

Note

Statistical analysis of Curiosity data shows no evidence for a strong seasonal cycle of Martian methane

Edward Gillen^{a,1,2}, Paul B. Rimmer^{a,b,c,*,1}, David C. Catling^d

^a Cavendish Astrophysics, Cavendish Laboratory, JJ Thomson Ave, Cambridge CB3 0HE, UK

^b Department of Earth Science, Downing St, Cambridge CB2 3EQ, UK

^c Laboratory for Molecular Biology, Francis Crick Ave, Cambridge CB2 0QH, UK

^d Department of Earth and Space Sciences & cross-campus Astrobiology Program, University of Washington, Box 351310, Seattle, WA 98195, USA

ABSTRACT

Using Gaussian process regression to analyze the Martian surface methane Tunable Laser Spectrometer (TLS) data reported by Webster et al. (2018), we find that the TLS data, taken as a whole, do not indicate seasonal variability. Enrichment protocol CH₄ data are consistent with either stochastic variation or a spread of periods without seasonal preference.

1. Introduction

For the past ~15 years, purported measurements of methane on Mars have elicited excitement (Formisano et al., 2004; Krasnopolsky et al., 2004; Mumma et al., 2009), puzzlement (Lefèvre and Forget, 2009; Lefèvre, 2019) and skepticism (Zahnle et al., 2011). Using methane's 3.3 μm absorption band, reports during the years 2004–2019 using remote spectroscopy ranged from non-detections, to purported detections at mean levels ~10–15 ppbv, to claimed enhancements of up to 50 ppbv locally on Mars. The remote techniques have included spectroscopy from ground-based telescopes (Krasnopolsky et al., 2004; Krasnopolsky and Krasnopolsky, 2012; Mumma et al., 2009; Villanueva et al., 2013) and orbital spectroscopy by the Planetary Fourier Spectrometer (PFS) on the Mars Express orbiter (Formisano et al., 2004; Geminale et al., 2011; Giuranna et al., 2019). In addition, a weak signal from the 7.7 μm band in data from the Thermal Emission Spectrometer (TES) on the Mars Global Surveyor orbiter purportedly showed seasonally variable methane at levels of 0–30 ppbv (Fonti and Marzo, 2010). However, the data were revisited and the result was shown to be consistent with a non-detection (Fonti et al., 2015).

Aforementioned reports of remote methane detection and its purported variability have caused controversy (Zahnle et al., 2011). Claims of variability are confronted by the lifetime of methane on Mars of ~300 years, which is based on firmly established chemical kinetic data (Summers et al., 2002). The gas phase photochemistry of methane is

well characterized both in the laboratory and in nature because of methane's importance to the chemistry and climate of the terrestrial atmosphere (e.g., Burkholder et al., 2015; Prather et al., 2012). Thus, claims of “methane plumes” that rapidly vanish and methane variability are extraordinary and require extraordinary evidence to be accepted (Lefèvre, 2019; Lefèvre and Forget, 2009; Zahnle et al., 2011).

Questions have been raised about the purported detections from TES and PFS because of limited instrumental sensitivity and/or measurement protocols, and the high-resolution remote detection of Mumma et al. (2009) has been attributed to over-nulled ¹³CH₄ terrestrial emission when purported martian methane lines would have been Doppler blue-shifted onto ~20 times bigger ¹³CH₄ terrestrial lines (Zahnle et al., 2011). Zahnle et al. claim that when the opposite Doppler shift occurred, putting putative Mars methane lines in a relatively uncontaminated part of the terrestrial spectrum, non-overlapping terrestrial ¹³CH₄ was modeled correctly (by construction) and nulled out. But see Villanueva et al. (2013) for an alternative view.

More recently, high-resolution and high signal-to-noise spectral data from the Trace Gas Orbiter (TGO) do not show any methane to a detection limit of 0.05 ppbv (Korablev et al., 2019). With over 100 soundings around the globe, these data are much more numerous than Curiosity Rover data (discussed below). Instruments on TGO – the NOMAD (Nadir and Occultation for MARS Discovery) and ACS (Atmospheric Chemistry Suite) – were specifically designed to look for trace levels of methane on Mars.

* Corresponding author.

E-mail address: pbr27@cam.ac.uk (P.B. Rimmer).

¹ These authors contributed equally.

² Winton Fellow.

Before TGO, it was hoped that in situ measurements by the Curiosity Rover TLS (part of the Sample Analysis at Mars (SAM) instrument package), would help resolve debates about martian methane (Webster and Mahaffy, 2011). Instead TLS measurements continue to generate debate, especially when juxtaposed with the recent TGO non-detection. The TLS instrument has terrestrial methane in its foreoptics chamber, which must be exactly subtracted out in data reduction to derive accurate methane estimates (Webster et al., 2013). Specifically, following pump-outs of the foreoptics chamber on sols 400 and 1000, the methane content in the foreoptics chamber increases, which (Webster et al., 2018, hereafter W18) explain by “small amounts of methane released from epoxy and cabling within the chamber” (W18, Supplemental Material). Background terrestrial methane affects signal-to-noise (Zahnle and Catling, 2019). Importantly, unlike TLS or ground-based telescopic measurements, the sunlight spectra of TGO instruments during solar occultations by the martian atmosphere are not looking through any terrestrial methane, and so provide sensitive, uncontaminated data (Korablev et al., 2018, 2019).

TLS methane estimates are derived from two measurement protocols: (1) “direct-ingest” where intake of martian air is directly measured, and (2) “enrichment” where CO₂ is removed from the sample of martian air to concentrate the methane and other unreactive gases to provide more sensitive methane measurements (Webster et al., 2015, 2018). The enrichment protocol produces an estimated mean background methane level of 0.4 ± 0.16 ppbv (W18). However, over time, individual methane estimates from the enrichment protocol are scattered in their magnitude. In viewing this scatter, it is claimed by W18 that there is a “strong seasonal cycle” in these methane data. Unlike the enrichment protocol results, the direct protocol results provide derived methane estimates that have occasional spikes of ~ 5 – 10 ppbv methane (W18).

Establishing whether there is really a seasonal cycle is important because of the conflict between TGO global non-detection to 0.05 ppbv (Korablev et al., 2019), a recent report of a 15.5 ± 2.5 ppbv pulse of methane in 2013 from the Mars Express PFS (Giuranna et al., 2019), and results from Curiosity Rover. The purported seasonal cycle of methane on Mars has led to speculation that chemical sinks of methane are occurring, which are unique to Mars and generally near-surface (reviewed by Lefèvre, 2019), and that this rationalizes the conflict between Curiosity and TGO data. However, it is difficult to see how putative strong, near-surface destruction can be reconciled with the Mars Express or earlier telescopic detections. An attempt to model a seasonal methane cycle uses a destruction timescale for methane that is a free parameter, resulting in methane destruction ~ 1000 times faster than known chemistry, along with thermodynamic methane absorption significantly different from lab-based measurements (Moore et al., 2019).

Surface chemistry on dust or in soil (Atreya et al., 2007) or physical adsorption in soil have generally been invoked to explain methane variations. The former are hypothetical and unique to Mars, while laboratory measurements of the latter (Gough et al., 2010), are insufficient for the required sink (Meslin et al., 2011; Hu et al., 2016). Other hypotheses to explain the Curiosity detection of methane include aquifers that release sporadic methane spikes with no seasonality, or seasonal variation due to the presence of life (Hu et al., 2016). Thus, determining whether methane varies seasonally or not matters for ascertaining the validity of some hypotheses.

The motivation for this study is that the seasonal cycle that has been claimed for the Curiosity methane estimates derived from the enrichment protocol has not been shown to be statistically valid. Here, we examine the time series of the TLS data with statistical methods to see if the data are best explained by stochastic behavior or periodic variability, and, if the latter, whether seasonal cyclicality is favored.

W18 stacked all the TLS enriched methane data from different Mars years into one Mars year (their Fig. 1B). However, attempting to demonstrate periodicity by choosing a preferred period *a priori* and stacking the data into that period is not statistically robust. The reason

is that the data might be better explained by non-periodic stochastic variation or may vary with some other period.

Fortunately, robust statistical techniques to test for periodic signals in sparse data exist and are widely used. For example, such statistical methods are routinely adopted to search for periodic signals in scattered astrophysical data from stars, either to determine the periods of stellar rotation or to infer the presence of orbiting planets. Here, we use Gaussian process (GP) regression, which is a statistical method that has been used in many data-driven fields, e.g., biology, chemistry, physics, finance and data science, amongst others (e.g., Rasmussen and Williams, 2006).

We apply GP regression to the time series of Curiosity Rover Mars methane estimates to assess whether they show a statistically robust seasonal cycle. For our analysis, we use the values and error bars reported by W18: specifically, globally-inferred values from the enriched protocol that they used to claim seasonality (in their Table 1) and globally-inferred values from the direct protocol (in their Table S2).

2. Methods: statistical package

We wish to evaluate whether the variations in the surface methane level from Curiosity are statistically consistent with periodic or stochastic behavior. The basic dataset consists of a time series of methane abundance estimates, which are limited in number, e.g., 10 and 20 data points for the enriched and direct datasets, respectively. An objective, conservative approach begins with no *a priori* reason to assume that the methane estimates vary periodically or non-periodically, or stochastically or non-stochastically, i.e., we make no assumptions about which periods or timescales the data are allowed to vary on (beyond our priors; see below). This approach allows us to interrogate the data in an unbiased, statistically correct way where the analysis and results are driven by the dataset itself, as opposed to weighting a specific period, e.g. annual or diurnal.

Gaussian processes (GPs) offer a flexible framework for performing Bayesian inference on functions. A GP is a non-parametric form of model that defines a distribution over functions. GPs can be thought of as a way to model a dataset by parameterising the covariance between pairs of data points, as opposed to explicitly defining a functional form of model to fit the data. As GPs define a distribution over functions, when we vary the parameters of a GP (called hyperparameters), we move through function space rather than the traditional parameter space of parameteric models. GPs have been applied to similar problems before, e.g., modeling the atmospheric CO₂ concentration measurements made at Mauna Loa, Hawaii (e.g. Rasmussen and Williams, 2006). Tests show that an uninformed model approach, where no specific periods (e.g. annual) are weighted a priori, can find evidence of annual (i.e., seasonal) variation in the CO₂ level (as well as a long term trend of increasing CO₂ level with time). This suggests that uninformed GP models, which we use here, can infer genuine periodicity in a dataset. A detailed description of GP regression is beyond the scope of this Note, so we refer the interested reader to Roberts et al. (2012) and Rasmussen and Williams (2006) for a general and thorough introduction, respectively.

For this work, we use the Celerite GP package (Foreman-Mackey et al., 2017; Foreman-Mackey, 2018) through the exoplanet toolkit (Foreman-Mackey and Barentsen, 2019)³. Given the sparsely sampled TLS data and possibility of periodic variation in the methane level on Mars, we opted to use a GP kernel that is akin to a stochastically-driven damped simple harmonic oscillator (SHO). Depending on the level of damping, this kernel can vary smoothly and periodically, or stochastically (with rougher aperiodic variations). It is therefore a useful tool to assess both the presence and significance of variability within a dataset. Essentially, the SHO GP will explore functions that are consistent with

³ <https://github.com/dfm/exoplanet>

the data: if the data show periodic variability, and are sufficiently constraining, then the GP will favor periodic models. However, if the data show variability, but lack the required level of constraint to assert that it is periodic, then the GP will explore both periodic and aperiodic variability models. Based on the GP posterior distribution, it is possible to assess the level of evidence for or against periodic variability within a dataset, as well as providing an estimate and uncertainty for any such period(s).

The hyperparameters of this model are: ω_0 , the frequency of the undamped oscillator, Q , the quality factor of the oscillator, and S_0 , which is proportional to the power at ω_0 (see Foreman-Mackey et al., 2017). We used the natural logarithm of these parameters with the following priors: $\ln(\omega_0) \sim \mathcal{U}(\ln(2\pi/5000), \ln(2\pi/30))$, $\ln(Q) \sim \mathcal{U}(-5, 5)$ and $\ln(S_0) \sim \mathcal{N}(0, dr^2)$, where the standard deviation dr is the range of methane concentrations spanned by the data. Our choice of prior for $\ln(S_0)$ seeks to avoid overfitting for short periods and large amplitudes. We also fit for the mean of the data with prior $\mu \sim \mathcal{N}(dm, dr^2)$, where dm is the mean methane concentration, and a jitter term with $\ln(\sigma^2) \sim \mathcal{U}(-30, 10)$. We note that, with so few data-points, the analysis presented here will be sensitive to these priors, which will also apply to any model (see Section 4 for discussion of alternative priors). To perform posterior inference, we implemented gradient-based Markov-chain Monte Carlo (MCMC) with PyMC3 (Salvatier et al., 2016) using No U-Turn Sampling (NUTS; Hoffman and Gelman, 2014). We ran 5 independent chains of 100,000 steps, which typically yielded around ten thousand effective samples for model evaluation.

3. Results

We assess the claim of strong seasonal variation in Mars' background methane by applying the GP regression described above to methane estimates from Curiosity TLS data (W18). We perform three sets of models on the TLS data: 1. considering all data (i.e. direct and enriched); 2. considering the direct data alone; and 3. selecting only the enriched data, which W18 focused on to claim strong seasonal variation. It is important to note that we model the data in time to search for the presence of periodicity in the data, rather than folding the data on an already-determined period and then searching for trends in this phase space. The former allows for unbiased inference while the latter does not.

Fits to both the full and enriched datasets are shown in Fig. 1. Fig. 1a shows 200 realizations of the GP model fit to all data (direct and

enriched), which show the kinds of models that are consistent with the data. Taken as a whole, the distribution of models shows no clear periodicity in the background methane, seasonal or otherwise. Fig. 1b shows 200 realizations of the GP model fit to only the enriched data, which W18 focused on. We find that many of the individual models do favor relatively smooth variations, but the period of variability is not well constrained due to: 1. the sparseness of the data (10 data points spread over 1136 sols); and 2. the fact that the data do not cover even two full Martian years (1337 sols). This can be seen in the individual draws from the GP posterior distribution, most notably around 800 sol where the data are unable to constrain the GP distribution to a singular form of variation. It is worth noting that we incorporated a white noise jitter term into our GP model. This allows the GP to inflate the observational uncertainties, under penalty, by adding an additional error term in quadrature to the formal error estimates. As stated, any such inflation is penalized by the fit, so the GP will only do this if it is deemed necessary to explain the data, i.e. if the GP finds the formal errors to be underestimated.

As introduced in Section 2, we explore the posterior parameter space of the TLS data using a gradient-based MCMC algorithm. This allows us to efficiently sample a wide range of models, and hence periods, and from this, assess which periods are consistent with the data. The posterior distributions of these consistent periods can then be constructed to help us understand which period(s) are most favored by the data. Fig. 2 shows the probability density of models as a function of variability period and/or timescale. We make this distinction because some of the GP models are not strictly periodic and hence we consider their variations to follow a characteristic ‘timescale’ rather than a formal period. Fig. 2a shows the posterior period distribution for all models: this includes both periodic and stochastic models, and hence we explicitly label the horizontal axis as ‘variability period/timescale’. The distributions show that there are a range of variability periods and timescales consistent with all datasets (enriched, direct and both together), and that variability on the Martian year (668.6 sols) is not preferred over a wide range of other periods or timescales of variation. The top half of Table 1 reports the median periods and 1σ uncertainties for these distributions, which highlight the wide range of allowed periods. We caution the reader not to interpret these values as reliable approximations for the full distributions, which are more complex (e.g., the period distribution for the direct data is multimodal) and themselves indicative estimates only.

In Fig. 2b we select only the models that display periodic behavior, which amounts to around half of the full distributions (see Table 1,

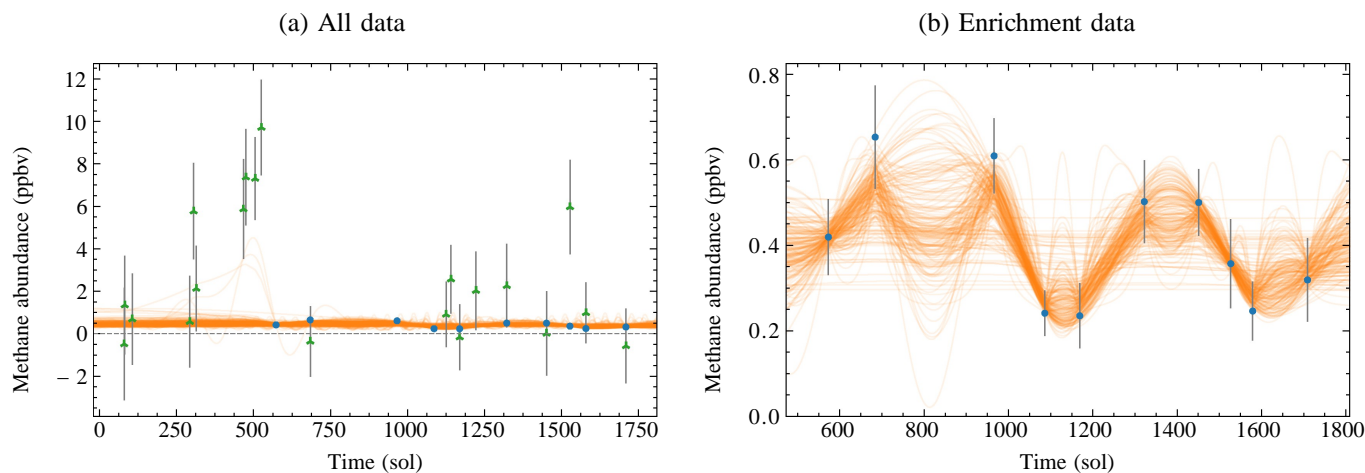


Fig. 1. Methane surface concentration [ppbv] vs. time [sol], using all data (a) and only the enrichment data (b), from Webster et al. (2018), where error bars are 1 standard error. Green triangles represent direct data and blue circles the enriched data. The methane estimates are fit using GP regression, where the orange lines represent 200 individual models drawn from the GP posterior distribution. These give some indication of the range of variability models that are consistent with the sparse TLS data. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

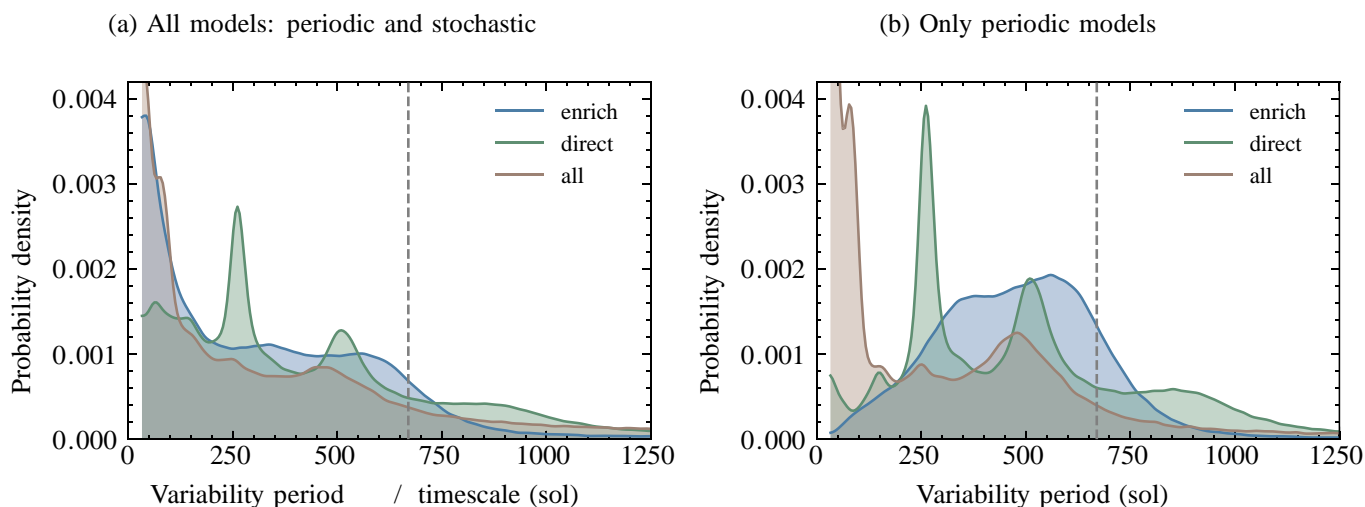


Fig. 2. Posterior period distributions from the GP model. Higher probability density corresponds to a more strongly favored period for methane variability. Three posteriors are shown: using all the data (red), only the direct data (green), and only the enriched data (blue). Panel a shows the results when including all GP models and panel b shows the posterior period distribution when selecting only those models which we classify as periodic. The vertical dashed grey line indicates the variability period expected if there were seasonal variation of methane on the Martian surface. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 1

Posterior GP variability periods/timescales from fitting different combinations of the data: enriched, direct, and both together. Results for the full posterior distribution are shown (top) along with the results from selecting only models that show periodic behavior (bottom). The % of the full posterior distribution is indicated for all models highlighting the fraction that show periodic behavior. For comparison, note that a Martian year is 668.6 sols (equivalent to 686.98 Earth solar days). We highlight that the posterior period distributions (see Fig. 2) are typically complex/multimodal, and hence the median period and 1σ uncertainty cannot accurately capture the distribution shape. We primarily quote values here to highlight the magnitude of the allowed period ranges.

Dataset	Posterior GP period (sol) (median $\pm 1\sigma$ uncertainty)	% of posterior distribution
Full posterior distributions		
Enriched & direct	363 $^{+1117}_{-290}$	100
Direct	386 $^{+522}_{-249}$	100
Enriched	321 $^{+349}_{-246}$	100
Selecting only periodic models		
Enriched & direct	341 $^{+588}_{-281}$	44
Direct	491 $^{+414}_{-241}$	55
Enriched	498 $^{+202}_{-205}$	47

bottom half). Given the possibility of stochastic variation on top of periodic modulation, as well as the sensitivity to our priors, we use a loose criterion for periodicity, which corresponds to models with $Q > 1/\sqrt{2}$. Even so, we note that selecting only the periodic models is a strong assumption, which is akin to making a 50:50 guess as to whether the observed variations are periodic in nature or not. Nonetheless, if we assume that periodic variability is in fact present, we again see that variability on the Martian year (i.e. seasonal variation) is not strongly supported by the data. Considering all data together, short periods (≤ 100 sols) are most strongly favored, but periods up to ~ 600 sols are plausible. Seasonal variation is not preferred. This is consistent with what W18 suggest for all data: the combined direct and enriched dataset are best explained by stochastic processes. The same can be said of the direct data, which seem to favor periods ~ 250 or 500 sols, although we note that this is at least in part driven by the sparseness of the data.

When selecting only the enriched data, the posterior distribution in Fig. 2b shows that a broad range of variability periods are consistent with the data. While there are a subset of models possessing variability

periods that are consistent with the Martian orbital period, this is by no means preferred by the data. Although a seasonal variation cannot be ruled out, if we consider only periodic models that lie within 5% (33 sols) of the Martian orbital period, only 9% of the periodic models are consistent with seasonal variation. Given that only about half of all models are periodic, only 5% of all models are seasonal with 5% tolerance on the period; or alternatively, 95% of models represent either stochastic variation or variation that is on a period other than seasonal. Finally, if periodic variability is present in the enriched data, we find that it is more likely to have a period somewhere between ~ 300 sols and the Martian orbital period (668.6 sols), than it is to vary on the Martian orbital period itself. We note that a Lomb-Scargle periodogram suggests a similar period and range to our periodic GP models, and also does not favour seasonal variation.

W18 observe that the sol 965 enriched methane datapoint is higher than their assumed single modal seasonal variation, and speculate that this may be due to a methane spike. In other words, W18 invoke a stochastic process to explain the 965 datapoint, and our model already accounts for stochastic processes. It is inappropriate to run our analysis with the 965 datapoint removed because this point is not a statistical outlier of the enriched dataset: it is less than 0.855 ppbv, which is the third quartile plus 1.5 times the upper interquartile range of the data - the standard statistical rule for a high outlier (Moore et al., 2009, p. 38). Nonetheless, if the sol 965 datapoint is removed, the result is that half of the models remain stochastic and the fraction of all models that have periods within 5% (~ 33 sols) of a Martian year increases from 5% to 7%.

4. Discussion

Many models can adequately explain the Curiosity methane data and their parameters cover a large range that must be better constrained by the addition of more surface data before anything definitive about the variability of background methane can be claimed. Given the small amount of data, the values and uncertainties of individual data points become important. In this regard, it is worth noting that reports of TLS CH_4 and their uncertainties have seen changes from paper to paper for certain martian sols (e.g., the enrichment factor has changed by a factor of 1.09 between Webster et al. (2015) and W18, while the sol 306 direct in situ value has changed from -2.21 ± 0.94 ppbv (Table 1, Webster et al., 2013) to 5.78 ± 2.27 ppbv due to removal of

systematic error (Table 1, Webster et al., 2015; Table S2, W18). This suggests that alternative treatments of systematic and random errors in the TLS data could cause further revision of derived CH₄ values. This motivated our choice to allow the GP to inflate the observational uncertainties, under penalty, if it deemed necessary for individual models.

As noted, this analysis is sensitive to the priors because this small amount of data is not constraining. We explored the effect of our priors by relaxing the prior on $\ln(S_0)$, as well as testing modified Jeffreys priors on S_0 (with ‘knee’ values between 0.01 and 1) and on our jitter term σ (with ‘knee’ values between 10-3 and the mean methane concentration uncertainty, above which the posterior distributions smear out). Within these ranges, we obtain qualitatively similar posterior results, with the main differences being at short $\lesssim 150$ sol periods; these also do not favour seasonal variation over other periodic or stochastic variation.

We must wait and see whether future data will reveal seasonal variation in background surface methane, or whether future data, such as future TGO observations in the nadir mode, will better constrain the concentration of surface methane on Mars (Korablev et al., 2019). Here, we simply assess whether the TLS data reported in W18 are evidence favoring the claim of strong seasonal variability over stochastic processes or variability with a different period. We find that they are not.

The sparse Curiosity enrichment data appear by eye to exhibit seasonal variability when the data are phase-folded within a one year period, which effectively weights the data in favor of an annual period, but in-depth statistical analysis is necessary for determining with any confidence whether the data themselves actually favor annual periodicity. Results of statistical analyses show no evidence favoring strong seasonal variation of methane over stochastic variability or variation over a wide range of other periods. The enriched data are roughly equally consistent with either stochastic processes or periodic variability where the latter includes an annual cycle but that cycle is not favored by the data over the many other periods. Indeed, a simple Bayesian model comparison between a flat line and a sine model does not favour either, and the sine model's period range does not favour seasonal variation.

More time series data would be desirable because the current enriched dataset comprises 10 measurements spread over 1136 sol, i.e. 1.7 Martian years. To assess whether a periodic signal is present in a dataset with a given period, a dataset covering a timespan of three or more times that specific period is often considered minimal. Furthermore, to reject other potential shorter periods or stochastic variation, a sufficiently high temporal density of measurements is desirable. The exact timespan required and its density of data depend on the complexity of the underlying processes, e.g., whether there is a single well-defined process (e.g., seasonal variation) or whether there are other stochastic factors, and the relative strengths of the various factors.

In conclusion, the statistical analysis in this paper finds that the hypothesis of “strong seasonal variability” in Mars' surface methane is unsupported by the Curiosity TLS data. This is because the data are too sparse over too limited a timespan to favor a seasonally cyclic explanation of the data over alternative hypotheses of stochastic variation or variation with other periods.

Acknowledgments

We thank Suzanne Aigrain and Kevin Zahnle for helpful discussions and comments, respectively, Vinesh Rajpaul for assistance with Bayesian model comparison, and two anonymous reviewers. This paper was initiated while DCC was a Leverhulme Trust Visiting Professor at the University of Cambridge. PBR and DCC thank the Simons Foundation for funding (SCOL awards 599634 to PBR and 511570 to DCC). EG gratefully acknowledges support from the David and Claudia Harding Foundation in the form of a Winton Exoplanet Fellowship.

References

- Atreya, S.K., Mahaffy, P.R., Wong, A.-S., 2007. Methane and related trace species on Mars: origin, loss, implications for life, and habitability. *Planet. Space Sci.* 55, 358.
- Burkholder, J., Sander, S., Abbatt, J., et al., 2015. Chemical kinetics and photochemical data for use in atmospheric studies: evaluation number 18. In: *Tech. Rep. Jet Propulsion Laboratory, National Aeronautics and Space Administration, Pasadena, CA.*
- Fonti, S., Mancarella, F., Liuzzi, G., et al., 2015. Revisiting the identification of methane on Mars using TES data. *A&A* 581, A136.
- Fonti, S., Marzo, G.A., 2010. Mapping the methane on Mars. *A&A* 512, A51.
- Foreman-Mackey, D., 2018. Scalable backpropagation for Gaussian processes using celestial. *Res. Notes Am. Astron. Soc.* 2 (1), 31.
- Foreman-Mackey, D., Agol, E., Ambikasaran, S., Angus, R., 2017. Fast and scalable Gaussian process modeling with applications to astronomical time series. *AJ* 154, 220 Dec.
- Foreman-Mackey, D., Barentsen, G., 2019. Software Repository: Exoplanet v0.1.3. <https://doi.org/10.5281/zenodo.2536576>. Jan.
- Formisano, V., Atreya, S., Encrenaz, T., Ignatiev, N., Giuranna, M., 2004. Detection of methane in the atmosphere of Mars. *Science* 306, 1758.
- Geminale, A., Formisano, V., Sindoni, G., 2011. Mapping methane in Martian atmosphere with PFS-MEX data. *Planet. Space Sci.* 59, 137.
- Giuranna, M., Viscardi, S., Daerden, F., et al., 2019. Independent confirmation of a methane spike on Mars and a source region east of gale crater. *Nat. Geosci.* 1.
- Gough, R.V., Tolbert, M.A., McKay, C.P., Toon, O.B., 2010. Methane adsorption on a Martian soil analog: an abiogenic explanation for methane variability in the martian atmosphere. *Icarus* 207, 165.
- Hoffman, M.D., Gelman, A., 2014. The no-u-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *J. Mach. Learn. Res.* 15 (1), 1593.
- Hu, R., Bloom, A.A., Gao, P., Miller, C.E., Yung, Y.L., 2016. Hypotheses for near-surface exchange of methane on Mars. *Astrobiology* 16, 539.
- Korablev, O., Montmessin, F., Trokhimovskiy, A., et al., 2018. The atmospheric chemistry suite (ACS) of three spectrometers for the exomars 2016 trace gas orbiter. *Space Sci. Rev.* 214, 7.
- Korablev, O., Vandaele, A.C., Montmessin, F., et al., 2019. No detection of methane on Mars from early exomars trace gas orbiter observations. *Nature* 568, 517.
- Krasnopolsky, V.A., Krasnopolsky, V.A., 2012. Search for methane and upper limits to ethane and SO₂ on Mars. *Icarus* 217, 144.
- Krasnopolsky, V.A., Maillard, J.P., Owen, T.C., 2004. Detection of methane in the Martian atmosphere: evidence for life? *Icarus* 172, 537.
- Lefèvre, F., 2019. The enigma of methane on Mars. In: Cavalazzi, B., Westall, F. (Eds.), *Biosignatures for Astrobiology*. Springer, pp. 253.
- Lefèvre, F., Forget, F., 2009. Observed variations of methane on Mars unexplained by known atmospheric chemistry and physics. *Nature* 460, 720.
- Meslin, P.Y., Gough, R., Lefèvre, F., Forget, F., 2011. Little variability of methane on Mars induced by adsorption in the regolith. *Planet. Space Sci.* 59, 247.
- Moore, D.S., McCabe, G.P., Craig, B.A., 2009. *Introduction to the Practice of Statistics*, 6th. W. H. Freeman, New York.
- Moore, J.E., Gough, R.V., Martinez, G.M., et al., 2019. Methane seasonal cycle at gale crater on Mars consistent with regolith adsorption and diffusion. *Nat. Geosci.* 12, 321.
- Mumma, M.J., Villanueva, G.L., Novak, R.E., et al., 2009. Strong release of methane on Mars in Northern Summer 2003. *Science* 323, 1041.
- Prather, M.J., Holmes, C.D., Hsu, J., 2012. Reactive greenhouse gas scenarios: systematic exploration of uncertainties and the role of atmospheric chemistry. *GRL* 39, L09803.
- Rasmussen, C.E., Williams, C.K.I., 2006. *Gaussian processes for machine learning*. Gaussian Processes for Machine Learning MIT Press.
- Roberts, S., Osborne, M., Ebdon, M., et al., 2012. Gaussian processes for time-series modelling. *Phil. Trans. R. Soc. A* 371 (1984), 20110550.
- Salvatier, J., Wiecki, T.V., Fonnesbeck, C., 2016. Probabilistic programming in Python using pymc3. *PeerJ Comput. Sci.* 2, e55.
- Summers, M.E., Lieb, B.J., Chapman, E., Yung, Y.L., 2002. Atmospheric biomarkers of subsurface life on Mars. *GRL* 29, 2171.
- Villanueva, G.L., Mumma, M.J., Novak, R.E., et al., 2013. A sensitive search for organics (CH₄, CH₃OH, H₂CO, C₂H₆, C₂H₂, C₂H₄), hydroperoxy (HO₂), nitrogen compounds (N₂O, NH₃, HCN) and chlorine species (HCl, CH₃Cl) on Mars using ground-based high-resolution infrared spectroscopy. *Icarus* 223, 11.
- Webster, C.R., Mahaffy, P.R., 2011. Determining the local abundance of Martian methane and its ¹³C/¹²C and D/H isotopic ratios for comparison with related gas and soil analysis on the 2011 Mars Science Laboratory (MSL) mission. *Planet. Space Sci.* 59, 271.
- Webster, C.R., Mahaffy, P.R., Atreya, S.K., et al., 2013. Low upper limit to methane abundance on Mars. *Science* 342, 355.
- Webster, C.R., Mahaffy, P.R., Atreya, S.K., et al., 2015. Mars methane detection and variability at Gale crater. *Science* 347, 415.
- Webster, C.R., Mahaffy, P.R., Atreya, S.K., et al., 2018. Background levels of methane in Mars atmosphere show strong seasonal variations. *Science* 360, 1093.
- Zahnle, K., Catling, D., 2019. The paradox of Mars methane. *Ninth International Conference on Mars 2019 (LPI Contrib. No. 2089)*, 6132. <https://www.hou.usra.edu/meetings/ninthmars2019/pdf/6132.pdf>.
- Zahnle, K., Freedman, R.S., Catling, D.C., 2011. Is there methane on Mars? *Icarus* 212, 493.