The task of the CO₂ model is to update the map predicted by the internal variability model with warming in response to the global average CO₂. Mathematically, the CO₂ machine learning model predicts the difference between the internal variability prediction and the truth. The CO₂ model prediction is added to the prediction from the internal variability model to return a predicted future temperature updated with warming in response to the amount of CO₂ in the atmosphere.

This model is trained using the piControl-WACCM, the piControl-CAM, the SSP2-4.5 and the CESM2-LE simulations. During training, the truth of the 1850 control simulations are artificially set to zero so that when there is no additional CO₂ in the atmosphere, the machine learning model learns not to update the prediction. Conversely, when there is additional CO₂ in the atmosphere, the machine learning model learns to update the prediction.

In Figure 4.3, the red line shows the predicted global mean near surface land temperature as predicted from the ARISE-SAI-1.5 simulation. This represents how the CO₂ machine learning model takes the predicted internal variability from the internal variability machine learning model and adds the warming in response to the global amount of CO₂. Note that the T_0 input is sampled

from ARISE-SAI-1.5 as the centering processing step allows for any temperature map to be used as an input. This means the CO₂ machine learning model prediction (red line) represents SSP2-4.5 data generated by sampling data from ARISE-SAI-1.5 for the initial temperature input.

4.3.3 AOD Machine Learning Model

Training the AOD machine learning model is completed using the piControl-WACCM, the piControl-CAM, the ARISE-SAI-1.5, and the CRASSULA simulations. The design of the AOD machine learning model has two inputs, an initial state of AOD anomalies (AOD_0) and the injection scenario (S). In order to account for the four simultaneous injections in ARISE-SAI-1.5 and to allow for SAI deployment scenarios with multiple injections, the AOD machine learning model takes up to four injections. From the initial AOD anomalies and the injection scenario, this model predicts a map of AOD anomalies for the following year. Note that the AOD model does not have an additional layer as the final layer in the machine learning model but instead directly predicts the annual AOD anomalies of the following year. The mean spatial average error of the AOD model is shown in Figure 4.2c.

The AOD machine learning model is the only model of the four that is autoregressive, allowing a predicted AOD to be used as the next initial AOD input. While training autoregressive models can have many challenges, the AOD response to SAI is distinct enough that the AOD machine learning model works well when run autoregressively.

Figure 4.4 compares the average error when running the AOD machine learning model autoregressively (orange line) versus if the user were to make the same predictions for one year at a time (purple line) for 30 years chunks. Calculating errors for 30 year chunks is repeated 10 times. The shading in Figure 4.4b shows the full spread of errors and the solid line represents the mean error. Due to the distinct AOD response to SAI, the skill of the autoregressive AOD machine learning model is similar to that of the AOD machine learning model predicting a single year at a time.

The autoregressive ability of the AOD machine learning model allows the full emulator to quickly explore SAI scenarios as injection scenarios can be identified by the user and an entire

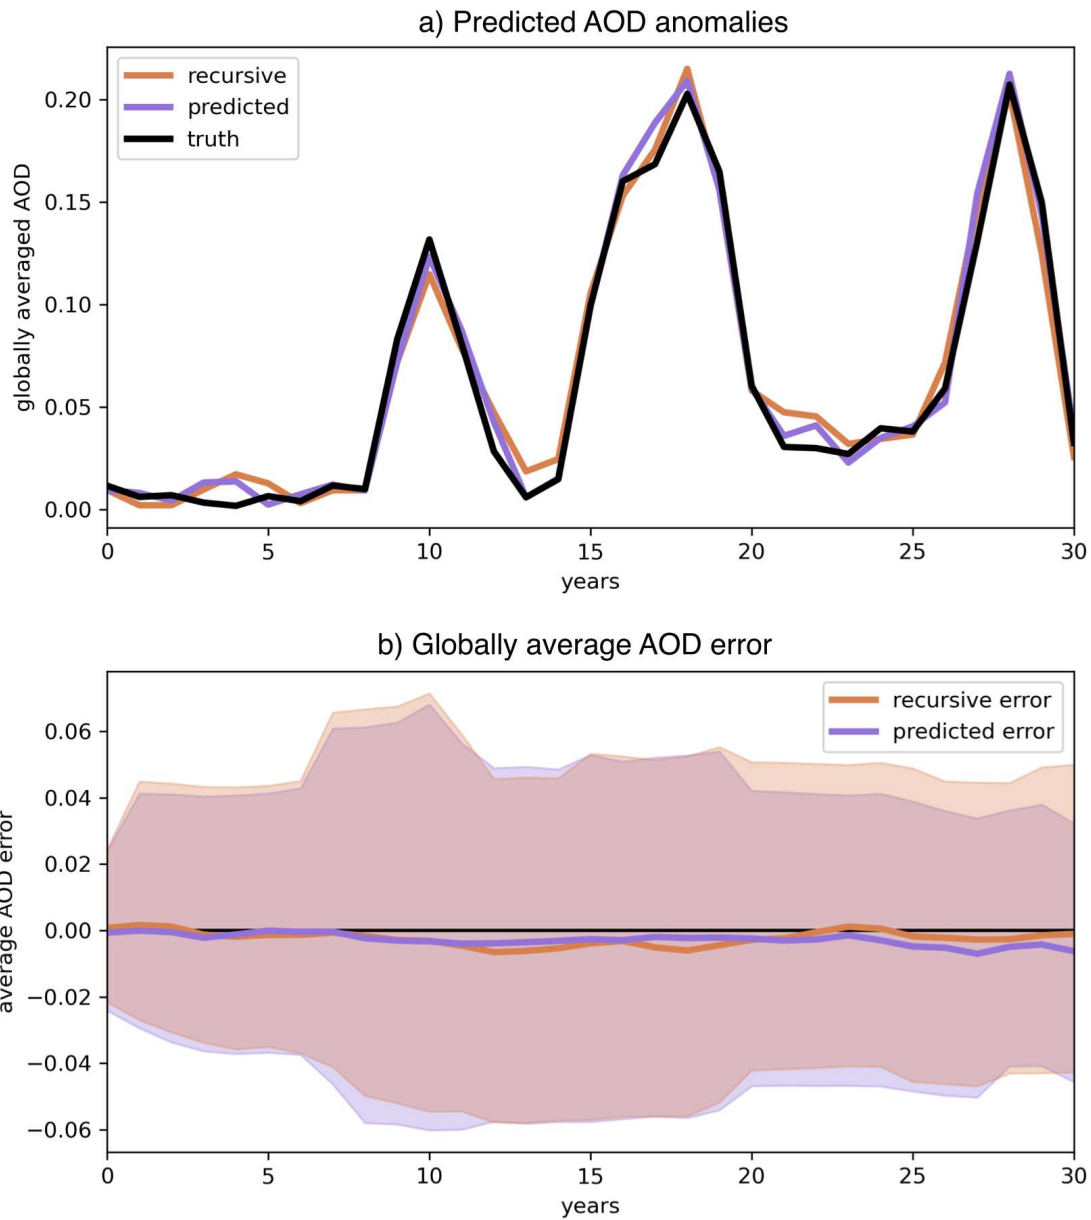


Figure 4.4: Globally average error from the AOD machine learning model using 10 sets of 30-years from CRASSULA to identify the distribution of errors. The purple line is the average error for using the AOD machine learning model to predict a single year with the shading representing the min and max error. The orange line shows the globally averaged error when using the AOD machine learning model autoregressively with the orange shading representing minimum and maximum error.

AOD timeseries can be generated. Figure 4.5 shows an example of two generated AOD timeseries using the AOD machine learning model autoregressively. Scenario one contains a single continuous 5Tg/yr injection (pink line) at 50°N, 180°E and Scenario two has the same 5Tg/yr injection with an additional second 5Tg/yr injection at 50°S, 180°E starting at year four (purple line).

For scenario 1, the globally averaged AOD remains constant across the 10 years (Figure 4.5c) with most of the AOD response constrained to the Northern Hemisphere (Figure 4.5a). Scenario 2 experiences an increase in globally averaged AOD in year five after the second deployment begins (Figure 4.5c) with greater AOD values in the Northern and Southern Hemispheres in response to injections in both hemispheres.

The locations and year of the second deployment selected for Figure 4.5 are selected here for demonstration, however, this represents how the AOD machine learning model allows the machine learning emulator to facilitate more strategically designed scenarios. Consider an experiment where a user initially generates only the first four years in Figure 4.5, representing a single actor deploying SAI at a single location continuously for four years. The SAI emulator can use this timeseries to predict the temperature response to this single point injection. The resulting temperature response can be used to consider human responses by asking questions such as, how might these temperature responses motivate additional SAI deployments? Answers to questions like this one will impact what the final six years will look like.

4.3.4 SAI Scenario Learning Model

The SAI machine learning model takes an initial annual temperature map (T_0), the prediction from the CO₂ model (T_{CO_2}), the initial state of AOD anomalies (AOD_0), and the predicted AOD anomalies (AOD_1). This machine learning model is trained using the piControl-WACCM, the SSP2.4.5, the ARISE-SAI-1.5, and CRASSULA simulations. The task of the SAI model is to update the map predicted by the CO₂ machine learning model in response to the SAI scenario by the same means as the CO₂ machine learning model. During training, the truth of the piControl-WACCM simulation and the SSP2.4 simulation are artificially set to zero so that when there is no

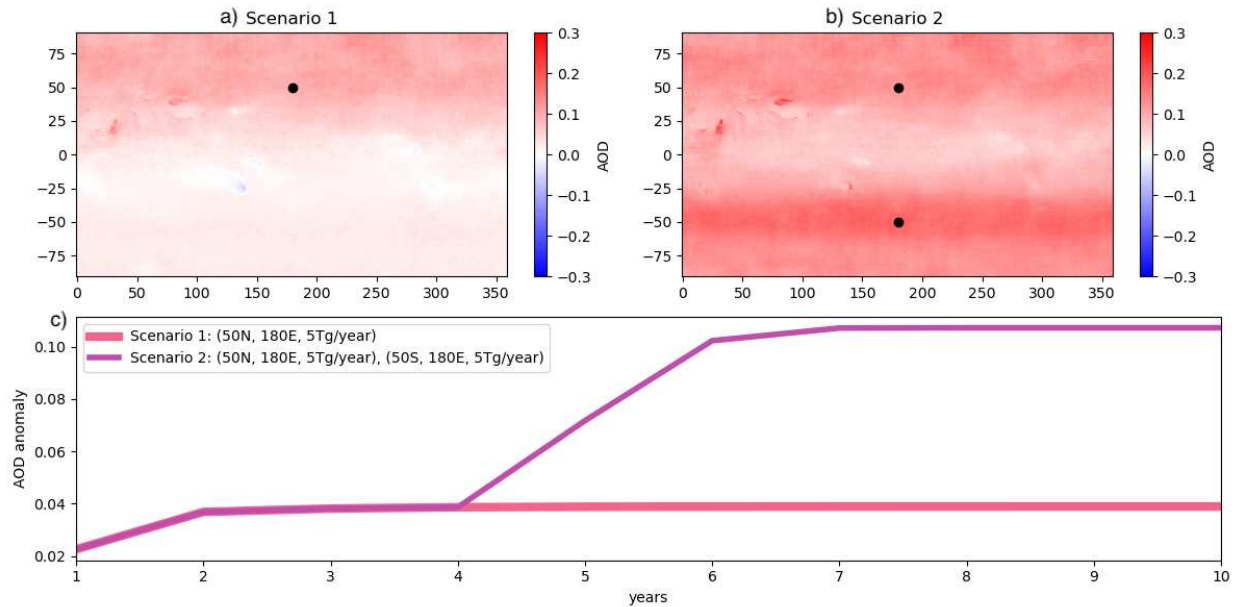


Figure 4.5: (c) The global mean change in AOD in response to a single continuous injection (pink line) compared to a scenario where a second party deploys SAI at a new location at year 4 (purple line). Panels (a) and (b) shows the AOD in year 10 for scenario 1 and scenario 2.

additional SO_2 in the atmosphere, the machine learning model learns not to update the prediction. Conversely, when there is additional SO_2 in the atmosphere, the machine learning model learns to update the prediction.

The mean spatial average error of the SAI machine learning model is shown in Figure 4.2d. The R^2 value for the SAI model is not as high as the R^2 value for the CO_2 machine learning model or for the AOD machine learning model. This is because the near surface temperature response to changes in AOD is noisier than the temperature response to CO_2 or changes in AOD to additional SO_2 in the stratosphere. We show in the next sections and in Appendix B that the SAI machine learning model can skillfully predict an average near surface temperature response to an SAI scenario.

The blue line in Figure 4.3 shows the output from the SAI machine learning model, the final output from the entire SAI emulator. Here, the blue line falls on top of the true global mean temperature (black line) from one of the ARISE-SAI-1.5 ensemble members. In summary, Figure 4.3 shows how the internal variability, CO_2 , AOD, and SAI machine learning models work together

and can be useful for scenarios with different patterns of internal variability, different levels of CO₂ emissions, and different SAI scenarios.

4.4 The Multi-Model Approach for Scenario Exploration

To use this machine learning emulator to explore SAI scenarios, the user needs to determine the four user defined inputs. The CO₂ PPM input represents the emission scenario of interest. The injection scenario relates to the injection scenario of interest. Data will need to be sampled from existing data for the initial AOD input and the initial temperature input. The initial AOD input can come from existing CESM2 data such as an 1850 control simulation, ARISE-SAI-1.5, or CRASSULA. It is recommended that an 1850 control simulation be used for the initial temperature input; however, any CESM2 data can be used.

The detailed process for generating a near surface temperature timeseries for a 10-year scenario using this machine learning emulator goes as follows.

1. Identify a 10-year injection scenario (injection rate, latitude, and longitude for the 10 years of the scenario) and the CO₂ emission scenario of interest.
2. Using an annual AOD map from existing data as the initial AOD input and the user-created injection scenario, a timeseries of AOD is generated by running the autoregressive AOD machine learning model.
3. 10 years of annual mean temperature from piControl-WACCM are used as the initial temperature inputs for the internal variability machine learning model.
4. Pass the predictions from the internal variability machine learning model and the 10 years of user-created CO₂ totals into the CO₂ machine learning model.
5. Finally, pass the prediction from the CO₂ machine learning model and the AOD timeseries into the SAI machine learning model. This will update the CO₂ prediction with the response associated with the AOD pattern.

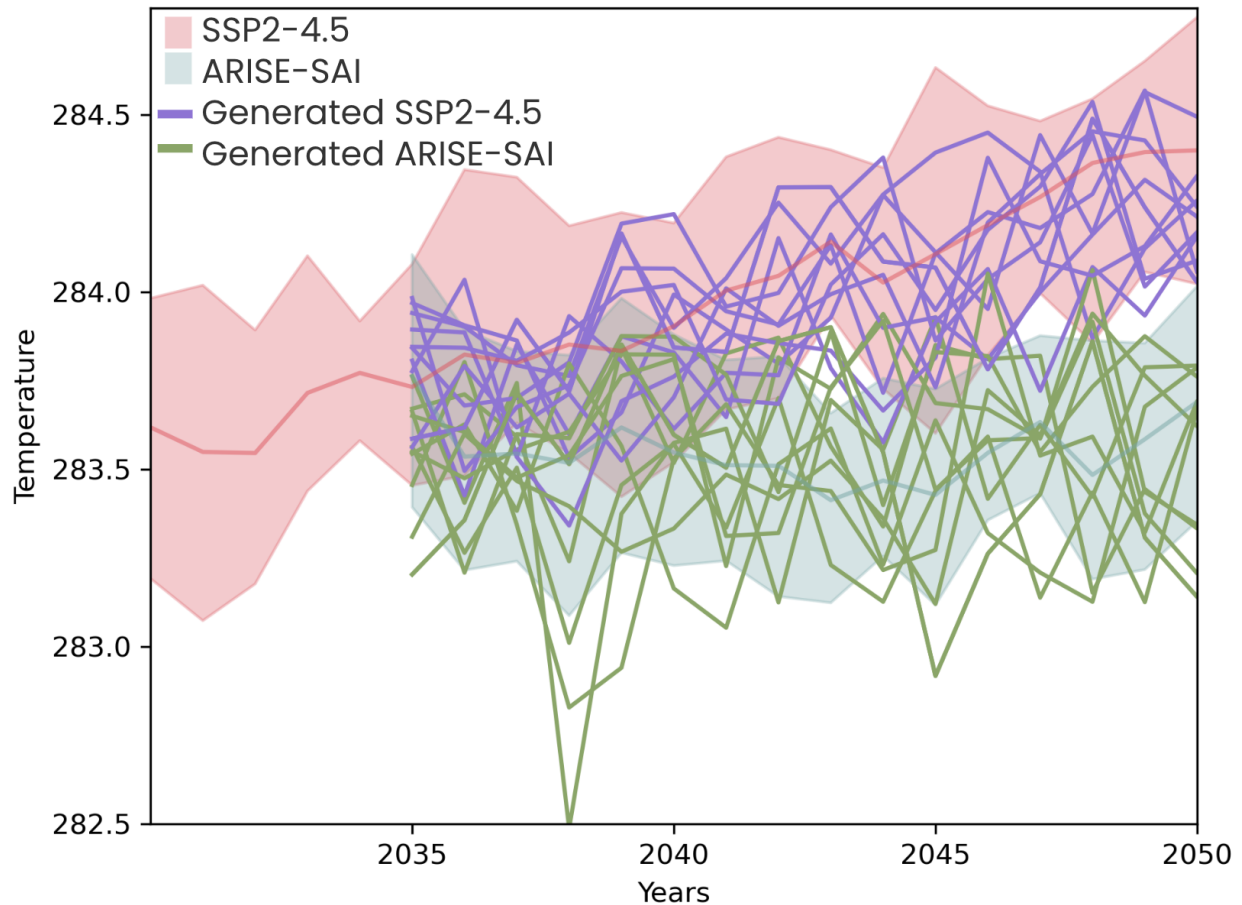


Figure 4.6: Global mean land temperature for SSP2-4.5 (red shading), ARISE-SAI-1.5 (green shading), data generated by the SAI emulator to replicate SSP2-4.5 (purple lines), and data generated by the SAI emulator to replicate AIRSE-SAI-1.5 (green lines). Shaded areas represent the range between the minimum and maximum values in SSP2-4.5 and ARISE-SAI-1.5.

With the steps outlined above, ensembles can be generated by sampling different annual temperature maps from existing data as the initial temperature input. These steps are used to generate the autoregressive data in Figure 4.4 and Figure 4.5 and for the data in the following Figures 4.6 and 4.7.

Take for example Figure 4.6 which shows both real (shading) and generated (solid lines) data for the SSP2-4.5 simulations and the ARISE-SAI-1.5 simulations. The machine learning generated data uses CO₂ values from SSP2-4.5, the SAI scenario from ARISE-SAI-1.5, and internal variability sampled from the piControl-WACCM simulation. The AOD input is generated using the recursive steps outlined above starting from the year 2034. The AOD machine learning em-

ulator begins drifting after 15 years and so only years 2035 to 2050 are generated here. The full ARISE-SAI-1.5 timeseries with the drift is shown in Figure B.6. By sampling internal variability from the piControl-WACCM simulation and using the CO₂ and SAI scenarios from already existing simulations, the machine learning emulator can generate data similar to that of the data simulated with CESM2 but quicker (Figure 4.6). The ability of this machine learning emulator to consider internal variability and responses to changes in CO₂ makes it a more useful tool for looking at future climate scenarios that involve SAI. Users that use this emulator for looking at future scenarios can also consider changes to CO₂ emissions as well as deployment scenarios.

Consider the work by Keys et al., (2022) which used the ARISE-SAI-1.5 simulations to calculate what regions are more likely to experience warming after deployment. This warming, when cooling is expected, can be perceived as a failure of the deployment. This experienced warming is referred to as *perceived failure*. This work finds that some of the wealthiest nations have the highest chances of experiencing perceived failure (Figure 4.7a adapted from Keys et al., (2022)), which could influence potential retaliation or motivate additional actors to deploy SAI or other climate intervention strategies. By using the ARISE-SAI-1.5 data as the initial temperature input (T₀), the machine learning emulator can sample the same internal variability and then calculate the probability of perceived failure following the same steps in Keys et al., (2022). Results from using the machine learning generated data is shown in Figure 4.7b. When sampling the same internal variability, the emulated data highlights many similar regions as shown in Keys et al., (2022) such as North America, Europe, and parts of Asia.

The authors of Keys et al., (2022) had only 10 ensemble members and concluded that the regions identified as most likely to experience perceived failure may be a result of the sampled internal variability in the ARISE-SAI-1.5 simulation.

Given the ability of the SAI emulator to replicate the results in Keys et al., (2022), the SAI emulator is now used to calculate the probability of perceived failure from a 10 member ensemble with different sampled internal variability. This indicates whether these regions most likely to experience perceived failure are a function of the sampled internal variability, or a possible response to

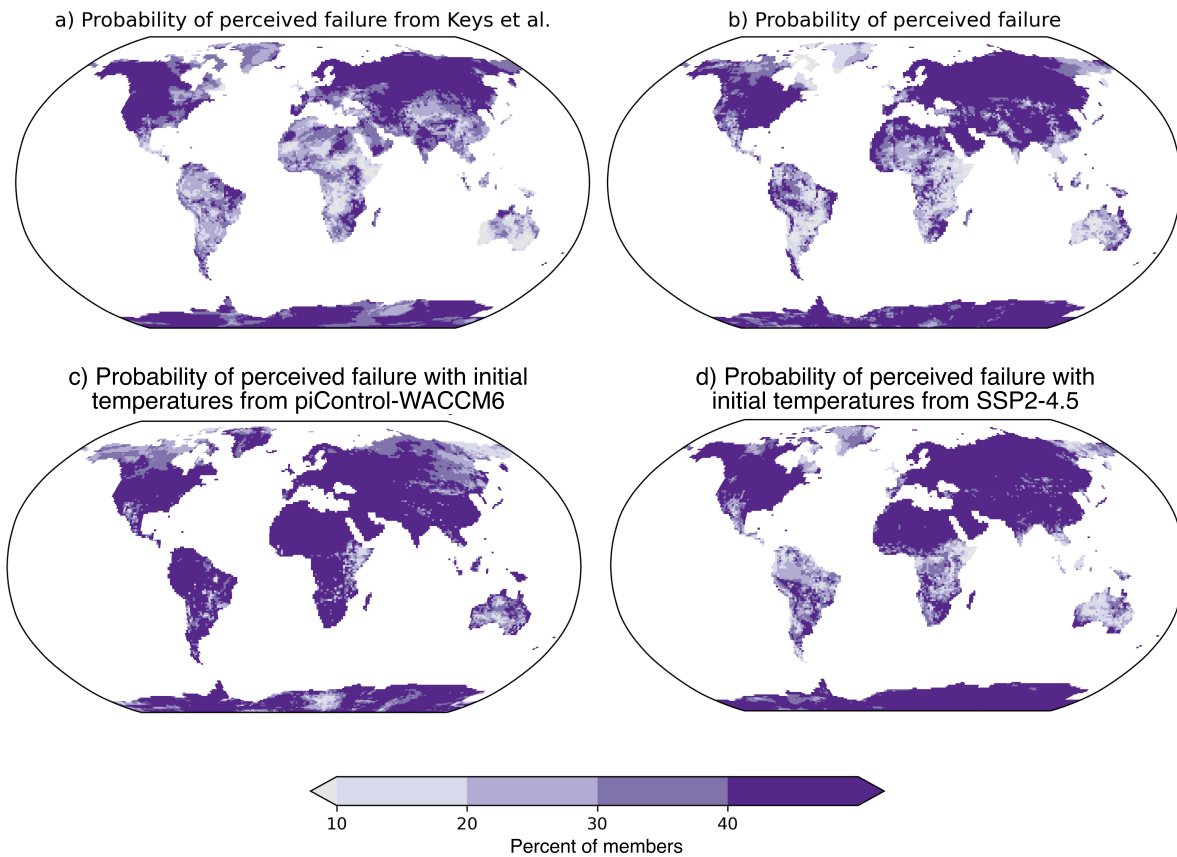


Figure 4.7: Probability ensemble members experience perceived failure in the 10 year period following an SAI deployment. Perceived failure is categorized as a positive trend in near surface temperature across the 10 years post deployment. (a) Probability of perceived failure of the 10 member ARISE-SAI-1.5 simulation adapted from Keys et al., (2022). (b) Probability of perceived failure from the 10 member AIRSE-SAI-1.5 as predicted by the SAI emulator (c) Probability of perceived failure sampling internal variability from the piControl-WACCM simulation. (c) Probability of perceived failure sampling internal variability from the SSP2-4.5 simulation.

SAI as learned by the SAI emulator. These ensembles will use the same CO₂ values from SSP2-4.5 and SO₂ values from one of the ten ARISE-SAI-1.5 ensemble members. The initial temperature is sampled from the piControl-WACCM simulation for Figure 4.7c and from the SS2-4.5 simulation for 4.7d.

When sampling initial temperature from the piControl-WACCM simulation, almost all regions experience a greater than 50% chance of experiencing perceived failure. When sampling initial temperature from the SSP2-4.5 simulation, similar areas appear as most likely to experience perceived failure. This difference in what regions are most likely to experience perceived failure indicate that the measure of perceived failure is dominated by the internal variability.

Calculating the probability of perceived failure from data generated with the machine learning emulator is a good way to represent how this SAI scenario emulator may be used to explore possible SAI impacts alongside internal variability. While the SAI emulator is not perfect in its recreation of the perceived failure work in Keys et al., (2022), it does motivate similar questions such as, how might these communities respond if an SAI deployment is perceived as a failure? The benefit of the machine learning emulator is that it is cheaper and faster than Earth system models.

4.5 Conclusion

This chapter introduces an initial version of a novel machine learning emulator. This machine learning emulator is designed for a specific task rather than emulating the entire atmosphere and it contains multiple smaller machine learning models. Importantly, one of the machine learning models is autoregressive which allows it to be used for scenario design. This multi-model approach provides an interpretable machine learning workflow to quickly explore future climate scenarios with different decisions surrounding SAI. The internal variability machine learning model allows scenarios to consider the impact of internal variability on possible human responses. The CO₂ machine learning model allows a user to create scenarios with changes in CO₂ emissions. Finally, the AOD machine learning model and SAI machine learning model allow a user to explore scenarios that contain SAI.

While this specialized emulator generates near surface temperature responses to SAI quickly and cheaply compared to Earth system models, it is not designed to replace Earth system models for SAI impact research or be used to identify physical responses to deployment strategies. Machine learning emulators rely on statistical relationships between inputs and outputs learned during the training process. The goal of this emulator is to aid in scenario design and therefore, training and evaluation focused on determining its skill and usability for these specific tasks. Additional emulators for different variables such as precipitation and atmospheric circulation could be trained with the same setup outlined here to explore atmospheric responses to different variables.

Chapter 5

Closing Remarks

It is important to prioritize lowering greenhouse gas emissions, however, it is also crucial to be prepared for the range of possible futures we may find ourselves in. As the remainder of the century is expected to bring ongoing and potentially intensifying climate change impacts, there is increasing motivation and incentive to research and understand methods of climate intervention [3–5].

The ways in which climate intervention may happen in the future are vast, with unique and impactful consequences [20, 26, 28]. These three projects are among a select few leading efforts to redefine the way in which climate intervention research is approaching scenarios of climate intervention. This dissertation works to better understand possible future scenarios that involve a deployment of a climate intervention known as SAI by accomplishing two goals. These goals, along with how the author sees this work continuing to develop, are described in the next two sections.

5.1 Future of SAI Scenarios

The first goal of this work is to investigate how climate intervention scenarios are currently designed and simulated. When researching future climates that involve climate intervention, much work goes into considering what scenarios are plausible and therefore worth researching in detail. What scenarios are plausible changes based on the person designing the scenario and evolves with changes in technology, politics, and economies. However, it is easy for a field to converge on a set of scenarios that then dominate the field of research [104]. For instance, the Shared Socioeconomic Pathways and the Representative Concentration Pathways are designed to simulate a variety of plausible climate futures with different predetermined future changes (i.e. changes in emissions, land use, and population to name a few). Despite the limited scope of these scenarios and the fact that some rely on carbon capture technology on a scale that does not exist yet [99], these

scenarios have become foundational for climate change research. While these scenarios provide benefits such as standardizing output from Earth system models, they may also restrict the plausible climate change futures studied since they are often misinterpreted to represent the complete set of possible climate change futures [104, 105].

Since atmospheric responses to SAI are more dependent on where and when SO_2 is injected than where and when CO_2 is emitted, and since climate intervention has many possible political and social impacts, it is important that scenarios involving SAI remain diverse and constantly challenged. The work discussed in this dissertation does this by quantifying the impact of internal variability on a control algorithm and accessing impacts shortly after a single actor deployment.

Chapter 2 quantifies the sensitivity of a commonly used control algorithm to internal variability. These control algorithms are often used to determine where and how much SO_2 to inject into the stratosphere in response to experienced warming. While this algorithm is used to show an example of a deployment strategy that minimizes negative consequences associated with SAI, it is important to question whether these algorithms are appropriate for a realistic deployment scenario. These algorithms have been documented to amplify and attenuate the frequency of internal variability within climate models and to be sensitive to the current state of the atmosphere-ocean system [29]. As mentioned previously, is it acceptable that more SO_2 is injected only because the atmospheric-ocean system is in an El Niño phase rather than a La Niña phase? This demonstrates that even commonly used SAI injection scenarios require additional consideration about their implication for real world validity.

Single actor scenarios with single point injections typically have unfavorable atmospheric responses, and therefore, are often categorized as an undesirable climate scenario. However, undesirable does not mean unlikely. With the increasing climate change impacts and technology developments, its worth understanding possible atmospheric responses to single actor deployments. Since single actor scenarios have little to no cooperation, this deployment strategy may have political and social impacts along with atmospheric impacts. The work shown in Chapter 3 indicates measurable atmospheric responses at one to two years after deployment, albeit with large injec-

tions. More importantly, this work highlights increased variability in global precipitation and near surface temperature in the second year after deployment with injections at 1Tg/year.

Analyzing atmospheric responses this quickly after deployment is novel because the signal-to-noise ratio of the response to internal variability is small. However, the size and composition of CRASSULA allow these possible impacts to be quantified in a new way. Future work in the single actor space should look at smaller injections and simulate data that does not have a continuous injection such as CRASSULA has. These possible atmospheric responses may play a crucial role in communication and mitigation strategies in the case of a single actor deployment.

The decision of whether to deploy SAI or not is one of complex social, environmental, and political consideration [26, 28], but also, the decision about the research surrounding climate intervention, including the scenario design is also a decision of complex social, environmental, and political consideration. Chapter 2 and Chapter 3 emphasize the need to question how SAI scenarios are simulated, what assumptions are made, and what possible scenarios are being ignored. These three goals can be accomplished by including multiple perspectives, experiences, and expertise into the design space which needs to continue to be an emphasis in the field.

5.2 Future of SAI Machine Learning Emulators

The second goal of this work is to design and train a machine learning emulator that will make it cheaper and quicker to explore a vast range of SAI scenarios. Through statistical evaluation and comparing emulator generated results to peer-reviewed results, this work shows how the machine learning emulator introduced in this dissertation is a tool that can help explore SAI scenarios. The author hopes that the specialization, interpretability, and autoregressive capabilities of this machine learning emulator will inspire additional emulators for different specialized tasks.

A major contributor towards the success of the SAI emulator is the variety of CESM2 training data already available and the tailored training data CRASSULA. As far as the author knows, CRASSULA is the first data simulated with an Earth system model specifically to train a machine learning model. This means that during the development of CRASSULA, important questions

about the data and emulator's future use had to be considered before any data was simulated. Some of these questions included:

1. How will the SAI emulator be used and how can the data support that?
2. What will the machine learning model learn if injections occur back-to-back?
3. How long should the injections occur to allow the machine learning model to learn possible nonlinearities?
4. How much data is needed to train this model?

This author hopes that CRASSULA is not the last dataset simulated specifically to train a machine learning model. This is because designing data specifically for training a machine learning model gives more freedom to the scientist to determine which assumptions are made and what details are necessary for machine learning training. However, given how much data is required to train a machine learning model and the environmental footprint of creating this data, it is the author's opinion that the data should be designed in a way to allow it to be used beyond training a machine learning model. For example, future versions of the SAI emulator introduced here may require additional data. Additionally, Chapter 3 concluded that additional data with smaller injections and single injections would provide valuable insight about the atmospheric responses identified. Therefore, additional data could have the primary goal of training a future version of the machine learning emulator and a secondary goal of further supporting the work in Chapter 3.

Climate scenarios allow us to systematically investigate possible futures that are found to be the most probable. Understandably, in doing so, assumptions are made and scenarios are ignored. Climate intervention research needs to stay open to new scenarios whether that comes from changes in technology, politics, or society. The author hopes that work here shows why being critical about how scenarios are simulated and about what scenarios are simulated is important when considering what a future with SAI might look like. This range of possible futures and the dependency climate intervention has on the deployment scenario makes it a prime use for emulators. The author believes that many more emulators, each with different goals and biases, will enter the climate

intervention field and allow more of these climate intervention scenarios to be studied. While emulators provide an obvious benefit, with a decision as impactful as climate intervention, the author believes it is important to not solely rely on climate intervention emulators. Just as we know that Earth system models do not perfectly replicate the real Earth system, emulators do not perfectly replicate the system they are trained on.

Bibliography

- [1] VP Masson-Delmotte, Panmao Zhai, SL Pirani, C Connors, S Péan, N Berger, Y Caud, L Chen, MI Goldfarb, and Pedro M Scheel Monteiro. Ipcc, 2021: Summary for policymakers. in: Climate change 2021: The physical science basis. contribution of working group i to the sixth assessment report of the intergovernmental panel on climate change. 2021.
- [2] Kevin Anderson, John F Broderick, and Isak Stoddard. A factor of two: how the mitigation plans of ‘climate progressive’ nations fall far short of paris-compliant pathways. *Climate Policy*, 20(10):1290–1304, 2020.
- [3] National Academies of Sciences, Medicine, and Others. *Reflecting sunlight: Recommendations for solar geoengineering research and research governance*. nap.nationalacademies.org, 2021.
- [4] White House Office of Science and Technology Policy. *Congressionally Mandated Research Plan and an Initial Research Governance Framework Related to Solar Radiation Modification*. Office of Science and Technology Policy. Washington, D.C., DC, June 2023.
- [5] United Nations Environment Programme. One atmosphere: An independent expert review on solar radiation modification research and deployment. 2023.
- [6] David G Victor, M Granger Morgan, Jay Apt, John Steinbruner, and Katharine Ricke. The geoengineering option: a last resort against global warming? *Foreign Aff.*, 88:64, 2009.
- [7] Scott Barrett. Solar geoengineering’s brave new world: Thoughts on the governance of an unprecedented technology, 2014.
- [8] Michael C MacCracken. The rationale for accelerating regionally focused climate intervention research. *Earth’s Future*, 4(12):649–657, 2016.

- [9] Govindasamy Bala, PB Duffy, and KE Taylor. Impact of geoengineering schemes on the global hydrological cycle. *Proceedings of the National Academy of Sciences*, 105(22):7664–7669, 2008.
- [10] Angus J Ferraro and Hannah G Griffiths. Quantifying the temperature-independent effect of stratospheric aerosol geoengineering on global-mean precipitation in a multi-model ensemble. *Environmental Research Letters*, 11(3):034012, 2016.
- [11] Ariel L Morrison, Debanjali Pathak, Elizabeth A Barnes, and James W Hurrell. Projected changes to arctic shipping routes after stratospheric aerosol deployment in the arise-sai scenarios. *Frontiers in Climate*, 6:1426679, 2024.
- [12] JF Tjiputra, A Grini, and H Lee. Impact of idealized future stratospheric aerosol injection on the large-scale ocean and land carbon cycles. *Journal of Geophysical Research: Biogeosciences*, 121(1):2–27, 2016.
- [13] Scott C Doney, Victoria J Fabry, Richard A Feely, and Joan A Kleypas. Ocean acidification: the other co2 problem. *Annual review of marine science*, 1(1):169–192, 2009.
- [14] S Tilmes, JH Richter, B Kravitz, DG MacMartin, AS Glanville, D Visionsi, DE Kinnison, and R Müller. Sensitivity of total column ozone to stratospheric sulfur injection strategies. *Geophysical Research Letters*, 48(19):e2021GL094058, 2021.
- [15] Daniele Visionsi, Andy Jones, James Haywood, Roland Séférian, Pierre Nabat, Olivier Boucher, Ewa Monica Bednarz, Ulrike Niemeier, et al. Stratospheric ozone response to sulfate aerosol and solar dimming climate interventions based on the g6 geoengineering model intercomparison project (geomip) simulations. *Atmospheric Chemistry and Physics*, 22(7):4557–4579, 2022.
- [16] Ben Kravitz, Douglas G MacMartin, Simone Tilmes, Jadwiga H Richter, Michael J Mills, Wei Cheng, Katherine Dagon, Anne S Glanville, Jean-Francois Lamarque, Isla R Simpson, et al. Comparing surface and stratospheric impacts of geoengineering with different so2

- injection strategies. *Journal of Geophysical Research: Atmospheres*, 124(14):7900–7918, 2019.
- [17] Simone Tilmes, Jadwiga H Richter, Michael J Mills, Ben Kravitz, Douglas G MacMartin, Francis Vitt, Joseph J Tribbia, and Jean-Francois Lamarque. Sensitivity of aerosol distribution and climate response to stratospheric SO₂ injection locations. *J. Geophys. Res.*, 122(23):12,591–12,615, December 2017.
- [18] Douglas G MacMartin, Ben Kravitz, Simone Tilmes, Jadwiga H Richter, Michael J Mills, Jean-Francois Lamarque, Joseph J Tribbia, and Francis Vitt. The climate response to stratospheric aerosol geoengineering can be tailored using multiple injection locations. *Journal of Geophysical Research: Atmospheres*, 122(23):12–574, 2017.
- [19] D G MacMartin, D Visoni, B Kravitz, J H Richter, T Felgenhauer, W R Lee, D R Morrow, E A Parson, and M Sugiyama. Scenarios for modeling solar radiation modification. *Proc. Natl. Acad. Sci. U. S. A.*, 119(33):e2202230119, August 2022.
- [20] Patrick W Keys and Curtis M Bell. Designing a scenario of unilateral climate intervention. 2024.
- [21] Wilfried Rickels, Martin F Quaas, Katharine Ricke, Johannes Quaas, Juan Moreno-Cruz, and Sjak Smulders. Who turns the global thermostat and by how much? *Energy economics*, 91:104852, 2020.
- [22] Daniel Martin Hueholt, Elizabeth A Barnes, and James W Hurrell. Exploring the influence of internal climate variability and forced change on arctic greening. In *AGU Fall Meeting Abstracts*, volume 2024, pages GC11A–04, 2024.
- [23] Holly Jean Buck. Geoengineering: Re-making climate for profit or humanitarian intervention? *Development and Change*, 43(1):253–270, 2012.
- [24] Curtis M Bell and Patrick W Keys. Strategic logic of unilateral climate intervention. *Environ. Res. Lett.*, 18(10):104045, October 2023.

- [25] Douglas G MacMartin, Peter J Irvine, Ben Kravitz, and Joshua B Horton. Technical characteristics of a solar geoengineering deployment and implications for governance. *Clim. Policy*, 19(10):1325–1339, November 2019.
- [26] Edward A Parson and Jesse L Reynolds. Solar geoengineering: Scenarios of future governance challenges. *Futures*, 133(102806):102806, October 2021.
- [27] Patrick W Keys, Elizabeth A Barnes, Noah S Diffenbaugh, James W Hurrell, and Curtis M Bell. Potential for perceived failure of stratospheric aerosol injection deployment. *Proceedings of the National Academy of Sciences*, 119(40):e2210036119, 2022.
- [28] Holly Jean Buck. Geoengineering: re-making climate for profit or humanitarian intervention?: Geoengineering: For profit or humanitarian intervention? *Dev. Change*, 43(1):253–270, 2012.
- [29] Douglas G MacMartin, Ben Kravitz, David W Keith, and Andrew Jarvis. Dynamics of the coupled human–climate system resulting from closed-loop control of solar geoengineering. *Clim. Dyn.*, 43(1):243–258, July 2014.
- [30] Alice F Wells, Matthew Henry, Ewa M Bednarz, Douglas G MacMartin, Andy Jones, Mohit Dalvi, and James M Haywood. Identifying climate impacts from different stratospheric aerosol injection strategies in ukesm1. *Earth’s Future*, 12(3):e2023EF004358, 2024.
- [31] Yan Zhang, Douglas G MacMartin, Daniele Visoni, Ewa M Bednarz, and Ben Kravitz. Hemispherically symmetric strategies for stratospheric aerosol injection. *Earth System Dynamics*, 15(2):191–213, 2024.
- [32] Charlotte Connolly, Emily Prewett, Elizabeth A Barnes, and James W Hurrell. Quantifying the impact of internal variability on the cesm2 control algorithm for stratospheric aerosol injection. *Earth’s Future*, 12(6):e2023EF004300, 2024.

- [33] Philip J Rasch, Paul J Crutzen, and Danielle B Coleman. Exploring the geoengineering of climate using stratospheric sulfate aerosols: The role of particle size. *Geophysical Research Letters*, 35(2), 2008.
- [34] Ben Kravitz, Ken Caldeira, Olivier Boucher, Alan Robock, Philip J Rasch, Kari Alterskjaer, Diana Bou Karam, Jason N S Cole, Charles L Curry, James M Haywood, Peter J Irvine, Duoying Ji, Andy Jones, Jón Egill Kristjánsson, Daniel J Lunt, John C Moore, Ulrike Niemeier, Hauke Schmidt, Michael Schulz, Balwinder Singh, Simone Tilmes, Shingo Watanabe, Shuting Yang, and Jin-Ho Yoon. Climate model response from the geoengineering model intercomparison project (GeoMIP). 118(15):8320–8332, August 2013.
- [35] B Kravitz, A Robock, S Tilmes, O Boucher, J M English, P J Irvine, A Jones, M G Lawrence, M MacCracken, H Muri, J C Moore, U Niemeier, S J Phipps, J Sillmann, T Storelvmo, H Wang, and S Watanabe. The geoengineering model intercomparison project phase 6 (GeoMIP6): simulation design and preliminary results. *Geosci. Model Dev.*, 8(10):3379–3392, October 2015.
- [36] Simone Tilmes, Jadwiga H Richter, Ben Kravitz, Douglas G MacMartin, Michael J Mills, Isla R Simpson, Anne S Glanville, John T Fasullo, Adam S Phillips, Jean-Francois Lamarque, Joseph Tribbia, Jim Edwards, Sheri Mickelson, and Siddhartha Ghosh. CESM1(WACCM) stratospheric aerosol geoengineering large ensemble project. *Bull. Am. Meteorol. Soc.*, 99(11):2361–2371, November 2018.
- [37] Jadwiga H Richter, Daniele Visioni, Douglas G MacMartin, David A Bailey, Nan Rosenbloom, Brian Dobbins, Walker R Lee, Mari Tye, and Jean-Francois Lamarque. Assessing responses and impacts of solar climate intervention on the earth system with stratospheric aerosol injection (ARISE-SAI): protocol and initial results from the first simulations. *Geosci. Model Dev.*, 15(22):8221–8243, November 2022.
- [38] G Danabasoglu, J-F Lamarque, J Bacmeister, D A Bailey, A K DuVivier, J Edwards, L K Emmons, J Fasullo, R Garcia, A Gettelman, C Hannay, M M Holland, W G Large, P H

- Lauritzen, D M Lawrence, J T M Lenaerts, K Lindsay, W H Lipscomb, M J Mills, R Neale, K W Oleson, B Otto-Bliesner, A S Phillips, W Sacks, S Tilmes, L Kampenhout, M Vertenstein, A Bertini, J Dennis, C Deser, C Fischer, B Fox-Kemper, J E Kay, D Kinnison, P J Kushner, V E Larson, M C Long, S Mickelson, J K Moore, E Nienhouse, L Polvani, P J Rasch, and W G Strand. The community earth system model version 2 (CESM2). *J. Adv. Model. Earth Syst.*, 12(2), February 2020.
- [39] Ben Kravitz, Douglas G MacMartin, Michael J Mills, Jadwiga H Richter, Simone Tilmes, Jean-Francois Lamarque, Joseph J Tribbia, and Francis Vitt. First simulations of designing stratospheric sulfate aerosol geoengineering to meet multiple simultaneous climate objectives. *Journal of Geophysical Research: Atmospheres*, 122(23):12–616, 2017.
- [40] Clara Deser, Adam Phillips, Vincent Bourdette, and Haiyan Teng. Uncertainty in climate change projections: the role of internal variability. *Clim. Dyn.*, 38(3):527–546, February 2012.
- [41] Kevin E Trenberth. The definition of el niño. *Bull. Am. Meteorol. Soc.*, 78(12):2771–2778, December 1997.
- [42] M Ho, AS Kiem, and DC Verdon-Kidd. The southern annular mode: a comparison of indices. *Hydrol. Earth Syst. Sci.*, 16(3):967–982, March 2012.
- [43] James W Hurrell and Clara Deser. North atlantic climate variability: The role of the north atlantic oscillation. *J. Mar. Syst.*, 79(3):231–244, February 2010.
- [44] RE Holasek, S Self, and AW Woods. Satellite observations and interpretation of the 1991 mount pinatubo eruption plumes. *Journal of Geophysical Research: Solid Earth*, 101(B12):27635–27655, 1996.
- [45] Edward Hanna and Thomas E Cropper. North atlantic oscillation. In *Oxford Research Encyclopedia of Climate Science*. 2017.

- [46] Gwendal Rivièrè and Marie Drouard. Understanding the contrasting north atlantic oscillation anomalies of the winters of 2010 and 2014. *Geophysical Research Letters*, 42(16):6868–6875, 2015.
- [47] Ramón Fuentes-Franco, David Docquier, Torben Koenigk, Klaus Zimmermann, and Filippo Giorgi. Winter heavy precipitation events over northern europe modulated by a weaker nao variability by the end of the 21st century. *npj Climate and Atmospheric Science*, 6(1):72, 2023.
- [48] Ryan L Fogt and Gareth J Marshall. The southern annular mode: variability, trends, and climate impacts across the southern hemisphere. *Wiley Interdisciplinary Reviews: Climate Change*, 11(4):e652, 2020.
- [49] N P Gillett, T D Kell, and PD Jones. Regional climate impacts of the southern annular mode. *Geophysical Research Letters*, 33(23), 2006.
- [50] Karl Johan Åström and Richard M Murray. *Feedback systems: an introduction for scientists and engineers*. Princeton university press, 2021.
- [51] Andrew Jarvis and David Leedal. The geoengineering model intercomparison project (GeoMIP): a control perspective. *Atmos. Sci. Lett.*, 13(3):157–163, July 2012.
- [52] J K Angell. Variation in global tropospheric temperature after adjustment for the el nino influence, 1958-89. *Geophys. Res. Lett.*, 17(8):1093–1096, July 1990.
- [53] JC Wilson, HH Jonsson, CA Brock, DW Toohey, LM Avallone, D Baumgardner, JE Dye, LR Poole, DC Woods, RJ DeCoursey, et al. In situ observations of aerosol and chlorine monoxide after the 1991 eruption of mount pinatubo: effect of reactions on sulfate aerosol. *Science*, 261(5125):1140–1143, 1993.
- [54] Gregg JS Bluth, Scott D Doiron, Charles C Schnetzler, Arlin J Krueger, and Louis S Walter. Global tracking of the so2 clouds from the june, 1991 mount pinatubo eruptions. *Geophysical Research Letters*, 19(2):151–154, 1992.

- [55] D E Parker, H Wilson, and P D Jones. The impact of mount pinatubo on world-wide temperatures. *J. Appl. Meteorol. Climatol.*, 1996.
- [56] Andrew Jarvis and David Leedal. The geoengineering model intercomparison project (gemip): A control perspective. *Atmospheric Science Letters*, 13(3):157–163, 2012.
- [57] Aiguo Dai, John C Fyfe, Shang-Ping Xie, and Xingang Dai. Decadal modulation of global surface temperature by internal climate variability. *Nat. Clim. Chang.*, 5(6):555–559, April 2015.
- [58] Annarita Mariotti, Paolo M Ruti, and Michel Rixen. Progress in subseasonal to seasonal prediction through a joint weather and climate community effort. *npj Climate and Atmospheric Science*, 1(1):1–4, March 2018.
- [59] Kailai Xu and Eric Darve. Physics constrained learning for data-driven inverse modeling from sparse observations. *J. Comput. Phys.*, 453:110938, March 2022.
- [60] Wake Smith. The cost of stratospheric aerosol injection through 2100. *Environmental Research Letters*, 15(11):114004, 2020.
- [61] Jim M Haywood, Andy Jones, Nicolas Bellouin, and David Stephenson. Asymmetric forcing from stratospheric aerosols impacts sahelian rainfall. *Nature Climate Change*, 3(7):660–665, 2013.
- [62] Simone Tilmes, Douglas G MacMartin, Jan Lenaerts, Leo Van Kampenhout, Laura Muntjewerf, Lili Xia, Cheryl S Harrison, Kristen M Krumhardt, Michael J Mills, Ben Kravitz, et al. Reaching 1.5 and 2.0° c global surface temperature targets using stratospheric aerosol geoengineering. *Earth System Dynamics*, 11(3):579–601, 2020.
- [63] Yan Zhang, Douglas G MacMartin, Daniele Visioni, Ewa Bednarz, and Ben Kravitz. Introducing a comprehensive set of stratospheric aerosol injection strategies. *EGUsphere*, 2023:1–32, 2023.

- [64] Ewa M Bednarz, Daniele Visioni, Jadwiga H Richter, Amy H Butler, and Douglas G MacMartin. Impact of the latitude of stratospheric aerosol injection on the southern annular mode. *Geophysical research letters*, 49(19):e2022GL100353, 2022.
- [65] Ewa M Bednarz, Daniele Visioni, Ben Kravitz, Andy Jones, James M Haywood, Jadwiga Richter, Douglas G MacMartin, and Peter Braesicke. Climate response to off-equatorial stratospheric sulfur injections in three earth system models—part 2: Stratospheric and free-tropospheric response. *Atmospheric Chemistry and Physics*, 23(1):687–709, 2023.
- [66] A Gettelman, M J Mills, D E Kinnison, R R Garcia, A K Smith, D R Marsh, S Tilmes, F Vitt, C G Bardeen, J McInerney, H-L Liu, S C Solomon, L M Polvani, L K Emmons, J-F Lamarque, J H Richter, A S Glanville, J T Bacmeister, A S Phillips, R B Neale, I R Simpson, A K DuVivier, A Hodzic, and W J Randel. The whole atmosphere community climate model version 6 (WACCM6). *J. Geophys. Res.*, 124(23):12380–12403, December 2019.
- [67] Michael Allen. The huge carbon footprint of large-scale computing. *Physics world*, 35(3):46, 2022.
- [68] The carbon footprint of computational research. *Nature Computational Science*, 3(8):659–659, 2023.
- [69] Alan Robock. Volcanic eruptions and climate. *Reviews of geophysics*, 38(2):191–219, 2000.
- [70] Nicholas A Davis, Daniele Visioni, Rolando R Garcia, Douglas E Kinnison, Daniel R Marsh, Mike Mills, Jadwiga H Richter, Simone Tilmes, Chuck G Bardeen, Andrew Gettelman, et al. Climate, variability, and climate sensitivity of “middle atmosphere” chemistry configurations of the community earth system model version 2, whole atmosphere community climate model version 6 (cesm2 (waccm6)). *Journal of Advances in Modeling Earth Systems*, 15(9):e2022MS003579, 2023.

- [71] Georgiy Stenchikov, Thomas L Delworth, V Ramaswamy, Ronald J Stouffer, Andrew Wittenberg, and Fanrong Zeng. Volcanic signals in oceans. *Journal of Geophysical Research: Atmospheres*, 114(D16), 2009.
- [72] Sebastian Illing, Christopher Kadow, Holger Pohlmann, and Claudia Timmreck. Assessing the impact of a future volcanic eruption on decadal predictions. *Earth System Dynamics*, 9(2):701–715, 2018.
- [73] Hongwei Sun, Stephen Bourguet, Sebastian Eastham, and David Keith. Optimizing injection locations relaxes altitude-lifetime trade-off for stratospheric aerosol injection. *Geophysical Research Letters*, 50(16):e2023GL105371, 2023.
- [74] T Sakai, O Uchino, T Nagai, B Liley, I Morino, and T Fujimoto. Long-term variation of stratospheric aerosols observed with lidars over tsukuba, japan, from 1982 and lauder, new zealand, from 1992 to 2015. *Journal of Geophysical Research: Atmospheres*, 121(17):10–283, 2016.
- [75] Daniele Visioni, Ewa M Bednarz, Walker R Lee, Ben Kravitz, Andy Jones, Jim M Haywood, and Douglas G MacMartin. Climate response to off-equatorial stratospheric sulfur injections in three earth system models—part 1: Experimental protocols and surface changes. *Atmospheric Chemistry and Physics*, 23(1):663–685, 2023.
- [76] Syukuro Manabe, Kirk Bryan, and Michael J Spelman. Transient response of a global ocean-atmosphere model to a doubling of atmospheric carbon dioxide. *Journal of Physical Oceanography*, 20(5):722–749, 1990.
- [77] A Jarvis. The magnitudes and timescales of global mean surface temperature feedbacks in climate models. *Earth System Dynamics*, 2(2):213–221, 2011.
- [78] Kevin E Trenberth, Aiguo Dai, Gerard Van Der Schrier, Philip D Jones, Jonathan Barichivich, Keith R Briffa, and Justin Sheffield. Global warming and changes in drought. *Nature Climate Change*, 4(1):17–22, 2014.

- [79] Angus J Ferraro, Eleanor J Highwood, and Andrew J Charlton-Perez. Weakened tropical circulation and reduced precipitation in response to geoengineering. *Environmental Research Letters*, 9(1):014001, 2014.
- [80] Douglas G MacMartin, Katharine L Ricke, and David W Keith. Solar geoengineering as part of an overall strategy for meeting the 1.5 c paris target. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 376(2119):20160454, 2018.
- [81] Clara Deser, Adam Phillips, Vincent Bourdette, and Haiyan Teng. Uncertainty in climate change projections: the role of internal variability. *Climate dynamics*, 38:527–546, 2012.
- [82] V Aquila, CI Garfinkel, PA Newman, LD Oman, and DW Waugh. Modifications of the quasi-biennial oscillation by a geoengineering perturbation of the stratospheric aerosol layer. *Geophysical Research Letters*, 41(5):1738–1744, 2014.
- [83] Andy Jones, Jim M Haywood, Adam A Scaife, Olivier Boucher, Matthew Henry, Ben Kravitz, Thibaut Lurton, Pierre Nabat, Ulrike Niemeier, Roland Séférian, et al. The impact of stratospheric aerosol intervention on the north atlantic and quasi-biennial oscillations in the geoengineering model intercomparison project (geomip) g6sulfur experiment. *Atmospheric Chemistry and Physics*, 22(5):2999–3016, 2022.
- [84] Mark P Baldwin, LJ Gray, TJ Dunkerton, K Hamilton, PH Haynes, WJ Randel, JR Holton, MJ Alexander, I Hirota, T Horinouchi, et al. The quasi-biennial oscillation. *Reviews of Geophysics*, 39(2):179–229, 2001.
- [85] Ewa M Bednarz, Amy H Butler, Daniele Visioni, Yan Zhang, Ben Kravitz, and Douglas G MacMartin. Injection strategy—a driver of atmospheric circulation and ozone response to stratospheric aerosol geoengineering. *Atmospheric Chemistry and Physics*, 23(21):13665–13684, 2023.

- [86] S Tilmes, Daniele Richter, Visionsi, Andy Jones, James Haywood, Roland Séférian, Pierre Nabat, Olivier Boucher, Ewa Monica Bednarz, Ulrike Niemeier, et al. Stratospheric ozone response to sulfate aerosol and solar dimming climate interventions based on the G6 geo-engineering model intercomparison project (GeoMIP) simulations. *Atmos. Chem. Phys.*, 22(7):4557–4579, April 2022.
- [87] Martyn P Chipperfield and Slimane Bekki. Opinion: Stratospheric ozone–depletion, recovery and new challenges. *Atmospheric Chemistry and Physics*, 24(4):2783–2802, 2024.
- [88] Xuekun Fang, John A Pyle, Martyn P Chipperfield, John S Daniel, Sunyoung Park, and Ronald G Prinn. Challenges for the recovery of the ozone layer. *Nature Geoscience*, 12(8):592–596, 2019.
- [89] Edward A Parson, Virginia Burkett, Karen Fisher-Vanden, David Keith, Linda Mearns, Hugh Pitcher, Cynthia Rosenzweig, and Mort Webster. Global-change scenarios: their development and use. 2007.
- [90] Edward A Parson. Useful global-change scenarios: current issues and challenges. *Environmental Research Letters*, 3(4):045016, 2008.
- [91] Chenrui Diao, Patrick Keys, Curtis M Bell, Elizabeth A Barnes, and James Wilson Hurrell. A model study exploring the decision loop between unilateral stratospheric aerosol injection scenario design and earth system simulations. *Authorea Preprints*, 2024.
- [92] Jared Farley, Douglas G MacMartin, Daniele Visionsi, and Ben Kravitz. Emulating inconsistencies in stratospheric aerosol injection. *Environmental Research: Climate*, 3(3):035012, 2024.
- [93] Douglas G MacMartin and Ben Kravitz. Dynamic climate emulators for solar geoengineering. *Atmos. Chem. Phys.*, 16(24):15789–15799, December 2016.
- [94] Gokhan Danabasoglu. Near cesm2-waccm model output prepared for cmip6 cmip picontrol, 2019.

- [95] CESM. C mip6 cesm2 picontrol experiment with cam6, interactive land (clm5), coupled ocean (pop2) with biogeochemistry (marbl), interactive sea ice (cice5.1), and non-evolving land ice (cism2.1). 2019.
- [96] Veronika Eyring, Sandrine Bony, Gerald A Meehl, Catherine A Senior, Bjorn Stevens, Ronald J Stouffer, and Karl E Taylor. Overview of the coupled model intercomparison project phase 6 (cmip6) experimental design and organization. *Geoscientific Model Development*, 9(5):1937–1958, 2016.
- [97] Keith B Rodgers, Sun-Seon Lee, Nan Rosenbloom, Axel Timmermann, Gokhan Danabasoglu, Clara Deser, Jim Edwards, Ji-Eun Kim, Isla R Simpson, Karl Stein, Malte F Stuecker, Ryohei Yamaguchi, Tamás Bódai, Eui-Seok Chung, Lei Huang, Who M Kim, Jean-François Lamarque, Danica L Lombardozzi, William R Wieder, and Stephen G Yeager. Ubiquity of human-induced changes in climate variability. *Earth Syst. Dyn.*, 12(4):1393–1411, December 2021.
- [98] Peter A Bogenschutz, Andrew Gettelman, Cecile Hannay, Vincent E Larson, Richard B Neale, Cheryl Craig, and Chih-Chieh Chen. The path to cam6: Coupled simulations with cam5. 4 and cam5. 5. *Geoscientific Model Development*, 11(1):235–255, 2018.
- [99] Keywan Riahi, Detlef P van Vuuren, Elmar Kriegler, Jae Edmonds, Brian C O’Neill, Shinichiro Fujimori, Nico Bauer, Katherine Calvin, Rob Dellink, Oliver Fricko, Wolfgang Lutz, Alexander Popp, Jesus Crespo Cuaresma, Samir Kc, Marian Leimbach, Leiwen Jiang, Tom Kram, Shilpa Rao, Johannes Emmerling, Kristie Ebi, Tomoko Hasegawa, Petr Havlik, Florian Humpenöder, Lara Aleluia Da Silva, Steve Smith, Elke Stehfest, Valentina Bosetti, Jiyong Eom, David Gernaat, Toshihiko Masui, Joeri Rogelj, Jessica Strefler, Laurent Drouet, Volker Krey, Gunnar Luderer, Mathijs Harmsen, Kiyoshi Takahashi, Lavinia Baumstark, Jonathan C Doelman, Mikiko Kainuma, Zbigniew Klimont, Giacomo Marangoni, Hermann Lotze-Campen, Michael Obersteiner, Andrzej Tabeau, and Massimo Tavoni. The shared

- socioeconomic pathways and their energy, land use, and greenhouse gas emissions implications: An overview. *Glob. Environ. Change*, 42:153–168, January 2017.
- [100] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5–9, 2015, proceedings, part III 18*, pages 234–241. Springer, 2015.
- [101] Vagelis Plevris, German Solorzano, Nikolaos P Bakas, and Mohamed El Amine Ben Seghier. Investigation of performance metrics in regression analysis and machine learning-based prediction models. 2022.
- [102] Flavio Lehner, Andrew P Schurer, Gabriele C Hegerl, Clara Deser, and Thomas L Frölicher. The importance of enso phase during volcanic eruptions for detection and attribution. *Geophysical Research Letters*, 43(6):2851–2858, 2016.
- [103] Flavio Lehner and Clara Deser. Origin, importance, and predictive limits of internal climate variability. *Environmental Research: Climate*, 2(2):023001, 2023.
- [104] Patrick W Keys. The plot must thicken: a call for increased attention to social surprises in scenarios of climate futures. *Environmental Research Letters*, 18(8):081003, 2023.
- [105] Roger Pielke Jr and Justin Ritchie. Distorting the view of our climate future: The misuse and abuse of climate pathways and scenarios. *Energy Research & Social Science*, 72:101890, 2021.
- [106] Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*, pages 807–814, 2010.
- [107] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Appendix A

Supporting Information for Chapter 2

Figure A.1 demonstrates the process to calculate temperature anomalies associated with the Mt. Pinatubo eruption for global average temperature. To calculate the temperature patterns associated with the Mt. Pinatubo eruption used in Chapter 2, this process is completed at each grid point rather than for the global mean. Figure A.2 shows the temperature change from the 2035 base state to the 2045 base states. Figure A.3 shows which base state and internal variability patterns are used to make each figure in the manuscript. Figure A.4 shows the ENSO, NAO, and SAM driven portion of the injection as a function of index. Figure A.5 shows the percent change in total SO₂ injection as a function of two internal variability indices but using composites from the ARISE-SAI control simulation rather than the CESM Large Ensemble. Figure A.6 shows the deviations in T0, T1, and T2 for the ENSO and NAO internal variability patterns together. Figure A.7 shows the total SO₂ injection as a function of keeping either T0, T1, or T2 constant and varying the other two errors.


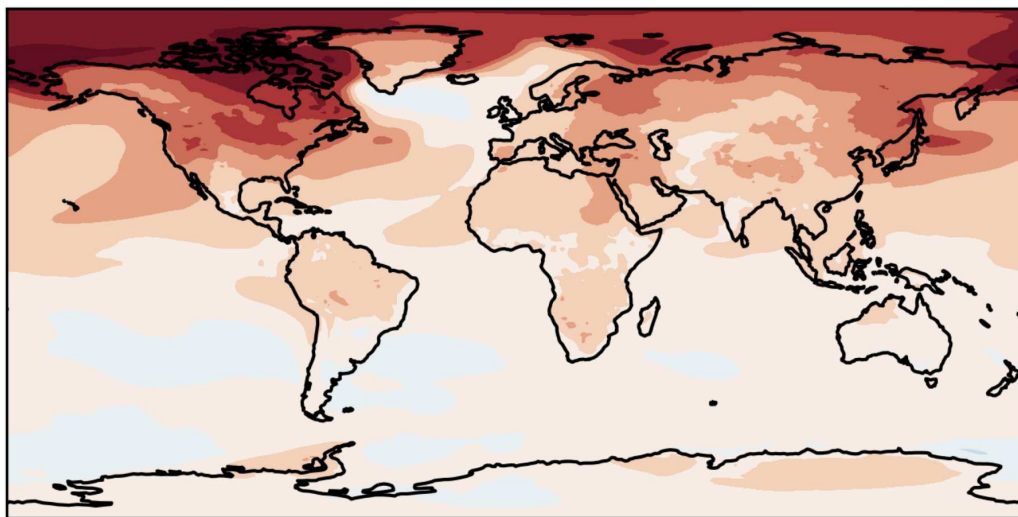


Figure A.1: Black line shows the global average temperatures. Red line is the linear trend fit to the global averaged temperatures 10 years prior to Mt. Pinatubo. The blue line is extrapolating the red line for an additional 2 years and is used to calculate the temperature anomalies associated with the Mt. Pinatubo eruption.

Base state 2045 - Base state 2035



Global mean temperature in 2035: 288.7K
 Global mean temperature in 2045: 289.1K




Figure A.2: Difference between the 2045 base state and 2035 base state.

	Base state	ENSO	NAO	SAM	Pinatubo
Figure 1a	2035	X	X		
Figure 1b	2045	X	X		
Figure 2a green line	2035	X			
Figure 2a orange line	2045	X			
Figure 2b green line	2035		X		
Figure 2b orange line	2045		X		
Figure 2c green line	2035			X	
Figure 2c orange line	2045			X	
Figure 3a	2035	X	X		
Figure 3b	2035	X		X	
Figure 3c	2035		X	X	
Figure 3d	2045	X	X		
Figure 3e	2045	X		X	
Figure 3f	2045		X	X	
Figure 4b green line	2035	X			X
Figure 4b orange line	2045	X			X
Figure 4c green line	2035		X		X
Figure 4c orange line	2045		X		X
Figure 4d green line	2035			X	X
Figure 4d orange line	2045			X	X

Figure A.3: Table showing which base state and internal variability patterns are used to make each of the figures in the manuscript.

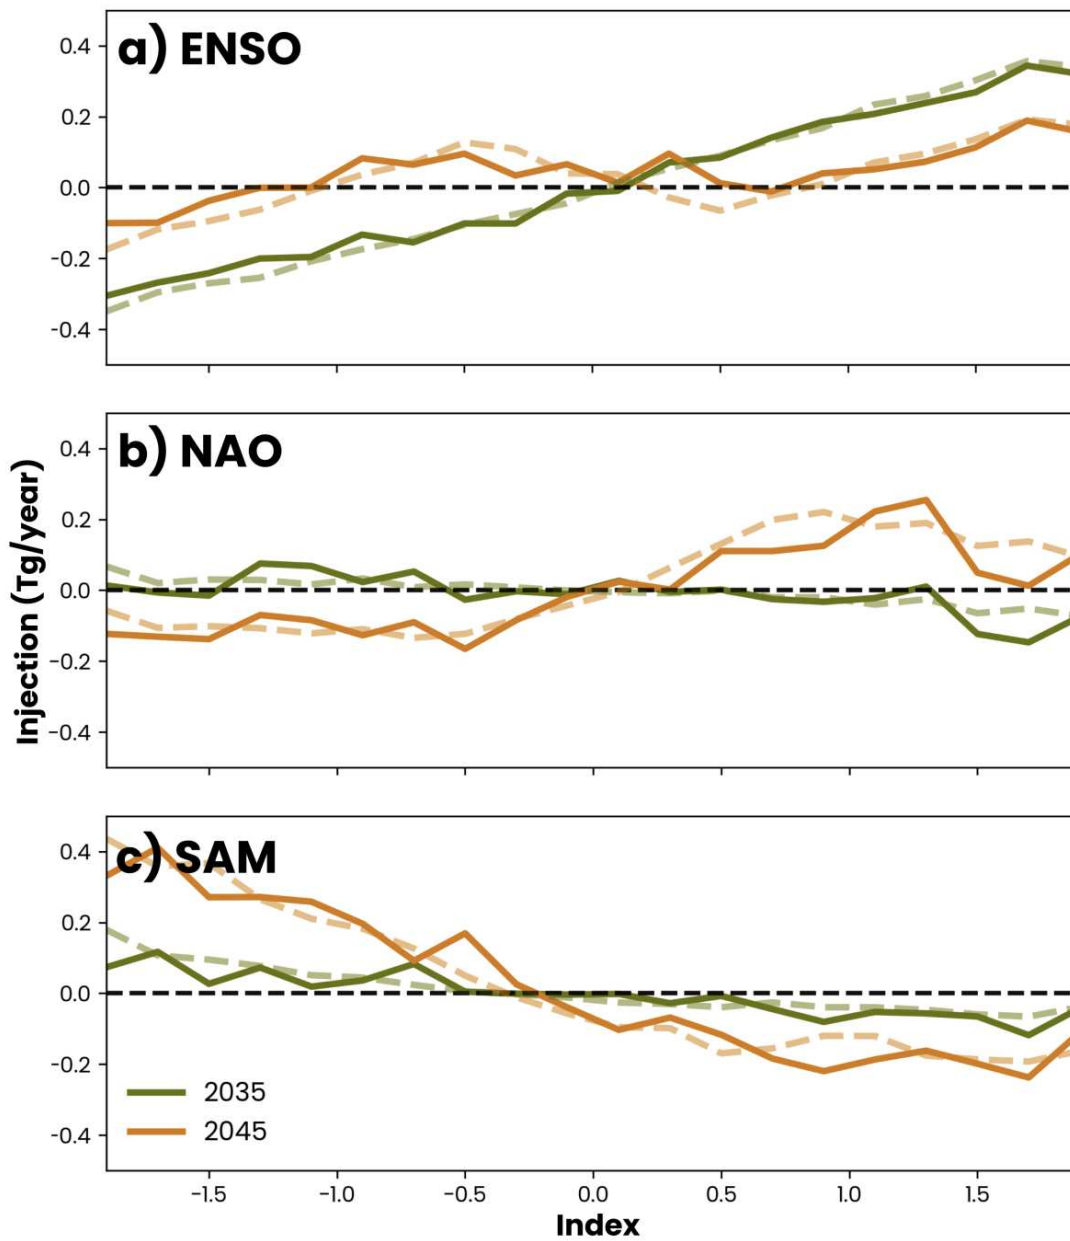


Figure A.4: The portion of the SO₂ injection (Tg/year) in response to (a) ENSO, (b) NAO, and (c) SAM using the base state years 2035 (green line) and 2045 (orange line). Solid lines use data from ARISE-SAI-CTRL and dashed lines use data from CESM2-LE.

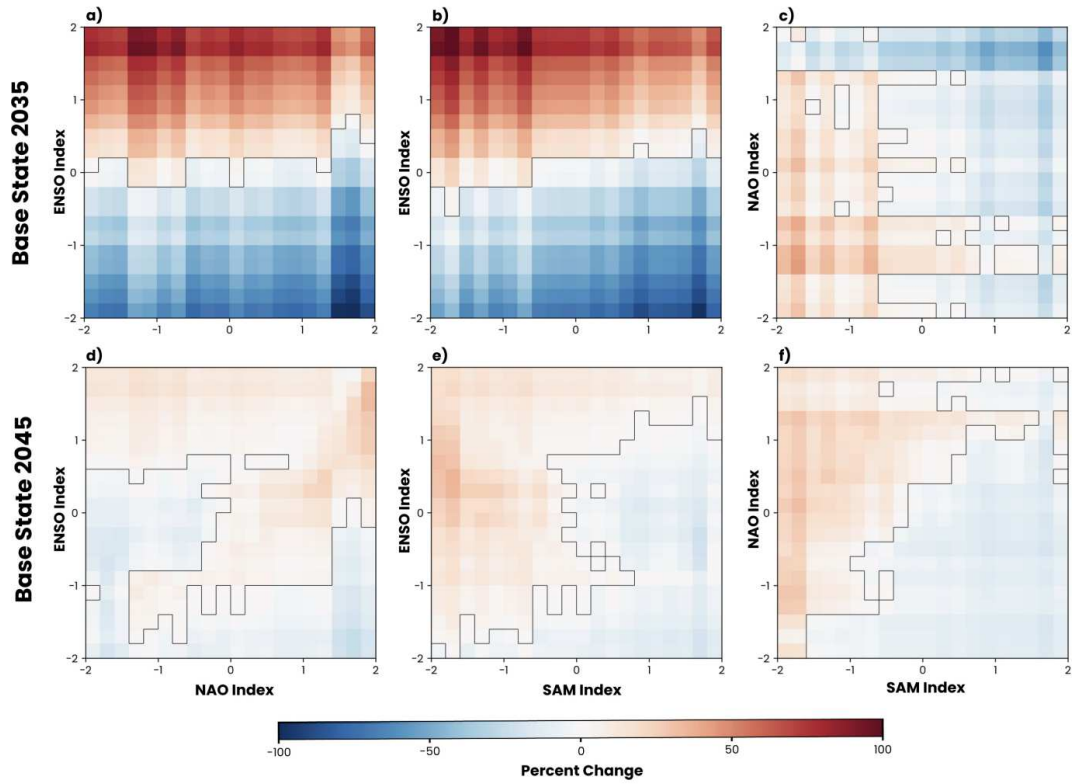


Figure A.5: Percent change in total SO₂ injection as a function of two internal variability indices but using composites from ARISE-SAI-CTRL. Top row uses the year 2035 base state and bottom row uses the year 2045 base state. Black line in each panel separates positive percent change (red shading) from negative percent change (blue shading).

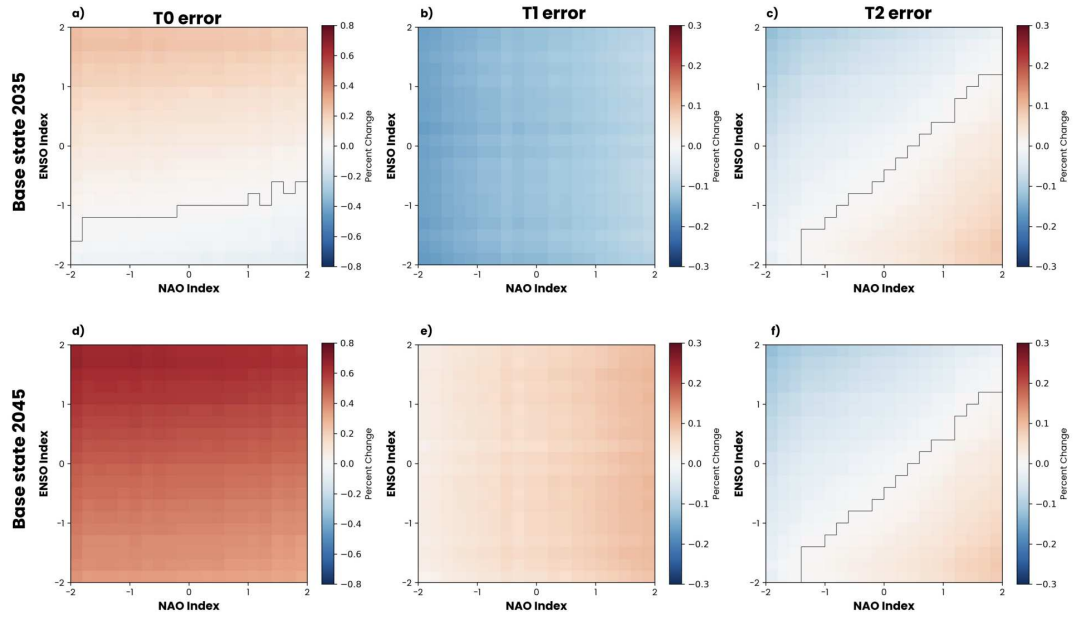


Figure A.6: Deviations in T0, T1, and T2 as a function of ENSO and NAO using composites from CESM-LE. Top row uses the year 2035 base state and bottom row uses the year 2045 base state. Black line in each panel separates positive error (red shading) from error (blue shading).




Figure A.7: Total SO_2 injection as a function of keeping either T0, T1, or T2 constant and varying the other two errors.

Appendix B

Supporting Information for Chapter 4

All machine learning models are trained with the RELU activation function [106] and using the Adam optimizer [107]. The random seed was set to 200. It was selected at the beginning of the project and was not changed other than to ensure no major sensitivity to the random seed selected. The learning rates for all machine learning models are .0001.

Following sections provide additional details about the machine learning models individually.

	Internal variability model	Carbon dioxide model	Aerosol optical depth model	Sulfur dioxide model
piControl-WACCM	X	X	X	X
piControl-CAM	X			
SSP2-4.5		X		X
CESM2-LE		X		
ARISE-SAI			X	X
CRASSULA			X	X

Figure B.1: Figure showing what simulations were used to train each individual machine learning models

B.1 Details about the internal variability machine learning model

This chapter covers the details for the internal variability machine learning model. Equation B.1 is the loss function used to train this model

$$Loss = \frac{\sum ((|Y_{true} - Y_{pred}|) * (|Y_{true}| + 1)^2 * weights_{lat} * mask_{land})}{n} \quad (B.1)$$

where Y_{true} is the truth and Y_{pred} is the prediction. n is the number of samples. $weights_{lat}$ weights the loss function by latitude and $mask_{land}$ is a mask of zeros and ones so that the model

can focus on learning the land temperature response. Figure B.2a show true and predicted global mean near surface temperature from the piControl-WACCM simulation. Figure B.2b shows the average error of the prediction compared to the error of a persistence prediction. Vertical lines indicate times when a persistence prediction outperforms the machine learning model prediction. The architecture of the internal variability machine learning model is found in Figure B.7.




Figure B.2: Panel (a) shows 50 years of global mean temperature from piControl-WACCM (black line) compared to predicted global mean temperature (red line). Panel (b) shows average errors of the machine learning model prediction (red line) compared to the error of a persistence prediction (purple line). Blue vertical lines show when a persistence prediction outperforms the machine learning model prediction.

B.2 Details about the CO₂ machine learning model

This chapter covers the details for the CO₂ machine learning model. Equation B.2 is the loss function used to train this model

$$Loss = \frac{\sum ((|Y_{true} - Y_{pred}|) * weights_{lat} * mask_{land})}{n} \quad (B.2)$$

where Y_{true} is the truth and Y_{pred} is the prediction. n is the number of samples. $weights_{lat}$ weights the loss function by latitude and $mask_{land}$ is a mask of zeros and ones so that the model can focus on learning the land temperature response. Figure B.3a show true and predicted global mean near surface temperature from the SSP2-4.5 simulation along with the prediction from the internal variability machine learning model. Figure B.3b shows the average error of the prediction compared to the error of a persistence prediction. Vertical lines indicate times when a persistence prediction outperforms the machine learning model prediction. The architecture of the CO₂ machine learning model is found in Figure B.8.

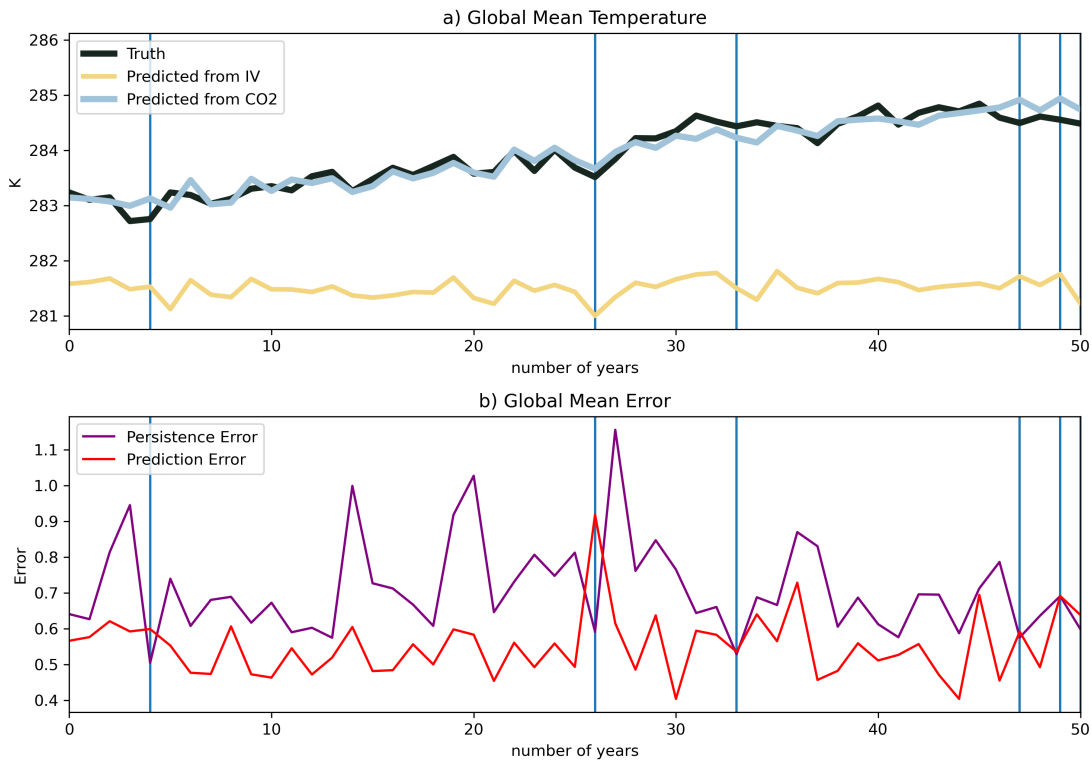


Figure B.3: Panel (a) shows 50 years of global mean temperature from piControl-WACCM (black line) compared to predicted global mean temperature (red line). Panel (b) shows average errors of the machine learning model prediction (red line) compared to the error of a persistence prediction (purple line). Blue vertical lines show when a persistence prediction outperforms the machine learning model prediction.

B.3 Details about the AOD machine learning model

This chapter covers the details for the AOD machine learning model. Equation B.3 is the loss function used to train this model. Figure B.4 shows the true and predicted global mean AOD from the CRASSULA simulation. The green line shows prediction where the SO_2 input is kept constant and only the initial AOD map changes. This shows that the model has learned that changes in AOD can persist into the following year but at reduced amounts. The blue line in Figure B.4 is the model prediction when the initial AOD map is held constant and only the SO_2 input is changed. First, the two-year injection period of CRASSULA is visible. More importantly this shows how AOD increases differently compared to the different injection scenarios. This figure shows how the SO_2 input and initial AOD map input work together when predicting the AOD for the following year.

The architecture of the AODmachine learning model is found in Figure B.9.

$$Loss = \frac{\sum ((Y_{true} - Y_{pred})^2 * weights_{lat})}{n} \quad (\text{B.3})$$

where Y_{true} is the truth and Y_{pred} is the prediction. n is the number of samples. $weights_{lat}$ weights the loss function by latitude.




Figure B.4: Predicted global average AOD (black line) compared to the true global average AOD (black line). Green line represents predictions when the SO_2 input is kept constant and only the initial AOD input is changed. The blue line represents predictions when the initial AOD input is kept constant and only the SO_2 input is changed.

B.4 Details about the SO₂ machine learning model

This chapter covers the details for the SO₂ machine learning model. Equation B.4 is the loss function used to train this model

$$Loss = \frac{\sum ((|Y_{true} - Y_{pred}|) * weights_{lat} * mask_{land})}{n} \quad (B.4)$$

where Y_{true} is the truth and Y_{pred} is the prediction. n is the number of samples. $weights_{lat}$ weights the loss function by latitude and $mask_{land}$ is a mask of zeros and ones so that the model can focus on learning the land temperature response. The architecture of the SAI machine learning model is found in Figure B.5. This model is not a typical U-net model. It contains two encoder branches that are concatenated at the center nodes. This is so the skip connection can provide information about the predicted AOD. When the skip connection included information about the temperature, the model learned to identify whether the data came from CRASSULA or ARISE-SAI-1.5, reducing the model's ability to generalize.

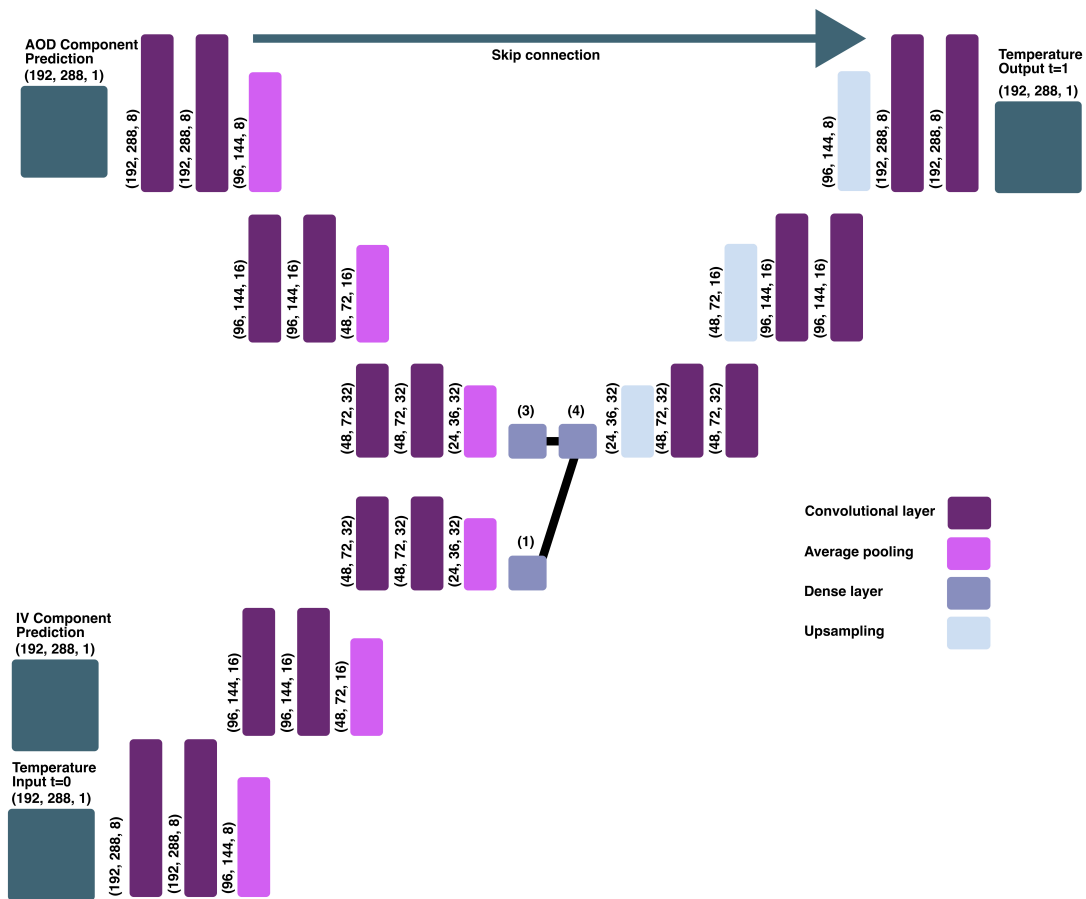


Figure B.5: Schematic of the SAI machine learning model architecture. Each layer is labeled with the input dimensions. Layers are colored to represent the type of layer. Dark purple in a convolutional layer. Pink is an average pooling layer. Purple is a dense layer. Light blue is a upsampling layer. Arrows indicate skip connection and Teal squares represent the inputs and outputs.

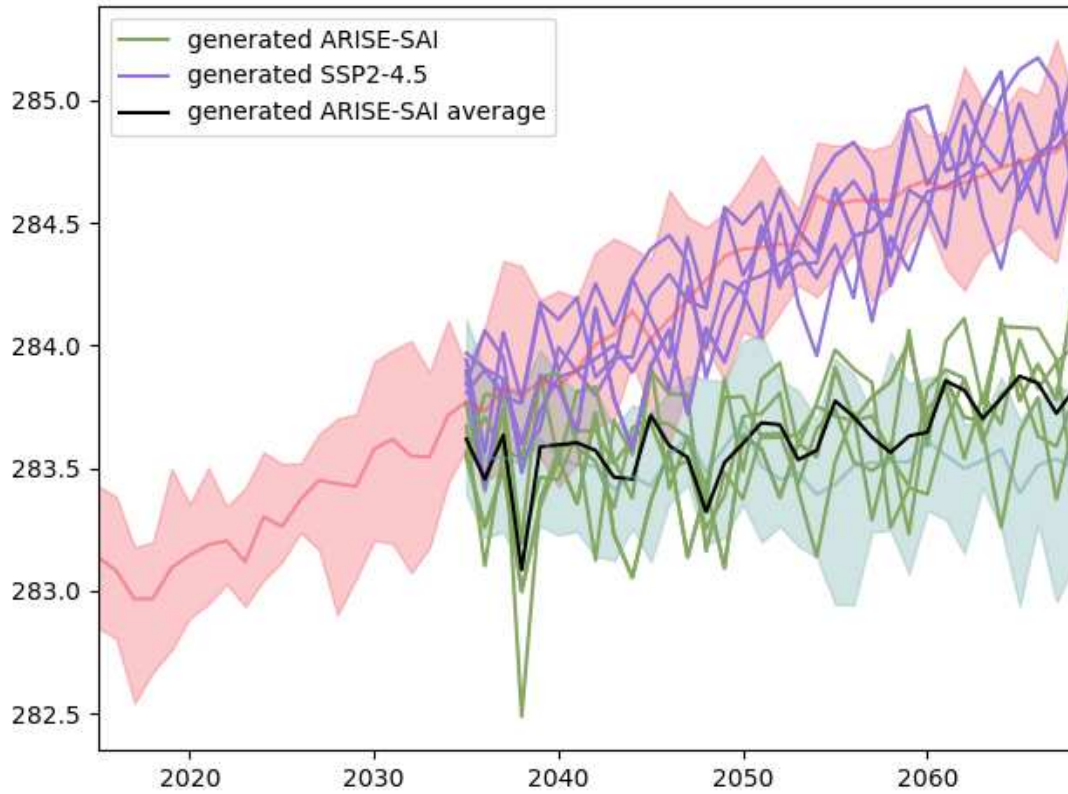


Figure B.6: Global mean land temperature for SSP2-4.5 (red shading), ARISE-SAI-1.5 (green shading), data generated by the SAI emulator to replicate SSP2-4.5 (purple lines), and data generated by the SAI emulator to replicate AIRSE-SAI-1.5 (green lines). Shaded areas represent the range between the minimum and maximum values in SSP2-4.5 and ARISE-SAI-1.5.

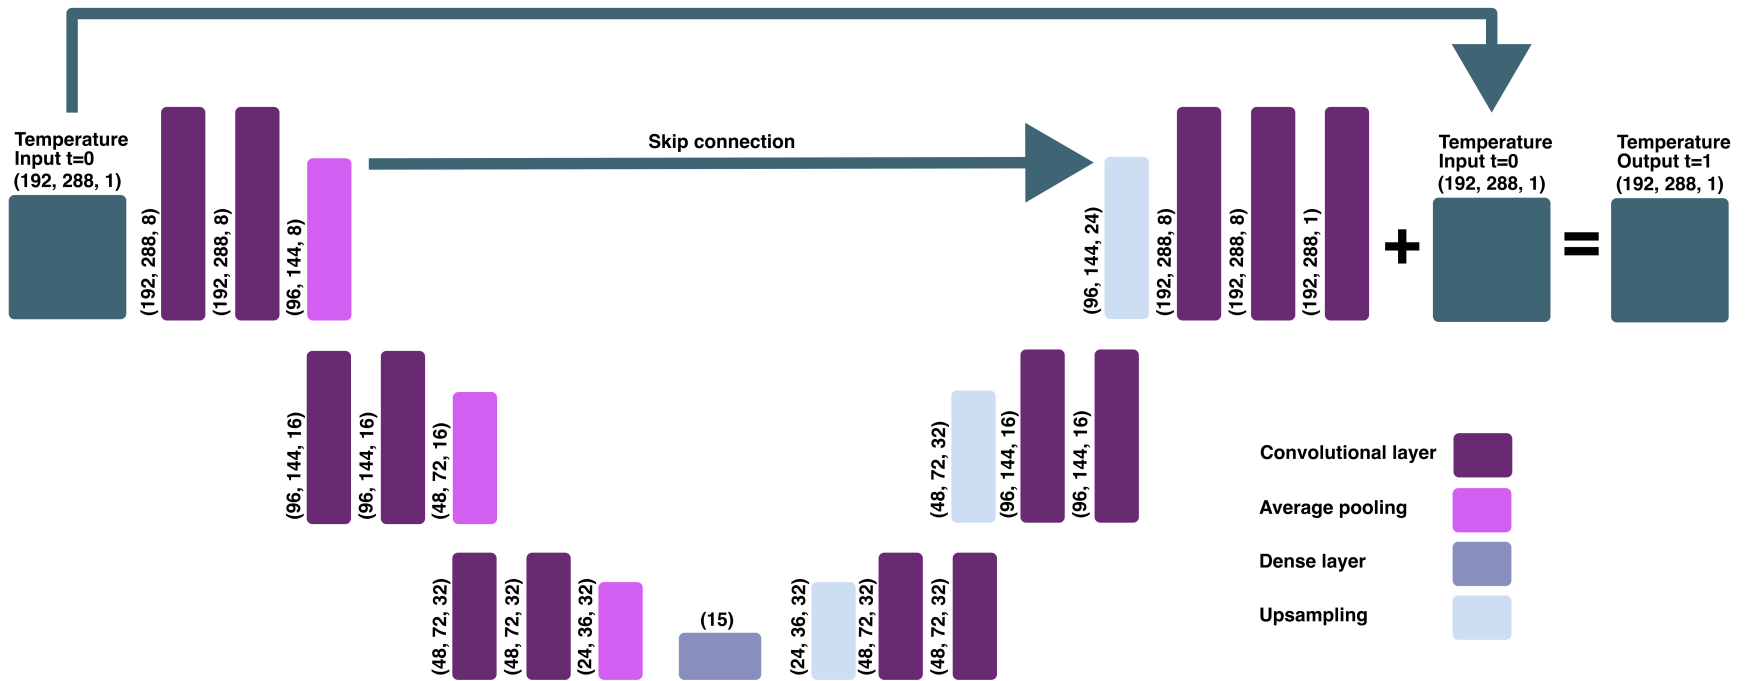


Figure B.7: Schematic of the internal variability machine learning model architecture. Each layer is labeled with the input dimensions. Layers are colored to represent the type of layer. Dark purple in a convolutional layer. Pink is an average pooling layer. Purple is a dense layer. Light blue is an upsampling layer. Arrows indicate skip connection and Teal squares represent the inputs and outputs.

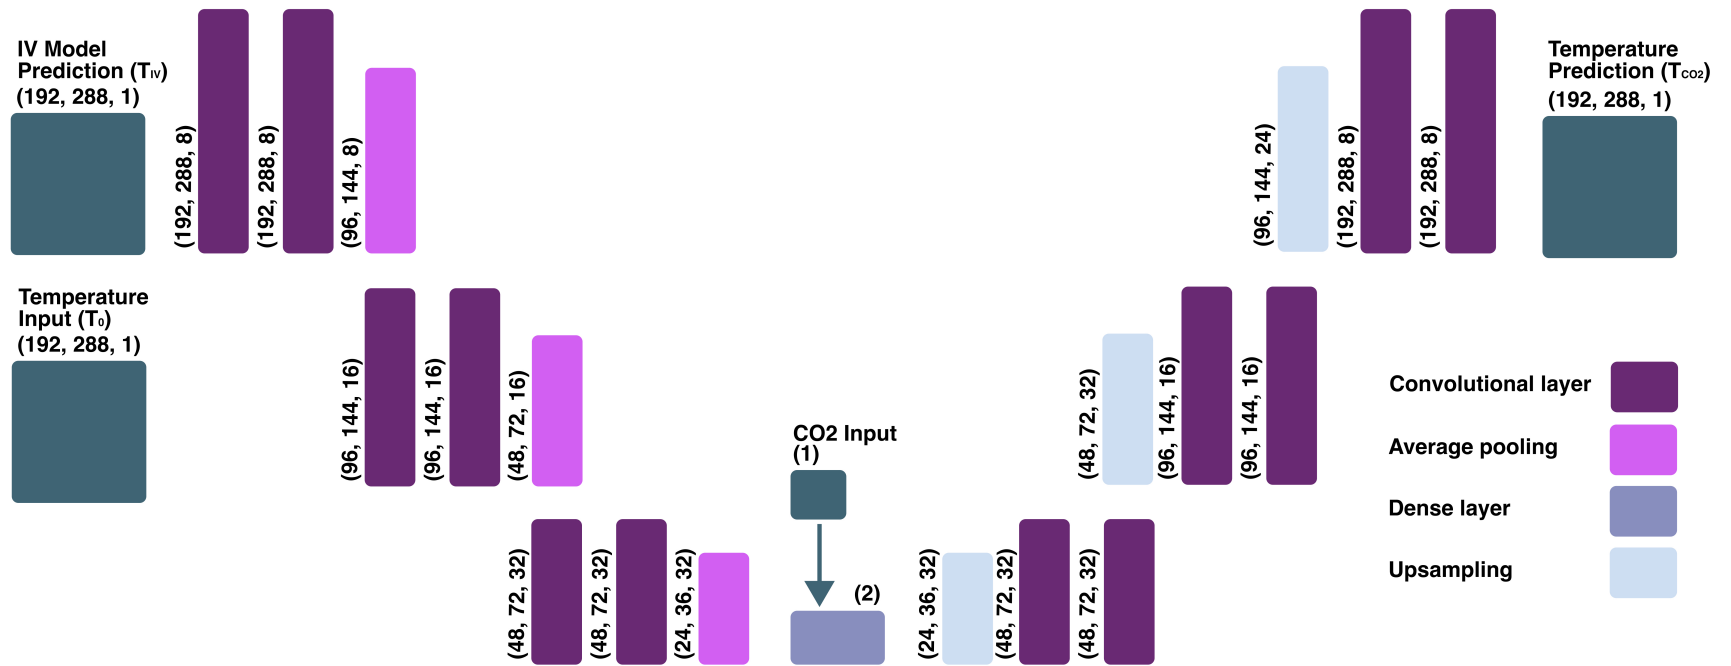


Figure B.8: Schematic of the CO₂ machine learning model architecture. Each layer is labeled with the input dimensions. Layers are colored to represent the type of layer. Dark purple in a convolutional layer. Pink is an average pooling layer. Purple is a dense layer. Light blue is an upsampling layer. Arrows indicate skip connection and Teal squares represent the inputs and outputs.

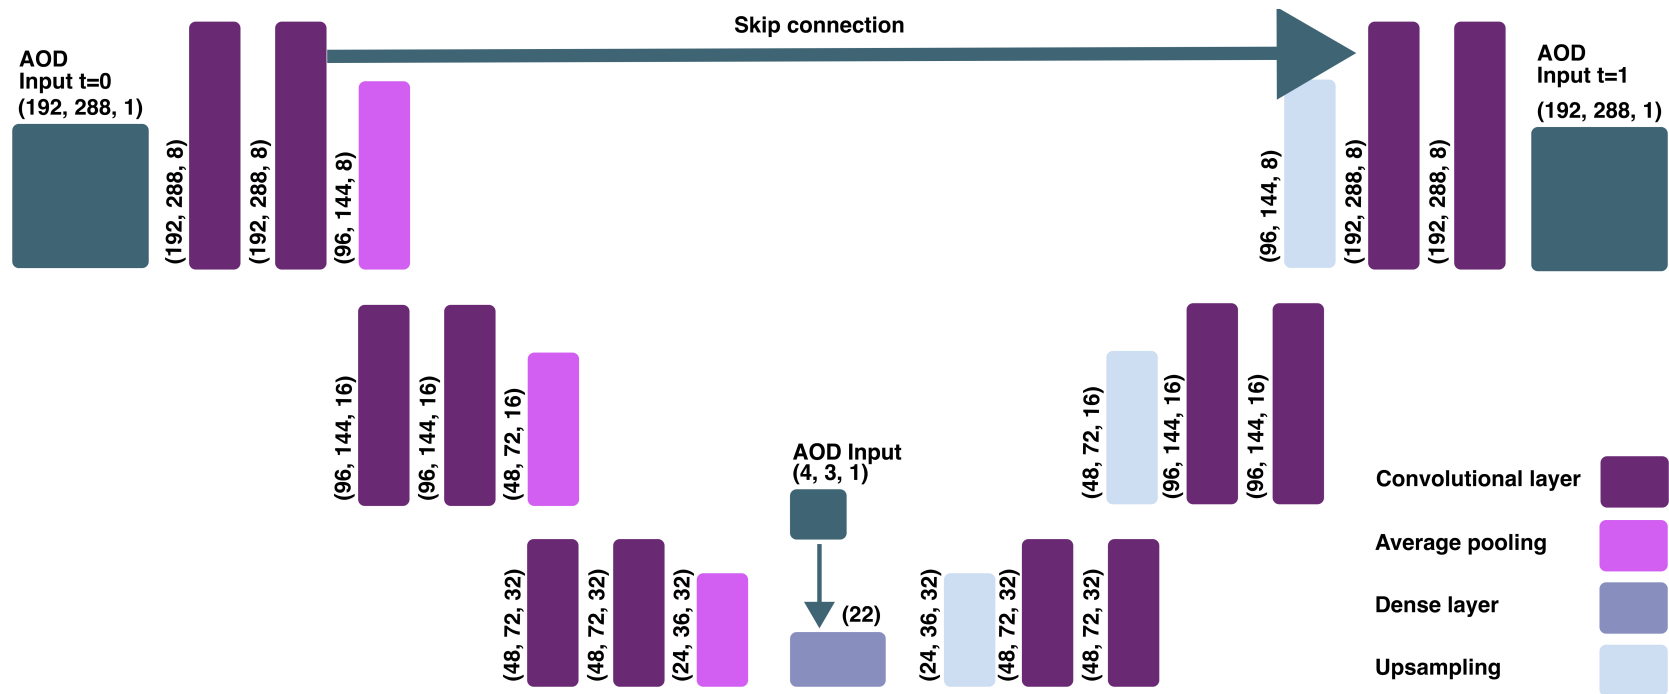


Figure B.9: Schematic of the AOD machine learning model architecture. Each layer is labeled with the input dimensions. Layers are colored to represent the type of layer. Dark purple in a convolutional layer. Pink is an average pooling layer. Purple is a dense layer. Light blue is a upsampling layer. Arrows indicate skip connection and Teal squares represent the inputs and outputs.