LA-UR-22-30051

Title: Inferring Water Content from Neutron Die-away Curves for Planetary Science Applications

Author(s):
Chase, Eve Adde
Grieve, Tristan Sumner
Shin, Tony Heong Shick
Nowicki, Suzanne Florence
Pinilla-Orjuela, Maria Isabel

Intended for: Report

Issued: 2022-11-18 (rev.2)

Inferring Water Content from Neutron Die-away Curves for Planetary Science Applications

Eve A. Chase
Avery S. Grieve
Tony H. Shin
Suzanne Nowicki
Maria Pinilla-Orjuela

LA-UR-22-30051
September 29, 2022

Los Alamos
NATIONAL LABORATORY

Contents

Acknowledgements . . . . **v**

1 Introduction **1-1**

2 Inferring Water Content from Neutron Die-away Curves **2-1**
  2.1 Uncertainty Quantification 2-2
    2.1.1 MCNP Statistical Uncertainty 2-2
    2.1.2 Poisson Uncertainty Assuming a Realistic Detector 2-3
  2.2 Example Water Content Recovery 2-4

3 Accounting for Variations in Boron Content **3-1**

4 Refining Water Content Estimates with Gamma Spectroscopy **4-1**

5 Conclusions & Future Work **5-1**

Figures

2-1 Neutron die-away curves for various water contents 2-1
2-2 m_W for water only 2-4
2-3 Scatterplot of water content estimates while only varying water content 3-1
3-1 Boron and chlorine cross sections 3-2
3-2 Die-away curves for varying boron content 3-3
3-3 Variations in m_W for different boron concentrations 3-4
3-4 Scatterplot of water content upper limit estimates, including variations in boron . . . 4-1
4-1 Gamma spectra for 5% and 20% water content 4-2
4-2 m_B for different water contents 4-4
4-3 Water content estimate from both neutron and gamma measurements 5-1

Acknowledgements

Template made by Tyler Borgwardt, then edited by Tony H. Shin on 05/17/2021 and Eve A. Chase on 04/27/2022. This research was supported by the Los Alamos National Laboratory (LANL) through its Center for Space and Earth Science (CSES). CSES is funded by LANL's Laboratory Directed Research and Development (LDRD) program under project number 20180475DR.

1 Introduction

Signatures of liquid water on planetary bodies may provide evidence for past, present, or future life elsewhere in the solar system. Multiple current and upcoming planetary science missions are prioritizing the search for water, using innovative technologies and creative algorithms. For this project, we explore the capabilities of an instrument similar to the Dragonfly spacecraft [1] to map water content on the surface of Titan, the largest moon of Saturn. We simulate Dragonfly's neutron observations for a variety of soil compositions and develop multiple techniques to extract water content from the resultant neutron data. Additionally, we supplement neutron observations with gamma spectroscopy in an attempt to further refine water content estimates. This work is part of a larger study to employ Gaussian process regression to estimate a map of water content along the surface of Titan [4, 7, 8]. Using the estimated map, an algorithm based on prediction difference mapping is employed to optimally move multiple Dragonfly detectors around the surface of Titan, fully autonomously.

Dragonfly is scheduled to land on Titan in 2034, then spend the remainder of its mission exploring the Shangri-La dune fields and Selk impact crater [1]. Multiple scientific instruments are onboard, including the Dragonfly Gamma-ray and Neutron Spectrometer (DraGNS), which contains a pulsed neutron generator, gamma-ray spectrometer, and neutron spectrometer. We are specifically interested in studying neutron die-away curves recorded at various locations on Titan's surface, assuming a range of water content values at each location. Die-away curves record the time-dependent flux of detected neutrons, following an initial neutron pulse at the source. These die-away curves primarily include down-scattered neutrons off of the soil sample.

We use Monte Carlo N-Particle®[1] Transport Code (MCNP 6.3) [9] to simulate neutron die-away curves and gamma-ray spectra for a variety of water compositions. The MCNP simulations utilize a simple geometry with a slab of soil material surrounded by an atmosphere region. The source neutrons are 14.1 MeV and are started at time=0 in the simulation.[2] The isotropic source is directionally biased towards the soil, and further variance reduction is achieved with forcing neutron collisions as well as using type 5 tallies. Due to directional biasing, particles are more frequently sampled in the direction of interest, at the cost of the statistical weight of the particle. This allows a source to remain isotropic in the weight-space of the problem, while more tracks are directed toward important regions. The MCNP type 5 tally is a deterministic estimate of particle flux (or fluence) at a point following a particle's birth and all collision events until the particle terminates [9]. This deterministic estimate is appropriate when there is little to no self-shielding effects from a detector; additionally, this estimate provides a smaller variance in reported tally bin means. Since details of the DraGNS detector are not publicly available, we assume a neutron source and detector broadly similar to the DAN experiment on-board the Mars Science Laboratory rover *Curiosity* [6].

Several previous planetary science missions relied on passive neutron detection – the detection of neutrons from cosmic ray spallation – for neutron-based composition studies. The DAN instrument was the first planetary science mission to employ active neutron detection, from a pulsed neutron generator, to infer surface composition. Litvak et al. 2008 [6] led the first study to simulate water content inference on Mars from die-away curves, assuming an otherwise fixed composition. Later, Hardgrove et al. 2011 [5] improved upon this method and also studied the effect of variations in chlorine and iron composition on Martian water content identification. Our study builds upon the

---

[1] MCNP® and Monte Carlo N-Particle® are registered trademarks owned by Triad National Security, LLC, manager and operator of Los Alamos National Laboratory.

[2] See Reference [4] for a discussion of more realistic pulse structure.

work of Hardgrove et al. 2011 [5] by exploring a wider range of chlorine compositions, in addition to a wide range of boron compositions. Additionally, we present a thorough investigation coupling neutron die-away curve data with gamma spectroscopy to further refine water content estimates.

This report is structured as follows: in Section 2 we present a method to infer water content from neutron detections for a fixed soil composition with only variations in water content. In Section 3 we vary boron-10 composition in addition to water content, and in Section 4 we present a method to combine gamma observations with neutron die-away curves to refine estimates on water content. In Section 5 we conclude and suggest directions for future work.

2 Inferring Water Content from Neutron Die-away Curves

Detailed analysis of neutron measurements may reveal underlying molecular compositions in soil samples, including water content of the soil. For this project, we assume a volume of planetary soil is analyzed through active neutron die-away measurements and subsequently attempt to extract water content from the observations. We assume water content refers to the percent of water by mass fraction, in a given sample.

We present an example of simulated neutron die-away curves for various water content percentages in Figure 2-1, where each color corresponds to a different water content ranging from 0% to 50% by mass. We assume an identical initial neutron pulse for each die-away curve, and assume a fixed lunar soil elemental composition for all isotopes other than water. Additionally, we only consider thermal neutrons, with energies below 0.4 eV. The die-away curves in Figure 2-1 are dominated by an interplay between neutron scattering and capture reactions. Higher water content correlates with an increase in neutron scattering, which increases the number of thermal neutrons in the system. With this increase in available thermal neutrons, the rate of thermal neutron capture also increases, resulting in fewer neutrons observed at later times.

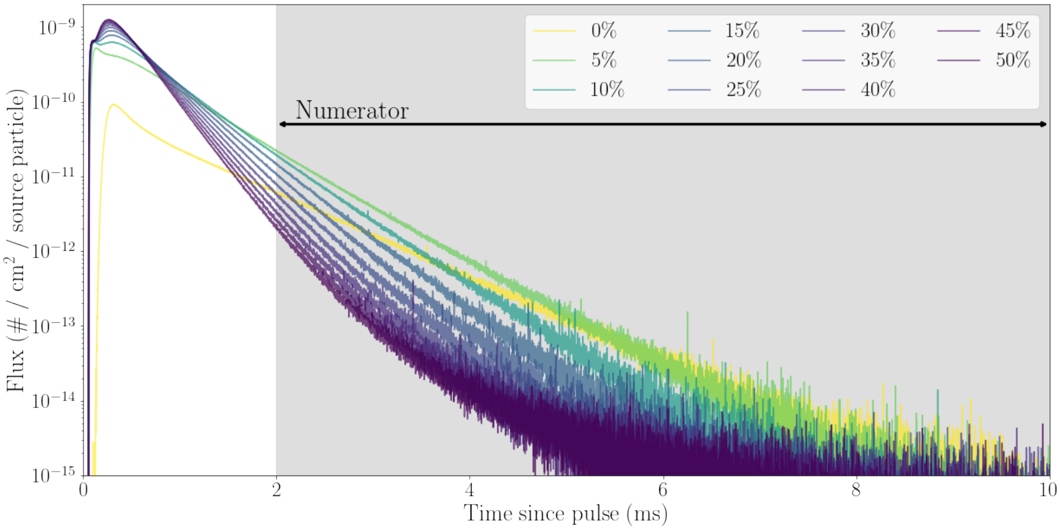

Figure 2-1: Thermal neutron die-away curves for various water contents, presented as percentage of water by atomic mass. The gray region indicates the integration bounds used in the numerator of Equation 2.

Using our die-away curves, we seek to define a single quantifiable metric that encompasses the full die-away behavior. Having one single value is useful for feeding this metric into an interpolation method, such as Gaussian process regression. Our water quantity metric has multiple necessary requirements: the metric must exhibit a monotonic relationship with water content percentage, so we can uniquely infer a single water content percentage to one metric value. Additionally, the water content should vary more strongly with lower water content values compared to high water content values, as we are more interested in differentiating lower water contents (i.e. the difference between no water and slightly more water) than large water contents.

With these factors in mind, we define a water content metric, m_W, which correlates to the integral of a normalized die-away curve. We normalize each die-away curve such that it integrates to one over the studied time frame. This follows the equations below,

$$n\left(t\right) = \frac{m\left(t\right)}{\int_0^\infty m\left(t\right) dt} = \frac{m\left(t\right)}{N}, \tag{1}$$

where N is the normalization constant, $m\left(t\right)$ is the unnormalized die-away curve (similar to those shown in Figure 2-1), and $n\left(t\right)$ is the normalized die-away curve. We define the neutron metric, m_W, as the ratio of the integral of one section of the die-away curve to the entire die-away curve, which can be represented in several forms:

$$m_W = \frac{\int_{t_0}^\infty m\left(t\right) dt}{\int_0^\infty m\left(t\right) dt} = \frac{\int_{t_0}^\infty m\left(t\right) dt}{N} = \int_{t_0}^\infty n\left(t\right) dt. \tag{2}$$

For ease, we often define

$$m_W = \frac{A}{B}. \tag{3}$$

The quantity t_0 refers to the lower integration bound for m_W. In Figure 2-1, we assume a t_0 of 2 ms, where the full shaded (gray) region corresponds to the numerator of the ratio of integrals in Equation 2. This value of t_0 was chosen to minimize uncertainty in water content estimates at lower values, while also allowing for optimal discrimination between runs with varying boron concentrations (see Section 3).

2.1 Uncertainty Quantification

We compute two sources of uncertainty on m_W: statistical uncertainty in the MCNP simulations (called σ_s) and Poisson counting uncertainty in recovering die-away curves, assuming a realistic detector (called σ_d). These two sources of uncertainty are combined with a quadrature sum:

$$\sigma_m = \left[\sigma_s^2 + \sigma_d^2\right]^{1/2}, \tag{4}$$

where σ_m is the total uncertainty on m_W.

2.1.1 MCNP Statistical Uncertainty

We leverage batch statistical analysis to compute σ_s. Since our die-away curve simulations are resolved in evenly spaced time bins of width Δt, we can rewrite the integrals in Equation 2 in the following form:

$$\int_a^b m\left(t\right) dt = \Delta t \sum_i^N m_i, \tag{5}$$

where m_i refers to the flux in time bin i. The variable N refers to the total number of time bins within the bounds of integration. The batch mean is then computed as

$$\overline{m} = \frac{1}{N}\sum_{i}^{N} m_i, \tag{6}$$

and the batch variance is described by the following equation:

$$\left(\sigma_{\overline{m}}\right)^2 = \frac{1}{N-1}\left[\left(\frac{1}{N}\sum_{i}^{N} m_i^2\right) - \overline{m}^2\right]. \tag{7}$$

The batch variance can be related back to A and B which compose the numerator and denominator of m_W (Equation 2), by altering the bins in each summation.

Batch standard deviation is converted to actual sample standard deviation with the following two equations:

$$\sigma_{s,A} = \Delta t \cdot \sigma_{\overline{m},A} \cdot N_A, \tag{8}$$

$$\sigma_{s,B} = \Delta t \cdot \sigma_{\overline{m},B} \cdot N_B, \tag{9}$$

where N_A and N_B are the number of time bins in A and B, respectively, and $\sigma_{\overline{m},A}$ and $\sigma_{\overline{m},B}$ are the corresponding batch standard deviations. We then compute the total MCNP statistical standard deviation as follows:

$$\sigma_s = \left[\left(\frac{\partial m_W}{\partial A}\right)^2 \left(\sigma_{s,A}\right)^2 + \left(\frac{\partial m_W}{\partial B}\right)^2 \left(\sigma_{s,B}\right)^2\right]^{1/2}, \tag{10}$$

which can be rewritten as

$$\sigma_s = \frac{A}{B}\left[\left(\frac{\sigma_{s,A}}{A}\right)^2 + \left(\frac{\sigma_{s,B}}{B}\right)^2\right]^{1/2}. \tag{11}$$

2.1.2 Poisson Uncertainty Assuming a Realistic Detector

We can compute Poisson counting uncertainties, by converting neutron die-away curve fluxes to realistic detector counts over time. We assume a pulsed neutron generator which produces 10^{10} neutrons at 14.1 MeV per second. Additionally, we assume a dwell time of 1800 s for a detector with 100% efficiency and an effective area of 1 cm^2. These are idealized efficiencies and effective areas, but are broadly consistent with typical detector values. These parameters are combined as follows, to convert die-away curve fluxes to counts:

$$d\left(t\right) = m\left(t\right)\left(10^{10}\ \text{neutrons/s}\right)\left(1800\ \text{s}\right)\left(1.0\ \text{efficiency}\right)\left(1.0\ \text{cm}^2\right), \tag{12}$$

where $m\left(t\right)$ is flux in units of # / cm^2 / source particle and $d\left(t\right)$ is the number of detected neutrons. We can then rewrite Equation 2 in terms of detector counts, as follows:

$$m_W = \frac{\int_{t_0}^{\infty} d\left(t\right)dt}{\int_{0}^{\infty} d\left(t\right)dt} = \frac{A_d}{B_d}. \tag{13}$$

We then assume both A_d and B_d have Poisson counting uncertainties of $\sigma_{d,A} = \sqrt{A_d}$ and $\sigma_{d,B} = \sqrt{B_d}$, respectively. These uncertainties can be combined as follows:

$$\sigma_d = \frac{A_d}{B_d}\left[\left(\frac{\sigma_{d,A}}{A_d}\right)^2 + \left(\frac{\sigma_{d,B}}{B_d}\right)^2\right]^{1/2}. \tag{14}$$

In general, larger values of $d(t)$, resulting from more neutrons, longer detection times, or increased effective area, yield a reduction in Poisson counting uncertainty.

2.2 Example Water Content Recovery

In Figure 2-2, we present the m_W values with their corresponding σ_m for each water content percentage. We can then use the relationship shown in Figure 2-2 to estimate water contents for new measurements of m_W. First, we interpolate three curves over the scatterpoints with cubic spline interpolation. Three curves are interpolated based on the measurements (black), 1σ lower limit on measurement (blue), and 1σ upper limit on measurement (red). We then compare these interpolated curves to a new m_W observation, shown as the magenta horizontal line with associated 1σ uncertainty. The orange shaded region represents the estimated water content for this m_W observation. The median inferred water content corresponds to the intersection of the black interpolated curve and median observed m_W. The lower limit on estimated water content corresponds to the intersection of the *upper* limit on the m_W measurement with the interpolated blue curve, for the lower fit, while the upper limit on estimated water content corresponds to the intersection of the *lower* limit on the m_W measurement with the interpolated red curve. For the example measurement shown in Figure 2-2, water content was estimated as $39.1^{+1.86}_{-1.77}$%, recovering the true value within 2σ.

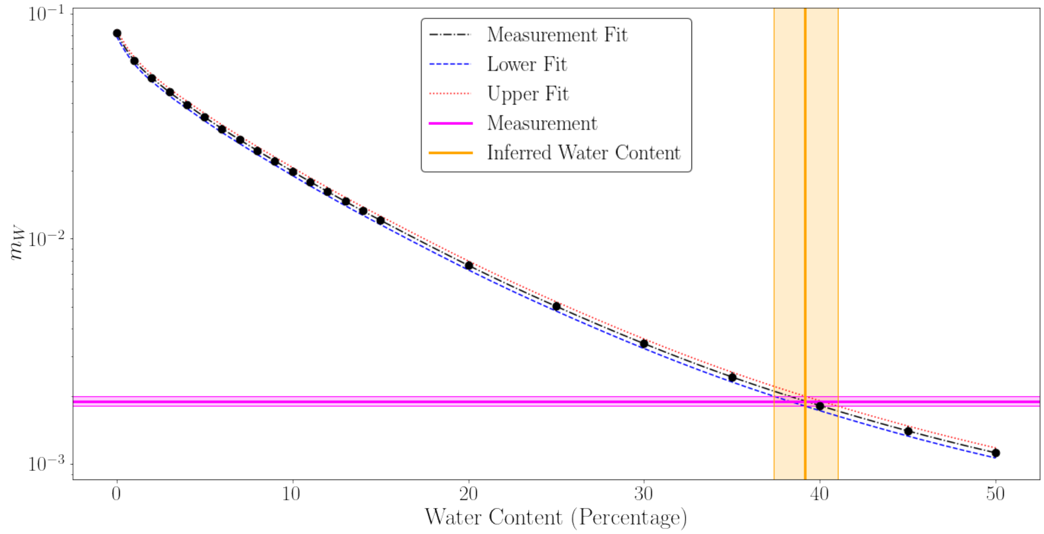

Figure 2-2: Neutron die-away curve metric, m_W, over water content. The horizontal magenta region represents m_W with 1σ uncertainty bands for a die-away curve with 37% water content. The vertical orange region presents the inferred water content of $39.1^{+1.86}_{-1.77}$%. The true value is recovered within 2σ.

We repeat this water content estimation procedure for various true water content percentages in Figure 2-3. This figure shows the estimated water content percentage, with uncertainty, given a true simulated water content. We bound resultant water content estimates between 0 and 50%. In all cases the median estimated water content results in the true value, represented by the black line.

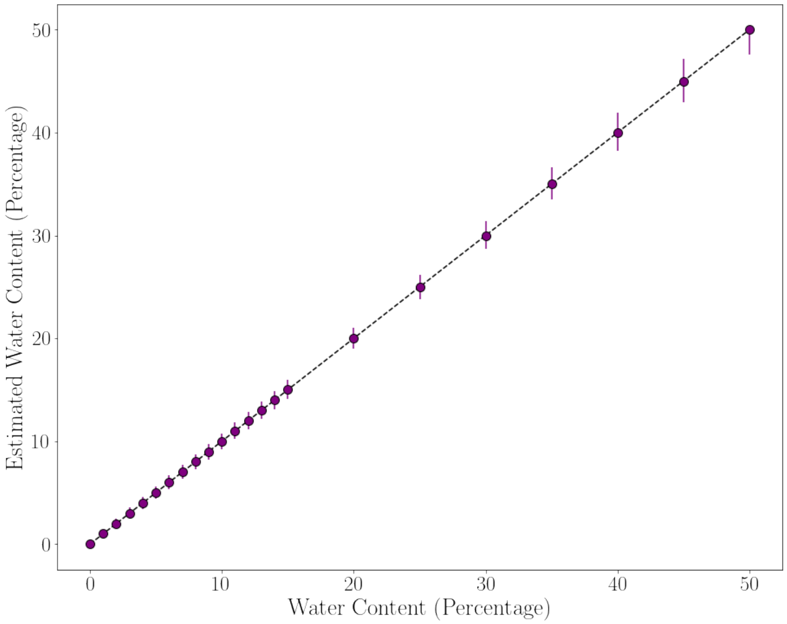

Figure 2-3: Estimated water content for each simulated value. Error bars represent 1σ uncertainty.

3 Accounting for Variations in Boron Content

The previous section presents an idealized case of water content estimation, where soil composition remained fixed relative to the proportion of water content in the soil. However, in actuality, the composition will be unknown and will deviate from the fiducial lunar sample. Therefore, in this section, we explore methods of unambiguously identifying water content, regardless of the composition of the remaining soil. We assume that mass spectroscopy is not feasible, consistent with instruments available on the Dragonfly spacecraft.

To explore this problem, we primarily look at variations in boron content in the soil, in addition to water content variation. Boron has a particularly large microscopic cross section for neutron absorption, with the number of captured thermal neutrons increasing as boron concentration increases. This is further exacerbated by the water concentration: increased water concentration results in more scatter, leading to more thermal neutrons available for capture by boron nuclei. Figure 3-1 presents the large boron-10 cross section as a function of the incident neutron energy, showing the substantially large boron cross section at thermal energies. We also present the cross section for chlorine-35, which also incurs measurable variation in neutron die-away curves with varying water content.[3] Boron produces the most significant variation in neutron interactions out of all elements realistically expected to appear in non-negligible quantities on planetary surfaces. Therefore, our analysis primarily focuses on boron variation. Additionally, exploring variation with only one element increases computational efficiency by reducing the number of MCNP simulations needed to fully explore the parameter space.

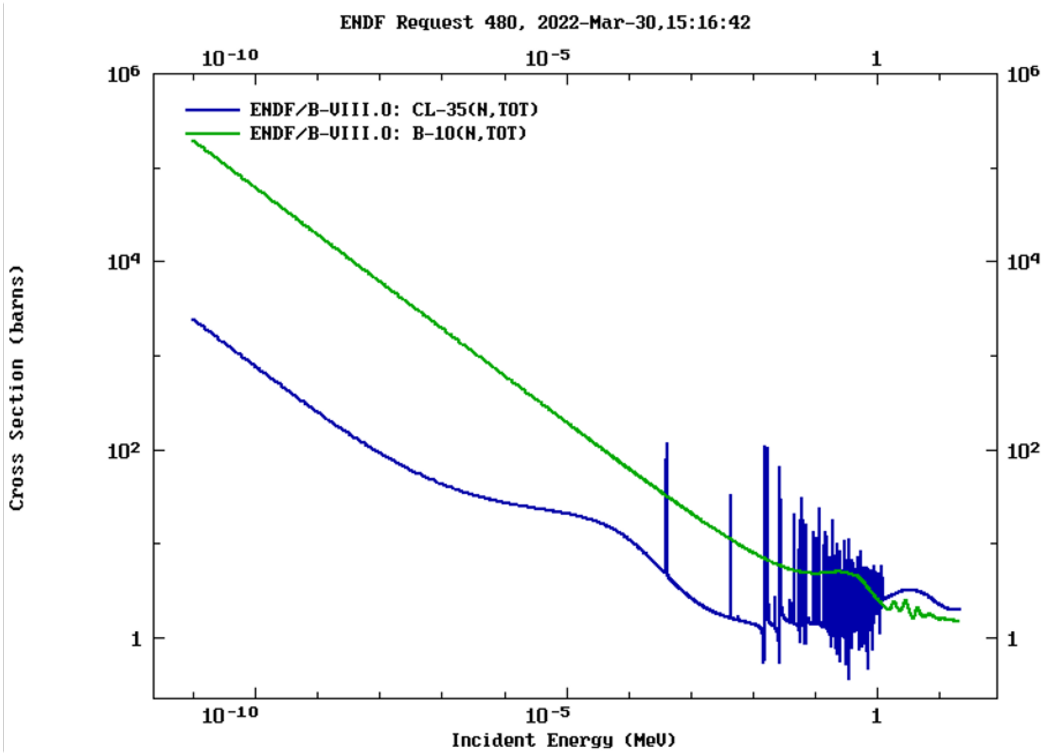

Figure 3-1: Total neutron interaction cross sections on boron-10 and chlorine-35. Both isotopes have large cross sections, with boron-10 exhibiting a substantially larger cross section at all incident thermal energies.

We assume that boron contents vary between 0 and 10,000 ppm, encompassing the maximum measured boron content for both lunar and Martian measurements [2, 3]. Figure 3-2 presents neutron die-away curves for varying boron content. In each panel, we fix the water content, vary the boron content, then set the remaining mass distribution to match the lunar soil composition. Die-away curves are dictated by a complex interplay between scattering and capture reactions from water, in addition to boron capture reactions.

Figure 3-3 displays m_W for every die-away curve in our simulation set, including variations in boron. The metric, m_W, is computed according to Equation 2, as described in Section 2, and uncertainty is smaller than the scatterpoints for all but the smallest m_W. For identical water contents, m_W significantly decreases with increasing boron content. As a result, the method described in Section 2 is no longer able to infer water content.

Instead, we use this previous method to infer an *upper limit* on water content, rather than estimating the water content itself. We assume a m_W measurement with uncertainty and compare

---

[3] We performed a similar analysis for variations in chlorine content, which is not presented in this document. Chlorine produces less variation in die-away curve behavior, and thus m_W, than boron, as supported by Figure 3-1. We only present boron variation in this document, in order to focus on the most drastic example.

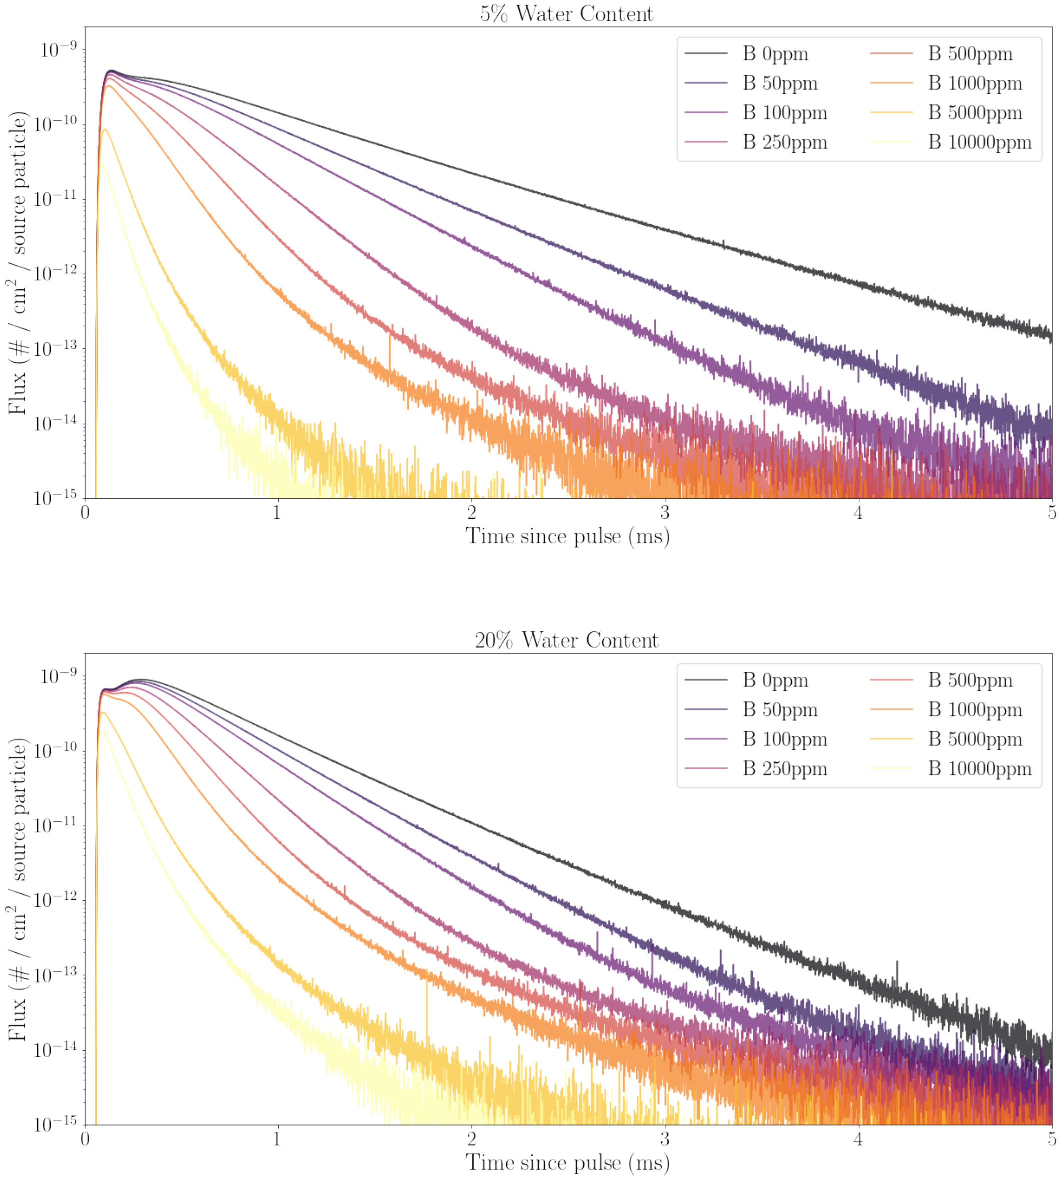

Figure 3-2: Thermal neutron die-away curves for varying boron content. The top (bottom) panel presents 5% (20%) water content.

that uncertainty to interpolated curves over the datapoints with 0 ppm boron. Then, we determine the inferred upper limit on water content, by recording the intersections of the measured m_W value with interpolated curves, taking uncertainty on m_W into account. Figure 3-3 presents an identical measurement to that presented in Figure 2-2; however, we now interpret the orange region as an upper limit on water content, meaning that the water content is estimated to be below $39.1^{+1.86}_{-1.77}\%$. Again, we restrict all inferred upper limits to values between 0% and 50%. Any m_W less than 0.0011 (50% water content with 0 ppm boron) will not intersect the interpolated curves, and results in an inconclusive upper limit estimate of $50.0^{+0.0}_{-0.0}\%$.[4] To interpret these fits as upper limits, we must assume that our baseline lunar soil composition results in the maximum m_W for any composition expected to occur on a planetary body of interest, such as Titan.

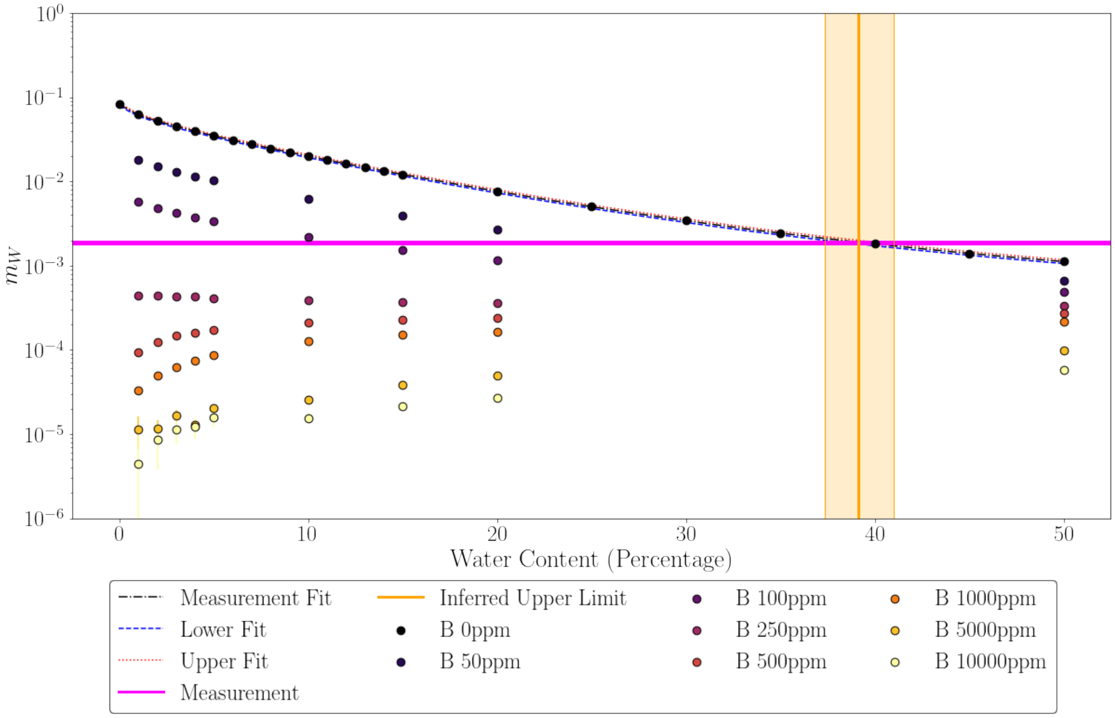

Figure 3-3: Variations in m_W for different boron concentrations. For fixed water content, m_W decreases with increasing boron content. Three curves (black, blue, red) are interpolated for scatterpoints with no boron. Values of m_W from measured die-away curves (magenta horizontal lines) are compared to the interpolated curves to produce estimates on the upper limit of water content (orange vertical lines).

We present a scatterplot of all upper limits inferred from every MCNP simulation in our data set in Figure 3-4. In all cases, the upper limit does not exclude the true water content value. There are numerous overlapping upper limit estimates for estimates of 50% water content, consistent with the large number of boron simulations which results in inconclusive water content constraints.

---

[4] We do not report uncertainty for inconclusive upper limit estimates.

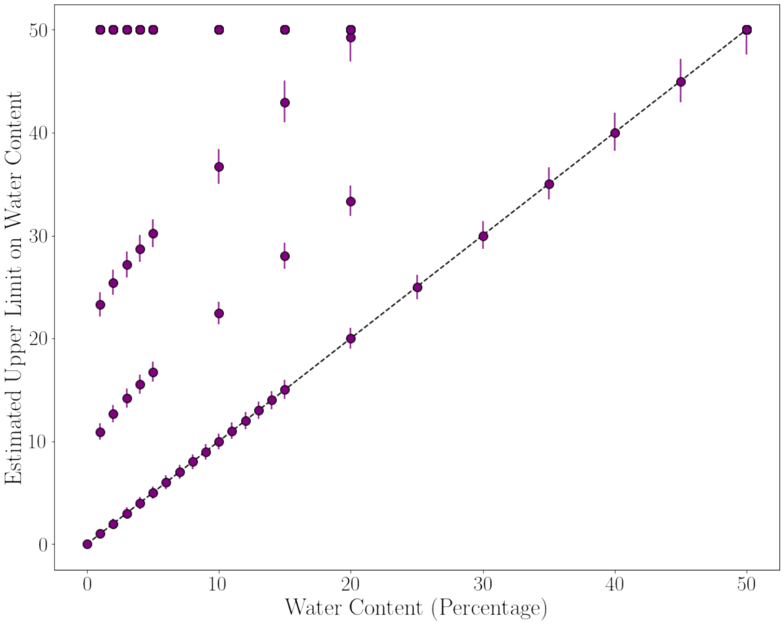

Figure 3-4: Estimated upper limit on water content for each simulation, including variation in boron content. Estimated values are bound between 0 and 50%. Inconclusive fits result in scatterpoints at 50%, plotted without error bars.

4 Refining Water Content Estimates with Gamma Spectroscopy

Water content estimates are improved by combining estimates from gamma spectroscopy with the neutron die-away curve analysis presented in the previous section. Gamma spectra are computed with MCNP simulations, using an identical configuration to the neutron die-away curve simulations, with an additional tally to capture gamma spectra produced from neutron interactions. We note that our simulated spectra are produced at far higher energy resolutions than would be anticipated on near-term planetary science instruments. Figure 4-1 presents two normalized gamma spectra, for 5% water content (top) and 20% water content (botton) for a variety of boron concentrations. Normalized spectra refer to gamma spectra where the spectral fluxes integrate to one over the full simulated energy range.

Figure 4-1 highlights gamma-ray detections between 0.4 MeV and 0.5 MeV, which encompassed the boron-10 energy peak at 0.477 MeV indicated by the gray dot-dashed lines. Higher boron concentrations result in a larger spectral flux measurement around 0.477 MeV, with the flux falling as boron concentration decreases. Normalization encompasses the effect of hydrogen spectral peaks at other simulated energies not included in the figure. Higher water contents lead to an increase in the quantity of available thermal neutrons due to scattering, thus resulting in more opportunity for boron capture reactions.

To capture the dependency on boron concentration, we define the boron content metric, m_B, as the height of the normalized spectral bin between 0.477 MeV and 0.478 MeV. Similarly to the procedure in Section 2.1.2, we must convert fluxes to detector counts to compute Poisson counting

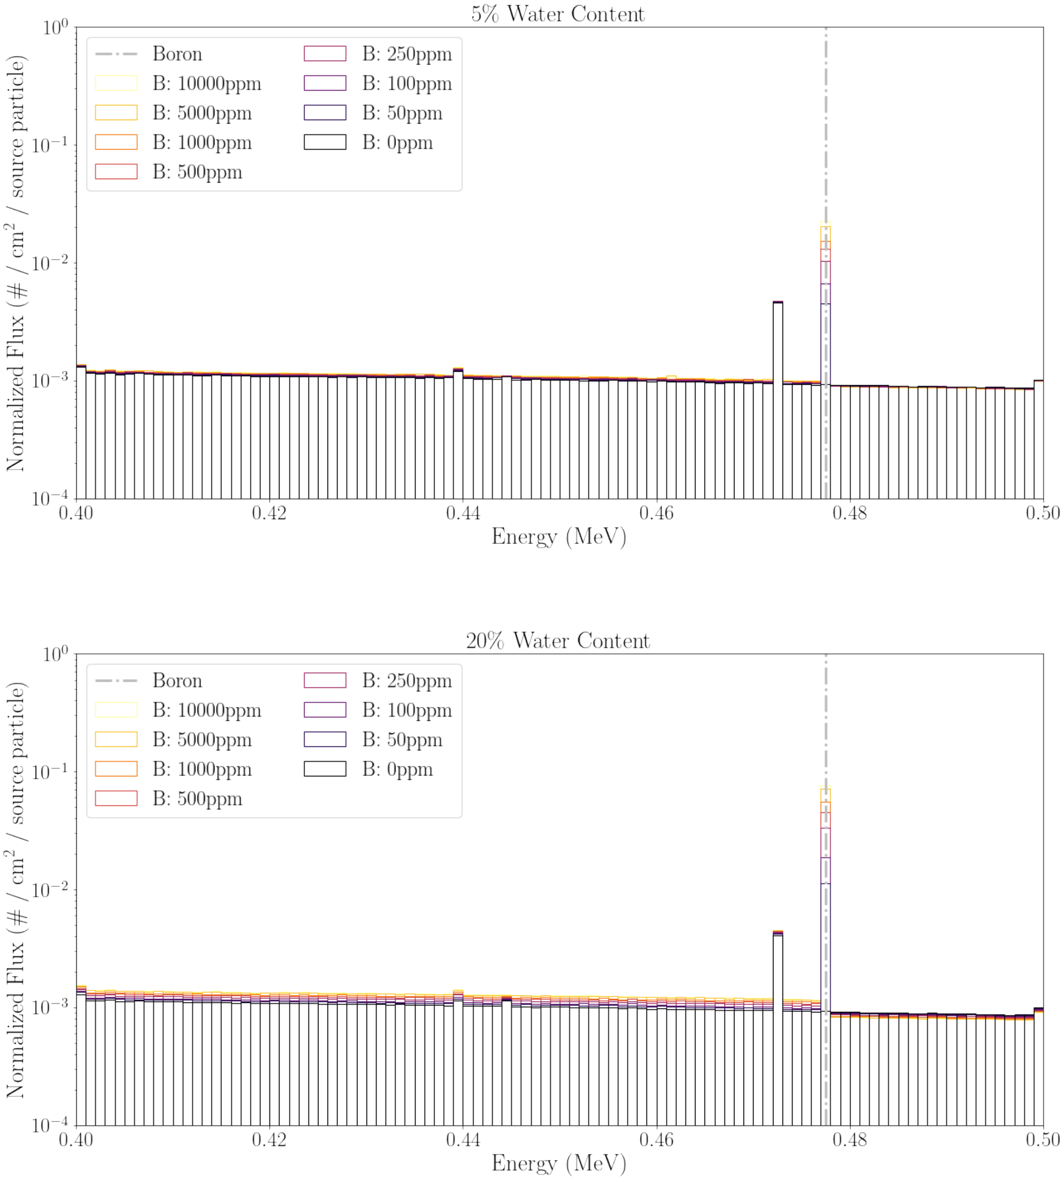

Figure 4-1: Normalized gamma spectra for 5% water content (top) and 20% water content (bottom), at a variety of boron concentrations. The vertical dot-dashed line indicates the boron peak at 0.477 MeV.

uncertainties. Batch statistical analysis is not feasible for gamma spectra, due to irregular energy binning, so we consider Poisson uncertainty as the main source of uncertainty. We assume an identical detector setup to Section 2.1.2 and convert flux to counts with the following equation:

$$d\left(E\right) = m\left(E\right)\Big(10^{10}\text{ neutrons/s}\Big)\Big(1800\text{ s}\Big)\Big(1.0\text{ efficiency}\Big)\Big(1.0\text{ cm}^2\Big), \tag{15}$$

where $m\left(E\right)$ is flux in units of # / cm^2 / source particle and $d\left(E\right)$ is the number of detected gamma-rays. We assume all gamma-rays are produced through neutron interactions.

We define m_B with the following equation:

$$m_B = \frac{d_k}{M}, \tag{16}$$

where d_k is the unnormalized gamma count in the energy bin between 0.477 MeV and 0.478 MeV. The variable M represents the normalization factor, computed as follows:

$$M = \sum_i d_i \Delta E_i, \tag{17}$$

where ΔE_i is each energy bin width and d_i is the gamma count in energy bin i. We assign Poisson counting uncertainties, such that $\sigma_M = \sqrt{M}$ and $\sigma_{d_k} = \sqrt{d_k}$. The final uncertainty on m_B is then represented by

$$\sigma_B = \frac{d_k}{M}\left[\left(\frac{\sigma_{d_k}}{d_k}\right)^2 + \left(\frac{\sigma_M}{M}\right)^2\right]^{1/2}. \tag{18}$$

Figure 4-2 presents m_B estimates for every water content and boron concentration simulated. The thick solid lines represent the median measurements, with uncertainty smaller than the width of the lines. Points for the same water content estimate are connected through linear interpolation. While m_B generally increases with boron content, water content variations dominate estimates of m_B.

We can refine water content estimates by comparing new m_B measurements to the curves in Figure 4-2. The thick magenta horizontal line indicates a new m_B measurement. In this example, the measurements correspond to a simulation with 18% water content and 85ppm boron. We then record intersection locations of the measured m_B with the median values of the curve for each water content value, corresponding to the feasible boron concentration for a given water content value. Uncertainty on the range of feasible boron values is determined by recording the intersections of the $m_B \pm \sigma_B$ measurement with $\pm\sigma_B$ estimates on each water content curve. We can place a lower limit on water content if a measured m_B does not intersect with a curve of potential m_B values for a given water content value at any boron concentration. In the example of Figure 4-2, water contents $\leq 2\%$ are ruled out by the data.

For each feasible water content value, we can interpolate a distribution of m_W as a function of boron content, given the distribution of scatterpoints shown in Figures 3-3 and 4-3. Then, we can infer a range of feasible m_W values, by comparing the range of feasible boron values for each water content to the distribution of m_W vs. boron content. In Figure 4-3, red scatterpoints indicate the range of m_W values consistent with the gamma data, incorporating Poisson uncertainty in m_B estimates. We can then compare the red scatterpoints to the magenta horizontal line, corresponding to the measured m_W value. Three cubic splines are used to interpolate m_W as a function of water content for the median, upper uncertainty, and lower uncertainty red scatterpoints. The magenta

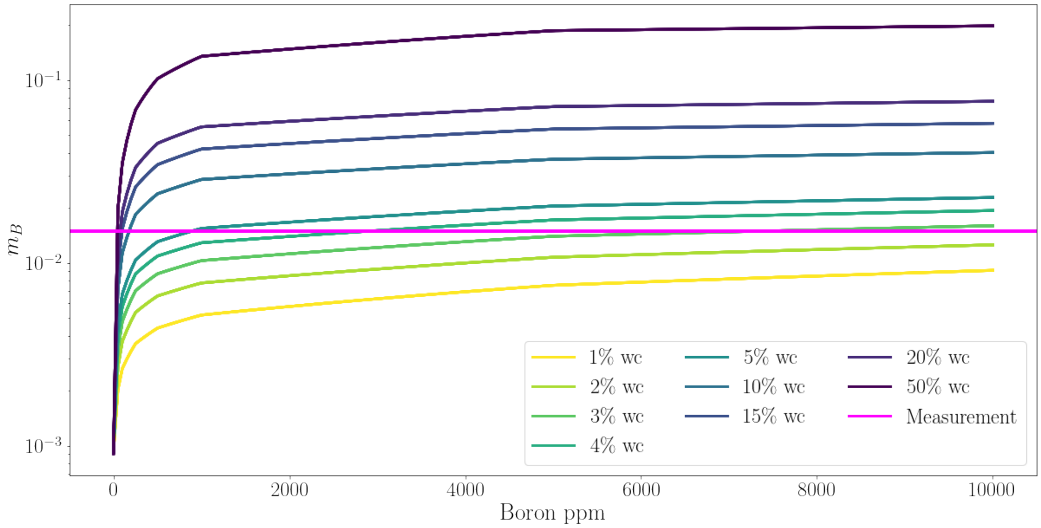

Figure 4-2: Variations in m_B with different boron concentrations, for a wide range of water contents. Uncertainty on m_B is smaller than the line width. The horizontal magenta line corresponds to m_B measured from an example gamma spectrum.

measurement is then compared to these interpolated curves, where we define any value that falls within the uncertainty bands of the interpolation as consistent with the data. We check that the estimated range of water content values is consistent with the upper limit on water content, inferred from only neutron detections.

In the example in Figure 4-3, water contents between 13.3% and 31.4% are estimated using both the gamma and neutron observations. This encompasses the true water content of 18% and is refined from the upper limit inferred from only neutron observations, shown in orange.

While this method is generally able to produce estimates on the water content, several problems still exist. Water content estimates are often weak or inconclusive, and, for some regions of parameter space, interpolation and estimation are not mathematically possible. The method only offers a partial solution to determining water content. **Gamma and neutron measurements are highly dependent and correlated, meaning that gamma measurements are insufficient to identify, let alone break, degeneracies in water content estimates from neutron measurements.**

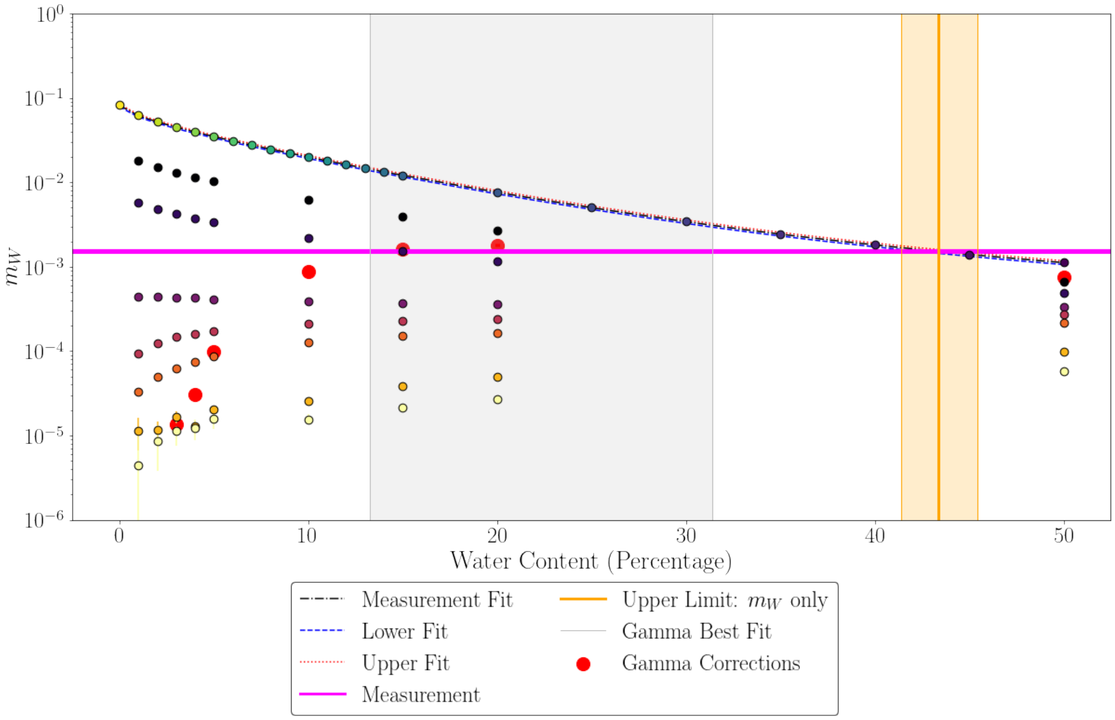

Figure 4-3: Estimate on water content (gray region) based on a combination of gamma and neutron measurements. The horizontal magenta line corresponds to the measured m_W value, while the vertical orange region is the upper limit inferred from only neutron data. Red scatterpoints indicate estimated m_W values at a given water content, based on the gamma data. The large gray region indicates the recovered water content between 13.3% and 31.4% at the 1σ level, including both gamma and neutron data. This encompasses the true water content of 18%. For almost all points, uncertainty is smaller than the scatterpoint.

5 Conclusions & Future Work

We present a statistical study coupling gamma observations with neutron die-away measurements in an attempt to capture water content. Our method successfully identifies water content for a large set of MCNP simulations, spanning multiple water compositions while fixing all other elemental contributions by-mass. However, our methods are unable to robustly identify water content when including variation in boron-10 or chlorine-35, which both exhibit substantial neutron cross sections leading to significant variation in observational signatures due to composition changes. In reality, planetary missions such as Dragonfly will attempt to infer water content for an unknown underlying soil composition, with variations in multiple elements beyond just boron or chlorine. The method presented in Section 4 is only able to weakly infer water content when boron-10 and water content are varied, and the proportion of all other isotopes remains constant. As a result, this method will not scale for soil samples where many elemental compositions are unknown.

There are multiple avenues to expand upon this work, and potentially provide stronger constraints on water content beyond the constraints presented in Section 4. Additional MCNP simulations will tighten water content estimates, especially by including simulations for variation in boron composition between 20% and 50% water content. Furthermore, more statistical techniques can be explored, including random forest or simple neural network regression techniques to estimate water content from multiple input features from neutron die-away curves and gamma spectroscopy. These machine learning methods will likely require a significantly larger sample of MCNP runs and, thus, require significant computational resources beyond the scope of this work.

[1] Jason W. Barnes et al. “Science Goals and Objectives for the Dragonfly Titan Rotorcraft Relocatable Lander”. In: 2.4, 130 (Aug. 2021), p. 130. DOI: `10.3847/PSJ/abfdcf`.

[2] D. Das et al. “Boron and Lithium in Calcium Sulfate Veins: Tracking Precipitation of Diagenetic Materials in Vera Rubin Ridge, Gale Crater”. In: *Journal of Geophysical Research (Planets)* 125.8, e06301 (Aug. 2020), e06301. DOI: `10.1029/2019JE006301`.

[3] Patrick J. Gasda et al. “In situ detection of boron by ChemCam on Mars”. In: 44.17 (Sept. 2017), pp. 8739–8748. DOI: `10.1002/2017GL074480`.

[4] A. S. Grieve et al. *Preliminary Investigation of Neutron Pulse Structure for Planetary Science Applications*. Report. Los Alamos National Laboratory, 2022.

[5] C. Hardgrove, J. Moersch, and D. Drake. “Effects of geochemical composition on neutron die-away measurements: Implications for Mars Science Laboratory’s Dynamic Albedo of Neutrons experiment”. In: *Nuclear Instruments and Methods in Physics Research A* 659.1 (Dec. 2011), pp. 442–455. DOI: `10.1016/j.nima.2011.08.058`.

[6] M. L. Litvak et al. “The Dynamic Albedo of Neutrons (DAN) Experiment for NASA’s 2009 Mars Science Laboratory”. In: *Astrobiology* 8.3 (June 2008), pp. 605–612. DOI: `10.1089/ast.2007.0157`.

[7] Tony H. Shin. *The Search for Water on Planetary Bodies using Neutron Science*. Report LA-UR-22-21910. Los Alamos National Laboratory, 2022.

[8] Tony H. Shin, Daniel T. Wakeford, and Suzanne F. Nowicki. “Multi-Sensor Optimal Motion Planning for Radiological Contamination Surveys by Using Prediction-Difference Maps”. In: *Applied Sciences* 12.11 (2022). ISSN: 2076-3417. URL: `https://www.mdpi.com/2076-3417/12/11/5627`.

[9] C. J. Werner et al. *MCNP 6.2 release notes*. Report LA-UR-18-20808. Los Alamos National Laboratory, 2018.