

SPF ICE: A Novel Approach to Model the Amount And Effectiveness of Silica to Preserve Glaciers Using Reinforcement Learning

Aadhav Prabu¹

¹Chattahoochee High School

November 22, 2022

Abstract

Glaciers cover nearly 10 percent of the earth's surface but are melting at an inexorable rate. Antarctica's Doomsday Glacier' is melting faster and could raise global sea levels by two feet. As three-quarters of the earth's fresh water is stored in glaciers, its melting depletes freshwater resources for millions of people. Glaciers also play a huge role in the climate crisis. Silica microspheres are promising materials to prevent glacier melting as it reflects most of the sun's radiation. When spread in layers over the glacier, it can slow the rate of melt and aid in new ice formation. However, currently, no modeling is available to show the amount of silica needed and its effectiveness in advance. This paper introduces a novel method SPF ICE that models the silica amount based on glacier's properties by testing reinforcement learning agents in a custom OpenAI Gym environment. The environment simulates a real-world model of a glacial setting using specific data, such as the glacier's mass balance, average accumulation, and ablation. After testing RL agents like DQN and SARSA, the proposed solution modeled the silica amount that reduced glacial melting by an average of 60.40% extending its lifetime by many years. The results indicate SPF ICE is a promising and cost-effective solution to curb glacier melting.

SPF ICE: A Novel Approach to Model the Amount And Effectiveness of Silica to Preserve Glaciers Using Reinforcement Learning

Aadhav Prabu¹

¹Chattahoochee High School, Johns Creek, GA, USA

Key Points:

- The reinforcement learning agents and a custom OpenAI Gym environment are used to determine ideal amounts of silica to preserve glaciers
- The RL agents, DQN and SARSA, act in an OpenAI Gym environment, which simulates the real-life conditions of a glacial setting
- The proposed solution reduced glacial melting by an average of 60.40% using the determined amount of silica

Abstract

Glaciers cover nearly 10 percent of the earth’s surface but are melting at an inexorable rate. Antarctica’s Doomsday Glacier’ is melting faster and could raise global sea levels by two feet. As three-quarters of the earth’s fresh water is stored in glaciers, its melting depletes freshwater resources for millions of people. Glaciers also play a huge role in the climate crisis. Silica microspheres are promising materials to prevent glacier melting as it reflects most of the sun’s radiation. When spread in layers over the glacier, it can slow the rate of melt and aid in new ice formation. However, currently, no modeling is available to show the amount of silica needed and its effectiveness in advance. This paper introduces a novel method SPF ICE that models the silica amount based on glacier’s properties by testing reinforcement learning agents in a custom OpenAI Gym environment. The environment simulates a real-world model of a glacial setting using specific data, such as the glacier’s mass balance, average accumulation, and ablation. After testing RL agents like DQN and SARSA, the proposed solution modeled the silica amount that reduced glacial melting by an average of 60.40% extending its lifetime by several years. The results indicate SPF ICE is a promising and cost-effective solution to curb glacier melting.

1 Introduction

1.1 Glacial Melting

Glaciers are receding and shrinking at a rapid phase. While melting glaciers are caused by climate change, glacier melting further increases the temperatures across the globe. The phenomenon is called ‘ice-albedo feedback’ (Curry et al., 1995). This feedback arises from the simple fact that ice is more reflective than land or water surfaces. Therefore, as global ice cover decreases, the reflectivity of Earth’s surface decreases, more incoming solar radiation is absorbed by the surface, and the surface warms. According to new research, the melting of glaciers as a result of climate change has even knocked the Earth off its axis (Deng et al., 2021).

Currently, the scale and speed of ice melting are extraordinary. In the summer of 2019, Arctic sea ice levels were tied for the second-lowest ever recorded. During a heat-wave in June, Greenland lost nearly 60 billion tons of ice in five days. In 2012, satellite observations reveal that melt occurred across 98.6% of the Greenland ice sheet (Nghiem et al., 2012). Climate models predict that Arctic summers could be close to ice-free in the next 70 years.

It is the glaciers and ice sheets that are absorbing the brunt of the climate crisis. Preserving glaciers is crucial for various environmental reasons, but the most important one is that it is the source of freshwater for millions of people around the world.

About three-quarters of Earth’s fresh water is stored in glaciers. Therefore, glacier ice is the second-largest reservoir of water and the largest reservoir of fresh water on Earth. Glaciers are critical to water management, fisheries, and flood prevention. With shrinking glaciers, less water will be available for nearby river systems when rainfall is low. In some parts of the world, millions of people could lose their primary water supplies. In the Pacific northwest US, if glaciers melted entirely, that could reduce the flow of certain watersheds by up to 15% in dry months of August and September (Menounos et al., 2019). In Asia, 700 million people will face water problems by 2100 due to melting glaciers in that region. The pace of retreat and loss of certain glaciers is most rapid within the Tropical Andes (Johansen et al., 2018). This melting of South America’s glaciers and ice fields poses a threat to water supplies and agriculture from Bolivia to Chile.

Rising sea levels can also introduce new or exacerbate existing saltwater intrusion into freshwater resources. Both groundwater and surface water sources are at risk in coastal

cities posing challenges for drinking water treatment facilities and water resource managers. Melting glaciers are contributing to rising sea levels flooding coastal cities throughout the world (Cazenave et al., 2018). Glaciers had predictable seasonal changes, losing mass in the summer and regaining it in the winter. In recent years they are losing more than they accumulate through new snowfall, ultimately adding more water to the oceans, leading to a rise in sea level. Global mean sea level has risen about 8–9 inches since 1880, with about a third of that coming in just the last two and a half decades, and from 2018 to 2019, global sea level rose 0.24 inches (Douglas et al., 2000).

1.2 Silicon Dioxide

Silica, silicon dioxide, is a compound of the two most abundant elements in Earth’s crust: silicon and oxygen. The mass of Earth’s crust is nearly 59 percent silica, and it is the main constituent of more than 95 percent of the identified rocks. Silica can reflect most of the radiation from the sun’s rays, making it an optimal option to prevent glacier melting. Also, it sticks to ice and water the moment it hits the surface. When sprayed over water, the reflective sand creates a white slush that mimics the reflective properties of ice meaning that heat from the sun can be reflected outward rather than being absorbed into the ice and sea. It is chemically unreactive, which means it is not prone to a chemical reaction. Since it is hydrophilic, it does not attract any oil-based pollutants. Sand silica can benefit the global silica cycle and ecosystems as long as its size is not large enough to be deemed harmful. Most silica microspheres average between 35–60 micrometers above the health risk threshold (Flörke et al., 2008). This choice is also safe in desired amounts for animals and ecosystems.

1.3 Proposed Solution

No single glacier or ice sheets are similar in mass balance, debris, density, thermal conductivity, and absorption; a single standard approach of preserving it with silica is not ideal. The proposed solution is novel as it reduces glacial melt using silica intelligently, taking into account all the characteristics and properties of the glacier. Furthermore, current machine learning algorithms only provide the evolution and mapping of glaciers (Bolibar et al., 2020) with no solution providing intervention and treatment of glacier melting like SPF ICE. With this novel approach using reinforcement learning, an area of machine learning, the glacial melt can be efficiently reduced by determining the amount of silica needed for adequate reflection of UV rays. As users enter specific properties of a glacier and additional metrics like temperature, average accumulation, and average ablation of the area, the algorithm will accurately determine the amount of silica desired to prevent the melt. Determining the amount of silica is crucial as it not only very cost-effective but also helps reduce any effects of silica on the environment.

2 Related Work

The Arctic Ice Project utilizes silica beads that were tested in Alaska have shown promising results. In a paper published by the American Geophysical Union, one field test reported a 15 to 20 percent increase in reflectivity due to the beads. In the Arctic, that could translate into a 1.5 degree Celsius temperature reduction, a 3-degree reduction in sea temperatures, and an increase in ice thickness up to 20 inches (Field et al., 2018). The Arctic project solution strategically applied in the Arctic can allow the world to buy up to 15 more years to decarbonize the economy and draw down greenhouse gas from the atmosphere (Chamberlin et al., 2020). Since there is no modeling to forecast the amount and effectiveness of silica, SPF ICE can help intelligently and quickly select the amount of silica through glacial modeling.

3 SPF ICE: Overview of Structure

The project’s software development occurred in two phases: the construction of the custom OpenAI Gym environment and the reinforcement learning (RL) agents, Deep Q Network, and SARSA.

3.1 OpenAI Gym Environment

3.1.1 Initialization and Timesteps

The first component is the custom OpenAI gym environment, which simulates the real-life conditions of a glacial setting. The initialization of every OpenAI gym environment consists of an observation space, action space, and the environment’s current state. In this scenario, the observation space is the observed mass balance of the glacier, and the action space is represented as a Box, an array of integers from 1-20, referring to the thickness of silica in centimeters. The current state of the environmental attributes to the annual melt rate of the glacier. Other factors defined are the season of the year, temperature, average accumulation, and average ablation. Users can enter values specific to their glacier or use predetermined values for prominent glaciers across the world like Matanuska, Mendenhall, Vatnajökull, and the Lambert glacier.

After the initialization of the environment is complete, each timestep must be defined, which aligns with the seasons of the year. Since glacial conditions are not similar during fall-winter and spring-summer, the timesteps are divided into these two groups. For the fall-winter timesteps, as snowfall is more likely, the accumulation rate is added to the average melt rate. Based on the amount of silica chosen by the RL agent, a new melt rate is calculated using the conductive heat flux formula: $Q_c = k\Delta T/h_d$. The melt rate is then calculated by $M = Q_c/L_f$ (Hock, 2005).

3.1.2 Reward System

The reward system is calculated based on how effective the silica is in preventing additional melt. For this research, it is hypothesized that silica could reduce glacial melt by greater than 50%. Therefore, if silica can reduce the melt rate by at least 50% than the current rate, a positive reward is given. However, to constrain the use of an excessive amount of silica, the size of the positive reward is inversely proportional to the thickness of silica. Using more silica yields a smaller positive reward, and using less silica returns a higher reward. If the silica cannot reduce the melt rate by at least 50% of the predicted melt rate, the silica is given a negative reward. The size of the negative reward is directly proportional to the amount of silica used. Using more silica returns a larger negative reward, and using less silica yields a smaller negative reward. In both cases, using more silica is punished more severely than using less silica.

As the ultraviolet radiation intensity increases during the spring-summer timestep, the additional ablation melt is added to the annual melt rate. The new melt rate and the reward system are calculated using the amount of silica similar to the fall-winter timestep. Each episode starts at the defined mass of the glacier and ends when the glacier’s mass reaches zero. After the end of each episode, the environment is reset.

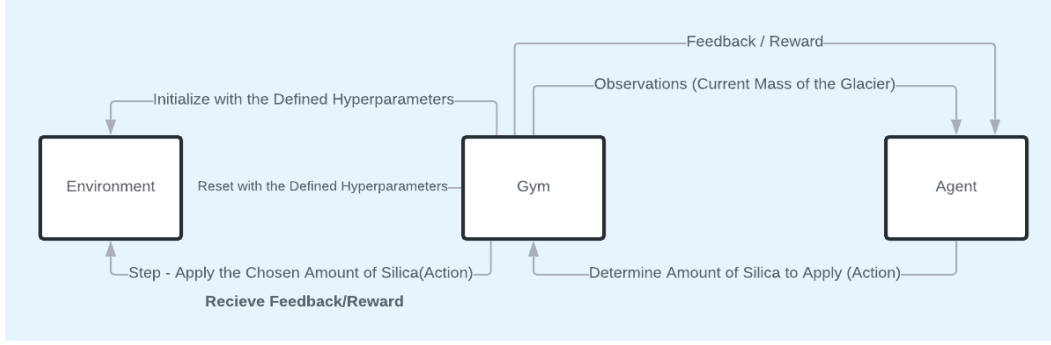


Figure 1: Interaction between RL Agent and Environment

3.2 DQN RL Agent

The second component of the software is the reinforcement learning agents: Deep Q Network and SARSA. DQN agent uses two neural networks: the main and target network. Both networks have the same architecture but use different weights to provide stability to the learning process. The neural networks map the input state to the (action, Q-value) pairs (Fan et al., 2020). The neural network architecture consists of five layers, including three hidden dense layers of 256 units with the ReLU activation function. The final layer has 20 units for each thickness of silica that could be applied. Using the current state as its input, DQN uses the Boltzmann policy to output the Q-values for all possible actions. The action associated with the highest Q-value is chosen. The agent's decisions or actions will affect the rewards it obtains. During each episode, Deep Q Network attempts to maximize the rewards that it receives.

3.3 SARSA RL Agent

The SARSA agent works differently from the DQN agent. It uses an on-policy learning algorithm, where in the current state (S), it chooses the best possible action (A) and receives a reward (R). It arrives in a new state (S1) and takes action (A1) in that state, creating the tuple (S, A, R, S1, A1). The Q values represent the possible reward in the next time step after taking the chosen action in the current state, plus the discounted future reward received in the next state. The Q-values are updated based on the action A1 taken in state S1. SARSA also attempts to maximize the rewards that it receives.

4 Proof Of Concept

4.1 Results

Below are results averaged over 32 episodes for comparing both agents to the mass of the glacier over time and their rewards to the most optimal policy using the following hyperparameters. These hyperparameters are approximate values of the Vatnajökull glacier found in southeastern Iceland.

Total Mass of Glacier = 3,000,000,000 Tonnes
 Total Area of Glacier = 8,000,000 Meters
 Average Accumulation = 17,500 Tonnes
 Average Ablation = 20,000 Tonnes
 Thermal Conductivity of Silica = 1.11
 Latent Heat of Energy of Silica = 50.55




Figure 2: Plot comparing the glacial melt with and without silica applied the DQN agent. Without silica, the glacier's lifetime is about 4058 time-steps or about 1014.70 years. Using the amount of silica applied by the DQN agent, the lifetime of the glacier is extended by 26549.55 years.




Figure 3: Plot comparing the glacial melt with and without silica applied the SARSA agent. Without silica, the glacier's lifetime is about 4048 time-steps or about 1012 years. Using the amount of silica applied by the SARSA agent, the lifetime of the glacier is extended by 17439.25 years.




Figure 4: Figure 4.a compares the average melt rate applying silica using DQN compared to the average melt rate without silica. Figure 4.b presents evidence that silica decreases the glacier's melt rate between 1500% and 1000%. Time-Steps, where accumulation occurred, were removed.




Figure 5: Figure 5.a compares the average melt of the glacier applying silica using the SARSA agent compared to the average melt rate without it. Figure 5.b presents evidence that silica decreases the glacier's melt rate between 1250% and 500%. Time-Steps, where accumulation occurred, were removed.

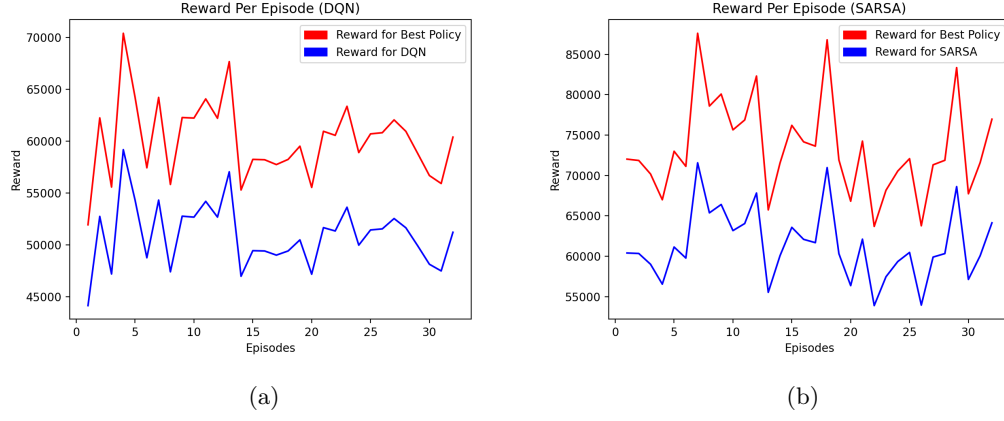


Figure 6: Both DQN and SARSA achieved high amounts of reward. Figure 6.a displays the average reward for DQN and Figure 6.b for SARSA. The SARSA agent typically achieved higher rewards than the DQN agent.




Figure 7: Plots showing the usage of silica over time. As the glacier continues to recede in size, less silica is used to prevent glacial melt. Figure 7.a shows the amount of silica used by DQN and 7.b for SARSA. The SARSA agent uses less silica than the DQN agent.

The agents were also tested changing the following hyperparameters, which are representative of cirque glaciers (small glaciers found in bowl-shaped depressions near mountains). The agents were only tested across 100 time-steps.

Total Mass of Glacier = 2000 Tonnes

185
186

Average Accumulation = 10 Tonnes
Average Ablation = 25 Tonnes




Figure 8: Figure 8.a compares the average melt rate applying silica using the SARSA agent (blue) compared to the average melt rate without it (red) across 100 episodes. As evident in Figure 8.b, using silica significantly reduced the melt rate.




Figure 9: Figure 9.a compares the average melt of the glacier applying silica using the DQN agent compared to the average melt rate without it. Figure 9.b presents evidence that silica decreases the glacier's melt rate between 58.75% and 70.35%.




Figure 10: The most optimal policy obtains the highest amount of reward possible in an episode. Below the best policy is indicated in red and the agents in blue. Figure 10.a displays the average reward for DQN and Figure 10.b for SARSA.

4.2 Discussion

As shown in the results section, both reinforcement learning agents reduced the glacier's melt rates substantially from the current melt rate. More importantly, silicon dioxide had a bigger impact on preserving the larger glacier compared to the smaller one. Both agents reduced the melt rate by nearly 1000% for the larger glacier compared to 60% for the smaller glacier. For the Vatnajökull glacier, the DQN agent performed better than the SARSA agent. DQN was able to extend the lifetime of the glacier by nearly 9,000 years because it reduced the melt rate significantly more than SARSA. However, DQN uses slightly more silicon dioxide than SARSA as indicated in Graph 8.a. For the smaller glacier, silicon dioxide had a smaller effect. The SARSA agent reduced the melt rate from an average of 175.54 to 64.29 using silicon dioxide. This difference amounts to an average decrease of 63.38% in melt rate. The DQN agent achieved similar results to the SARSA agent. Deep Q Network reduced the melt rate from an average of 172.83 to 68.44, an average decrease of 60.40% in the glacial melt. Although Deep Q Network performs slightly lower than the SARSA agent, DQN had less variability in percent difference than SARSA. This statistic could be critical, especially in real-world glacial environments, where minor differences in melt could lead to severe consequences, such as rising sea levels, habitat loss, and loss of glacier stability. Both agents were successful in their usage of silicon dioxide. They used silicon dioxide proportional to the melt rate of the glacier as shown in Figure 11.a and 11.b. SPF ICE is able to model the mass of the glacier over time and provide the amount of silica to preserve glaciers at a specific point in time.

For further analysis of the RL agents, their reward per episode can be compared to the most optimal policy. The best policy achieves the highest amount of reward per episode. Therefore, it is imperative that both the DQN and SARSA agents have policies that align most closely to the optimal policy. When comparing the RL agents to the most optimal policy, Deep Q Network performs better than SARSA. DQN can better maximize the amount of reward that it receives than SARSA. All of the graphs indicate that DQN is better suited in determining the correct amount of silica to reduce glacial melt across the world.

Future work includes gathering and testing more real-time glacier data from National Snow and Ice Data Center to augment the modeling to provide the best perfor-

mance in larger scale field tests. This data set consists of glacier regime parameters observed between 1945 and 2003. Data include annual mass balances, ablation, accumulation, and equilibrium-line altitude of mountain and subpolar glaciers outside the two major ice sheets. All available sources of information, such as publications, archived data, and personal communications, have been collected and include time series of more than 300 glaciers. Data have been digitized and quality checked (Dyurgerov, 2002). The data collected will be fine-tuned based on the year, categorized based on glaciers' specific properties such as average mass balance, accumulation, and ablation, and then inputted into the OpenAI Gym environment. Using these values, the RL agent will be tested over several episodes.

The values defined in the SPF ICE's OpenAI Gym are exact and are not variable to change over time. However, this is not typical in the real-world glacial setting. There will be more variability and randomness in a real-world glacial environment, increasing the glacial melt rate and reducing the lifetime of the glaciers. Modeling additional properties like temperature and solar reflectivity will make SPF ICE more accurate.

5 Conclusion

SPF ICE is a revolutionary and novel solution to preserve glaciers worldwide using reinforcement learning and OpenAI Gym to determine the amount of silica needed for adequate reflection of UV rays based on the season, temperature, and mass balance of the glacier without using thousands of time-consuming calculations. Determining the right amount of silica is crucial as it is cost-effective and helps reduce any adverse effects of silica on the natural environment. SPF ICE is customizable for any glacier as users can enter data specific to their glaciers, such as average mass balance, accumulation, and ablation. It is scalable to easily add more prevention techniques for glacial melting, modifying only specific formulas and values defined in the OpenAI Gym environment. With its immense benefits, this accurate and effective solution can be helpful for people around the world that depend heavily on glacial freshwater, prevent flooding of coastal cities, and helps reduce extreme weather events.

Acronyms

RL Reinforcement Learning
SARSA State-action-reward-state-action
DQN Deep Q Network

Acknowledgments

Data used in the research is from National Snow and Ice Data Center. This data set consists of glacier properties observed between 1945 and 2003. Data include annual mass balances, ablation, accumulation, and equilibrium-line altitude of mountain and subpolar glaciers. Dataset accessed [2021-04-09] at <https://nsidc.org/data/G10002/versions/1>.

References

- Bolibar, J., Rabatel, A., Gouttevin, I., Galiez, C., Condom, T., & Sauquet, E. (2020). Deep learning applied to glacier evolution modelling. *The Cryosphere*, 14(2), 565–584.
- Cazenave, A., Palanisamy, H., & Ablain, M. (2018). Contemporary sea level changes from satellite altimetry: What have we learned? what are the new challenges? *Advances in Space Research*, 62(7), 1639–1653. doi: 10.1016/j.asr.2018.07.017

- Chamberlin, D., Hirtzer, P., Sholtz, A., Player, T., & Field, L. (2020). Materials design considerations of hollow glass microspheres for albedo enhancement of arctic ice. *Earth and Space Science Open Archive*, 12. doi: 10.1002/essoar.10505054.1
- Curry, J. A., Schramm, J. L., & Ebert, E. E. (1995). Sea ice-albedo climate feedback mechanism. *Journal of Climate*, 8(2), 240–247. doi: 10.1175/1520-0442(1995)008<0240:siacfm>2.0.co;2
- Deng, S., Liu, S., Mo, X., Jiang, L., & Bauer-Gottwein, P. (2021). Polar drift in the 1990s explained by terrestrial water storage changes. *Geophysical Research Letters*, 48(7), e2020GL092114. doi: <https://doi.org/10.1029/2020GL092114>
- Douglas, B., Kearney, M. T., & Leatherman, S. P. (2000). Sea level rise: History and consequences. *Journal of Coastal Research*, 16(4).
- Dyrurgerov, M. (2002). Glacier mass balance and regime measurements and analysis, 1945–2003. *Glacier Mass Balance and Regime Measurements and Analysis*. Retrieved from <https://nsidc.org/data/G10002/versions/1> doi: 10.7265/N52N506F
- Fan, J., Wang, Z., Xie, Y., & Yang, Z. (2020, 10–11 Jun). A theoretical analysis of deep q-learning. In A. M. Bayen et al. (Eds.), *Proceedings of the 2nd conference on learning for dynamics and control* (Vol. 120, pp. 486–489). The Cloud: PMLR.
- Field, L., Ivanova, D., Bhattacharyya, S., Mlaker, V., Sholtz, A., Decca, R., ... Katur, K. (2018). Increasing arctic sea ice albedo using localized reversible geoengineering. *Earth's Future*, 6(6), 882–901. doi: <https://doi.org/10.1029/2018EF000820>
- Flörke, O. W., Graetsch, H. A., Brunk, F., Benda, L., Paschen, S., Bergna, H. E., ... et al. (2008). Silica. *Ullmann's Encyclopedia of Industrial Chemistry*. doi: 10.1002/14356007.a23.583.pub3
- Hock, R. (2005). Glacier melt: a review of processes and their modelling. *Progress in Physical Geography: Earth and Environment*, 29(3), 362–391. doi: 10.1191/0309133305pp453ra
- Johansen, K. S., Alfthan, B., Baker, E., Hespings, M., Schoolmeester, T., Verbist, K., et al. (2018). *The andean glacier and water atlas: the impact of glacier retreat on water resources*. UNESCO Publishing.
- Menounos, B., Hugonnet, R., Shean, D., Gardner, A., Howat, I., Berthier, E., ... Dehecq, A. (2019). Heterogeneous changes in western north american glaciers linked to decadal variability in zonal wind strength. *Geophysical Research Letters*, 46(1), 200–209. doi: <https://doi.org/10.1029/2018GL080942>
- Nghiem, S. V., Hall, D. K., Mote, T. L., Tedesco, M., Albert, M. R., Keegan, K., ... Neumann, G. (2012). The extreme melt across the greenland ice sheet in 2012. *Geophysical Research Letters*, 39(20). doi: 10.1029/2012gl053611