MNRAS 000, 1–4 (2023) Preprint 6 September 2023 Compiled using MNRAS LATEX style file v3.0

The Rotation of Magnetic White Dwarfs

Rachel A. Patton,[1,2] Ilaria Caiazzo,[3] Cole Johnston[4,5] Evan Bauer[6] and Ylva Götberg[7]
[1]*Department of Astronomy, The Ohio State University, 140 West 18th Ave, Columbus, OH 43210, USA*
[2]*Center for Cosmology and AstroParticle Physics, The Ohio State University, 191 West Woodruff Avenue, Columbus, OH 43210, USA*
[3]*TAPIR, Walter Burke Institute for Theoretical Physics, Mail Code 350-17, Caltech, Pasadena, CA 91125, USA*
[4]*Department of Astrophysics, IMAPP, Radboud University Nijmegen, P. O. Box 9010, 6500 GL Nijmegen, the Netherlands*
[5]*Institute of Astronomy, KU Leuven, Celestijnenlaan 200D, B-3001 Leuven, Belgium*
[6]*Center for Astrophysics | Harvard & Smithsonian, 60 Garden St., Cambridge, MA 02138, USA*
[7]*The observatories of the Carnegie institution for science, 813 Santa Barbara Street, Pasadena, CA 91101, USA*

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

The origin of magnetism in white dwarfs is presently unknown. Knowing the underlying distribution of properties of magnetic white dwarfs can constrain the cause of their strong magnetic fields; different sources correspond to different properties. While mass, field strength, temperature, age, and composition have been explored, the rotation periods of magnetic white dwarfs have not been measured en masse. In this project we aim to extract rotation periods for 842 magnetic white dwarfs using long-baseline photometry from the Zwicky Transient Facility. In a test sample of 80 white dwarfs, we find that nearly 80% of them have rotation periods below an hour, with the fastest ones rotating every few minutes.

Key words: white dwarfs – stars:magnetic field – stars:rotation

1 INTRODUCTION

About 10-20% of white dwarfs (WDs) exhibit magnetic fields (Ferrario et al. 2020, and references therein). Measured field strengths extend as high as 10^9 G and as low as 10^3 G, though this range could change with a detection method that was sensitive to lower field strengths. The origins of these fields are presently unknown. There are several proposed mechanisms by which a WD could develop a field, including retaining a fossil field from the WD's progenitor or a dynamo triggered by crystallization, common envelope evolution, or a WD-WD merger (Kawka 2018, 2020; Ferrario et al. 2020, and references therein).

These different mechanisms should produce magnetic white dwarfs (MWDs) with different properties. For example, MWDs which were merger products should be more massive on average than non-MWDs and be rapidly rotating. MWDs which experienced a common envelope event should have a close binary companion. MWDs which retained the magnetic field of their progenitors should be slowly rotating due to the enhancement of core-envelope coupling by the magnetic field close to the end of the star's life.

It follows that the way to distinguish between these proposed formation mechanisms is to aggregate the observed properties of MWDs and compare to what is predicted by the different theories. This is challenging in its own right due to a number of detection biases. For MWDs, the presence and strength of a magnetic field is detected spectroscopically, either through Zeeman splitting (e.g., Kemic 1974) of lines or spectropolarimetry (e.g., Landstreet & Bagnulo 2019). Zeeman splitting of lines becomes resolvable at field strengths above 10^6 G, whereas the latter is sensitive down to about 10^3 G but only works on bright sources. Furthermore, there are disparities between the number of MWDs with very strong fields (B > 10^6 G) and strong fields (10^3 < B < 10^6 G) because spectroscopic and spectropolarimetric surveys are not carried out in equal proportion. Spectroscopic surveys are far more common, thus the field strengths of detected MWDs are biased to larger values.

There are also photometric biases. WDs decrease in size as they increase in mass, which decreases their luminosity and makes them more difficult to detect. Older WDs are also more difficult to detect than younger WDs of equivalent mass because of drops in luminosity caused by cooling.

Nevertheless, some initial trends are emerging from the more well-characterized MWDs, namely that many of the most massive WDs known are magnetic (Bagnulo & Landstreet 2022), less-massive MWDs tend to be older (Bagnulo & Landstreet 2022), and the MWD mass-distribution peaks at a value higher than the mass-distribution of non-MWDs (Ferrario et al. 2020). More massive MWDs also tend to have stronger fields.

However, rotation, despite being another way to distinguish between magnetic origins, has only been measured for a few MWDs (Brinkworth et al. 2013; Ferrario et al. 2020). Yet it remains possible, as magnetism should produce variations on the surface of WDs which would modulate their light curves. In this project we seek to measure the rotation periods for as many MWDs as possible. Section 2 outlines our sample selection. In section 3 we discuss how to extract and verify the rotation periods and section 4 presents the preliminary period distribution for a subset of MWDs. Finally, we discuss the next steps of this project in section 5 and conclude in section 6.

2 SAMPLE SELECTION AND PERIODOGRAM CONSTRUCTION

We take light curves from Data Release 17 of the Zwicky Transient Facility (ZTF), a time-domain survey run at Palomar Observatory (Masci et al. 2019). ZTF is magnitude limited and observes in both Mould r and Sloan g bands, with a limiting magnitude of about 21 in both bands (Smith et al. 2014). DR17 includes 5 years of data. There are more than 1000 known MWDs (Dufour et al. 2017; Amorim et al. 2023, and references therein), 883 of which are visible to ZTF. Of the 883, 842 have ZTF light curves in at least one of the two bands. In addition to the r and g light curves, we also combine to the two light curves by normalizing to variations around 0. When searching for periodic modulation, it is the fluctuations that matter, not the magnitude values.

We also select a test sample of 80 WDs to verify the utility of several techniques to assess the robustness of our period measurements. These WDs were initially selected due to their high mass and short rotation periods but spectroscopic follow-up revealed that they all had magnetic fields. 77 of the 80 MWDs have forced-photometry measurements as well, which yields less noisy data and more accurate rotation period measurements.

ZTF tiles the entire northern sky (declination > -30°) roughly every two days but observes any given patch randomly in those two days, meaning that there are not uniform spacings between observations. This is an advantage because the recoverable periods are not Nyquist-limited. Instead, the limiting period is one minute, which is twice the exposure time of ZTF. As we expect at least some of the magnetic white dwarfs to be very rapidly rotating, sensitivity to short periods is critical.

The Lomb-Scargle periodogram (Lomb 1976; Scargle 1982; VanderPlas 2018) is able to approximate the Fourier transform of a set of data with non-uniform sampling. For our initial tests, we use the Lomb-Scargle periodogram implemented in Astropy due to its fast computation time. To make the periodogram, we must supply both the time-series data, including measurement errors, and a frequency grid to check.

Selecting the proper grid is critical to extracting the proper period. We choose a minimum frequency of the inverse of one half of the time baseline and a maximum frequency of 720 cycles per day. We want to sample at smaller frequency intervals than the minimum frequency, so an oversampling factor of 5-10 is typically adopted. We choose a frequency spacing of one fifth of the minimum frequency. The end result is a few million frequencies checked per light curve. Since the computational time for these periodograms scales as n log(n), where n is the number of data points, running many periodograms can quickly get computationally expensive.

3 PERIOD VERIFICATION

Ideally, if a WD is rotating and there are non-uniformities present on the surface of the WD to impart a signal in the light curve, its periodogram should be a delta function with a peak at a frequency of the inverse of the period. But there are several reasons why this is not the case. First, there is a baseline level of noise in the periodogram because of the noise and measurement errors in the photometry. Second, there are other sources of periodicity. The window function imparts its own signal in the light curve. The fact that observations are made daily imparts a one-day signal in the light curve. The changing sky brightness from lunar phases imparts a near-monthly signal. The Earth's orbit imparts an annual signal. There could be

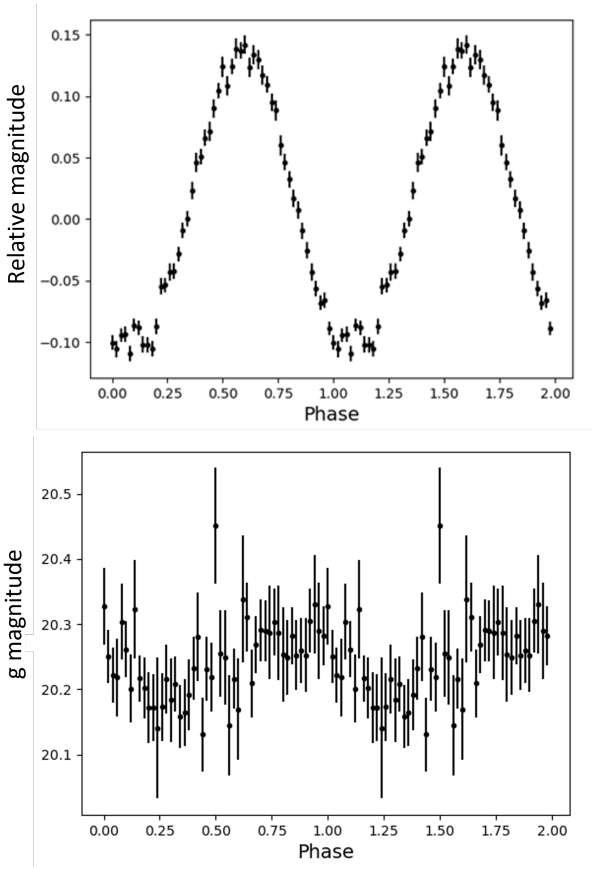

Figure 1. Phase-folded light curves for two MWDs in our sample, displaying a light curve with very obvious periodic variation (top) and a light curve where it is unclear whether there is a periodic signal (bottom). The lack of obvious variation in the bottom panel implies that either the data are too noisy or we extracted the incorrect period.

eclipses if the WD is in a binary. Third, the periodic signal may be too weak compared to the noise, so that the periodogram just appears to be noise with no strong signal. Lastly, the periodogram could have recovered an integer multiple of the true period and not the period itself.

Nevertheless, periodograms will always have a highest peak. We just have to assess the validity of that peak. There are several things we try. First, we phase-fold the light curves against the best-fit period. If the light curve collapses into a coherent shape, then we have extracted the correct period. If there is no obvious periodic modulation in the phase folded-light curve, then we have likely extracted the incorrect period. Some examples are shown in Figure 1.

The second assessment is the signal to noise, or peak significance. Because definitions of noise vary, we calculate noise in two ways: taking the root mean square (rms) of the entire periodogram and taking a rolling average of the noise, in this case taking steps in frequency of 5 cycles per day and averaging over a bin of 10 cycles per day. For our targets the significance values are comparable between the two, so we opt to just use the rms for the full sample. In our subset of 80 MWDs, we calculate a broad range of significance values, 5 to 185, with a peak around 15. We select 15 as our lower limit for a peak to be significant.

Figure 2 demonstrates why peak significance is a useful metric. A periodogram will always have a frequency of maximum power, whether it is a true periodic signal or just the strongest frequency

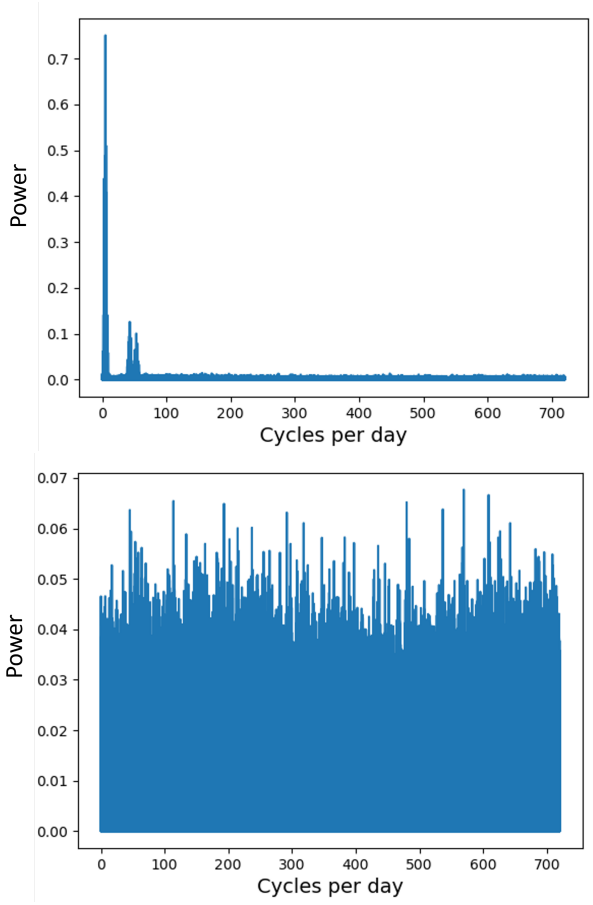

Figure 2. Caption

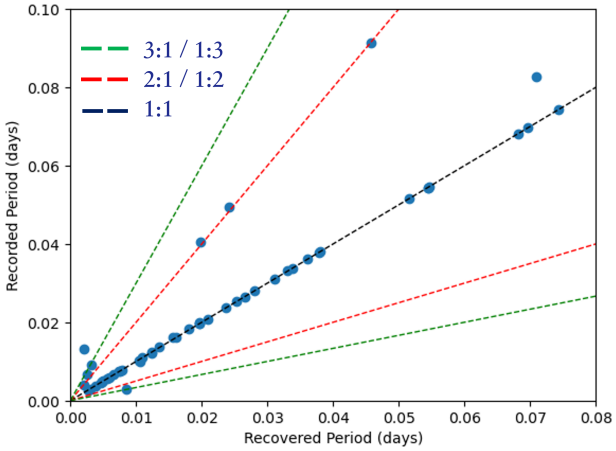

Figure 3. Reported versus recovered periods (in days) for a subset of 80 MWDs. Lines represent different integer multiples of the periods reported. The majority of periods which do not lie on the 1:1 line are aliases.

in the noise. The top panel of Figure 2 gives an example of what a true signal looks like. Not only is there a sharp peak, the power in the signal is at least an order of magnitude higher than what the average noise is. In the bottom panel, the highest peak does not stand out from many other frequencies. This could be a real signal buried in noisy data, or it is just that: noise. Peak significance gives us an initial sense of how trustworthy the signal is.

It is possible that the period with the strongest signal in the periodogram is not the true period but rather an integer multiple of the true period, an alias. Figure 3 shows the periods measured in our subset of 80, first months ago (the reported period) and now with a longer baseline of ZTF data (recovered period). Most values lie on the 1:1 line, yet of the periods which do not, the majority lie along an aliased line. For the full sample, we do not know the periods ahead of time, so we cannot construct a plot like Figure 3. However, we can phase-fold the light curves around a few aliased periods and see if the scatter reduces. Should the scatter drop, we will adopt the alias as the true period.

The strongest metric we have for assessing the validity of the period we recover is the false alarm probability, or the likelihood that the period we find is not the true period. Traditionally this is measured using a semi-analytic fit. However, this method breaks down at high frequencies, returning non-sensical values like 10^{-100} when a more appropriate value would be 0.01. Since we expect many of the MWDs to be rapidly rotating, we need sensitivity at high frequencies.

Instead, we will use a manual proxy for the false alarm probability in the form of Monte Carlo iterations. We will scramble the data 10,000 times, each time extracting a period and recording its significance. The false alarm probability is then the number of times we find a peak in the scrambled data with comparable or greater significance to the peak in the unscrambled data divided by 10,000. The higher this value, the lower the confidence in the period extracted from the original light curve.

Calculating 10,000 periodograms for 842 MWDs with millions of frequencies sampled per target results in a large computational expense. To do this efficiently, we need GPUs, whose parallel processing dramatically speeds up computation times. We use cuvarbase, a python library capable of running on GPUs which calculates, among other things, Lomb-Scargle periodograms. This enables us to run all periodograms in a few days.

4 PRELIMINARY ROTATION PERIODS

We are presently working on verifying rotation periods for the full sample of MWDs. However, we do have rotation periods extracted from combined light curves for the subset of 80 we were using as our test case. Figure 4 shows the distribution of their rotation periods.

There are several things to notice about this distribution. The first is that the periods go from minutes to days, with nearly 80% of targets having periods below a day. This is significant because the peak of the non-MWD period distribution is 1-2 days (Hermes et al. 2017) and the shortest period in that distribution was around an hour. For these MWDs, half of the sample would not even fall on the plot in Hermes et al. (2017). There is a significant bias in our distribution to short periods. However, this subset was selected to contain rapidly rotating MWDs, and it is possible that the final distribution of periods does not skew as much to such short periods.

The second thing to notice is that there are two peaks, one at one day and the other at 29.5 days. These are spurious signals, one from periodicity imparted by observing only at night almost every night and the other from the lunar phases; 29.5 days is the synodic period of the moon. The MWDs whose light curves showed the strongest modulation at these periods may still show variation at their respective rotation periods. We would have to remove the spurious period and then look again.

With these two false peaks removed, there is a more normal underlying distribution with a peak around 40 minutes and an extended tail up to periods of days. About 25% of the sample has a period be-

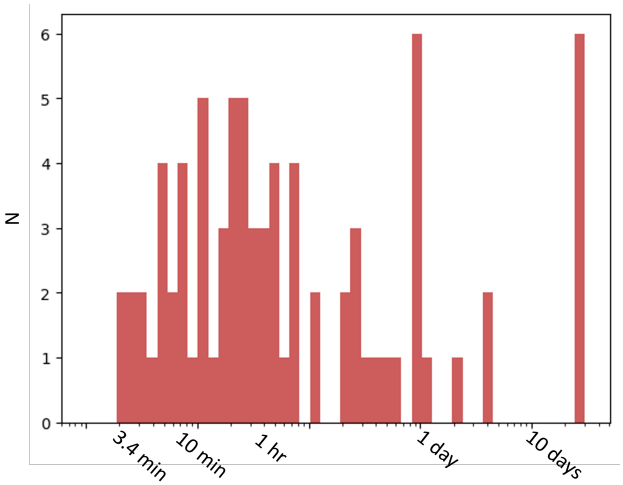

Figure 4. The distribution of rotation periods of 80 MWDs. Note that the x-axis is logarithmic and the labels have been replaced with minutes, hours, and days to make the time axis more interpretable.

tween 3.4 and 10 minutes. These very rapid rotators may come from WD-WD mergers, though it requires looking at trends with mass to determine that. Another third of the sample has periods between 10 minutes and an hour and another quarter have periods between an hour and a day. These large range in periods suggests that there are different sources of the MWDs rotation, which implies different origins of their magnetism. We think the distribution of periods for the full sample will provide good constraints on the origins of MWDs.

5 FUTURE WORK

We have verified that our period assessment techniques are adequate and have the initial, unverified periods for our entire sample. The remaining work includes checking for aliases and calculating the false alarm probability for the entire sample. Once we are confident in our period distribution, then we will look for trends against other known properties of our entire MWD sample, including mass and field strength. It is these trends which may be able to further constrain the origins of magnetism in WDs.

6 CONCLUSIONS

Different origins of magnetism in WDs should produce MWDs with different distributions of rotation periods. As such, measuring the rotation periods for as many MWDs as possible can place broad constraints on the origins of MWDs. Doing this requires careful verification of the periods extracted, including an assessment of peak significance, the presence of aliases as well as calculating a false alarm probability, and the removal of spurious periods, such as those at 1 and 29.5 days.

With the above assessments, we find a striking distribution of rotation periods in an initial subset of 80 MWDs. Nearly 80% of the sample rotates with periods shorter than a day, the typical rotation period of a non-MWD (Hermes et al. 2017). The presence of larger population of very rapidly rotating MWDs hints that WD-WD mergers could be a significant channel for forming MWDs. While these initial MWDs skew to shorter periods, we cannot assume that the distribution shape and the proportions of MWDs in each period range will remain the same in the full sample. We also anticipate many MWDs for which we cannot extract a rotation period due to noisy data and weak light curve modulation. Nevertheless, rotation is a key parameter in understanding the origin of magnetism in WDs.

ACKNOWLEDGEMENTS

This work was conducted as part of the 2023 Kavli Summer Program in Astrophysics, hosted by the Max Planck Institut für Astrophysik (MPA). Support and funding for this work comes from the Kavli Foundation and MPA. RAP would also like to thank Kevin Burdge for assistance with running periodograms on the GPU cluster.

DATA AVAILABILITY

Light curves from ZTF DR17 are publicly available. MWDs were selected from the Montreal White Dwarf Database (Dufour et al. 2017) and Amorim et al. (2023), which is also publicly available. Our results will remain proprietary until publication.

REFERENCES

Amorim, L. L., Kepler, S. O., Külebi, B., et al. 2023, ApJ, 944, 56. doi:10.3847/1538-4357/acaf6e
Bagnulo, S. & Landstreet, J. D. 2022, ApJ, 935, L12. doi:10.3847/2041-8213/ac84d3
Brinkworth, C. S., Burleigh, M. R., Lawrie, K., et al. 2013, ApJ, 773, 47. doi:10.1088/0004-637X/773/1/47
Dufour, P., Blouin, S., Coutu, S., et al. 2017, 20th European White Dwarf Workshop, 509, 3. doi:10.48550/arXiv.1610.00986
Ferrario, L., Wickramasinghe, D., & Kawka, A. 2020, Advances in Space Research, 66, 1025. doi:10.1016/j.asr.2019.11.012
Hermes, J. J., Gänsicke, B. T., Kawaler, S. D., et al. 2017, ApJS, 232, 23. doi:10.3847/1538-4365/aa8bb5
Kawka, A. 2018, Contributions of the Astronomical Observatory Skalnate Pleso, 48, 228. doi:10.48550/arXiv.1801.05602
Kawka, A. 2020, White Dwarfs as Probes of Fundamental Physics: Tracers of Planetary, Stellar and Galactic Evolution, 357, 60. doi:10.1017/S1743921320000745
Kemic, S. B. 1974, ApJ, 193, 213. doi:10.1086/153149
Landstreet, J. D. & Bagnulo, S. 2019, A&A, 623, A46. doi:10.1051/0004-6361/201834638
Lomb, N. R. 1976, Ap&SS, 39, 447. doi:10.1007/BF00648343
Masci, F. J., Laher, R. R., Rusholme, B., et al. 2019, PASP, 131, 018003. doi:10.1088/1538-3873/aae8ac
Scargle, J. D. 1982, ApJ, 263, 835. doi:10.1086/160554
Smith, R. M., Dekany, R. G., Bebek, C., et al. 2014, Proc. SPIE, 9147, 914779. doi:10.1117/12.2070014
VanderPlas, J. T. 2018, ApJS, 236, 16. doi:10.3847/1538-4365/aab766

This paper has been typeset from a TEX/LATEX file prepared by the author.