Evaluating Modern Tools and Large Language Models for Digitising Sea-Level Data from Marigrams

Christian Kenwright[1*], Ivan D. Haigh[2,3], Joanne Williams[4], Robert J. Nicholls[5], Steve McFarland[6], Marc Becker[6], Nurul Tazaroh[1], Seyedeh Fardis Pourreza Ahmadi[1], Lily Sharp[1], Callum Slade[1,7], Molly A. Phillips[5], Sam E. Pearson-Smith[8], Minnie J. Darby[1], Patrick Sharpe[8] and Duo Chan[1]

1. School of Ocean and Earth Science, University of Southampton, UK

2. Department of Civil, Environmental and Construction Engineering, University of Central Florida, Florida, USA

3. National Centre for Integrated Coastal Research, University of Central Florida, Florida, USA

4. National Oceanography Centre, Liverpool, UK

5. School of Engineering, University of Southampton, UK

6. Scottish Environment Protection Agency, Glasgow, UK

7. Natural History Museum, Cromwell Road, UK

8. Independent Researcher

*Corresponding Author (ck1g20@soton.ac.uk)

ORCID ID: 0009-0004-4425-2135

This manuscript has been submitted to EarthArXiv as a non-peer reviewed preprint. This manuscript has been submitted to the Geoscience Data Journal for peer review and publication

Abstract

Numerous graphical data digitisation tools exist; however, few have been applied to the recovery of historical sea-level marigrams (tidal charts). Using an early twentieth-century marigram from the Clyde Estuary, Scotland, containing one week of sea level observations, we evaluated the performance of 11 digitisation tools and 2 large language models (LLMs) in digitising marigrams. A multi-user testing framework was used to assess the accuracy and speed of each approach. Both LLMs tested, ChatGPT5.5 and Gemini 3.1Pro, were unsuitable for this application under the tested conditions, because they show errors exceeding 10% of the observed tidal range. Conventional digitisation tools performed substantially better, most digitisation tools successfully recovered sea-level data with errors below 1%. Despite this accuracy, these tools remained dependent on substantial human input. Even the fastest approach, DigitGraph, required on average 29 minutes to recover one week of sea-level data. Scaled to a century-long record, this would require more than one working year to complete, making application to large datasets challenging. Consequently, further development of dedicated, automated digitisation algorithms remains necessary.

Significance Statement

Historical sea-level records offer valuable insights into past sea levels, but much of the data remains on paper and is difficult to analyse. Here, we evaluate the capability of 11 digitisation tools and two large language models (LLMs) to digitise historical sea-level records. We find that LLMs do not produce results accurate enough for practical use. In contrast, most other tools achieve accuracy comparable to manual digitisation while requiring substantially less time. However, all accurate approaches remain too time-consuming for long records. While some tools can therefore improve digitisation efficiency, faster methods are still needed to enable the scalable digitisation of long historical tide records.

1.Introduction

Historical sea-level records are essential for understanding long-term sea-level trends, variability and extremes (Hogarth, 2014; Latapy et al. 2023; Haigh et al. 2025). Yet many records are inaccessible for machine-readable quantitative analysis as they are preserved in their original paper format or as scanned images (Talke and Jay, 2013; Latapy et al. 2023). For example, the recent TideVault project has catalogued 872 sub-daily sea-level records, representing ~15,264 years of sea-level observations, but currently 72.4% of these records remain undigitised (Kenwright et al., in review).

Among these records, marigrams are of particular interest because they preserve near-continuous sea-level observations, recorded as graphical traces (Figure 1a). Unlike images of tabulated data, which have previously been digitised using optical character recognition techniques (e.g. Bhaskar et al., 2025; Askarov et al., 2025; Gupta et al., 2015; Singh et al., 2012), extracting data from marigrams requires a different approach. Data must be inferred from graphical curves, often while separating overlapping daily traces and correcting for distortions caused by curled paper or uneven scanning.

Manually digitising marigrams is time-consuming and labour-intensive (McLoughlin et al. 2024). To make the process more efficient, a variety of digitisation tools have been developed and used, these include NUNIEAU for records from Marseille, France (Wöppelmann et al., 2014), Plot Digitizer for records from Williamstown, Australia (McInnes et al., 2024), and WebPlotDigitizer for records from Dún Laoghaire, Ireland (McLoughlin et al., 2024). Other tools also exist although they have not been used for marigrams. Kenwright et al. (In review) recently identified 20 potentially suitable digitisation tools, but their performance and suitability for marigram digitisation have not yet been systematically evaluated.

Additionally, the increasing availability and rapid development of LLMs (Large Language Models) (Zhao et al. 2023) may provide an alternative semi-automated approach to marigram digitisation. Although LLMs have previously been tested for extracting data from text, tables and graphs, (Dagdelen et al. 2024; Lovón-Melgarejo et al. 2025; Das et al. 2026), their applicability to marigram digitisation has yet to be assessed.

To address these two knowledge gaps, this study assesses the accuracy and speed of the digitisation tools identified by Kenwright et al. (in review), together with two LLMs, for extracting data from marigrams.

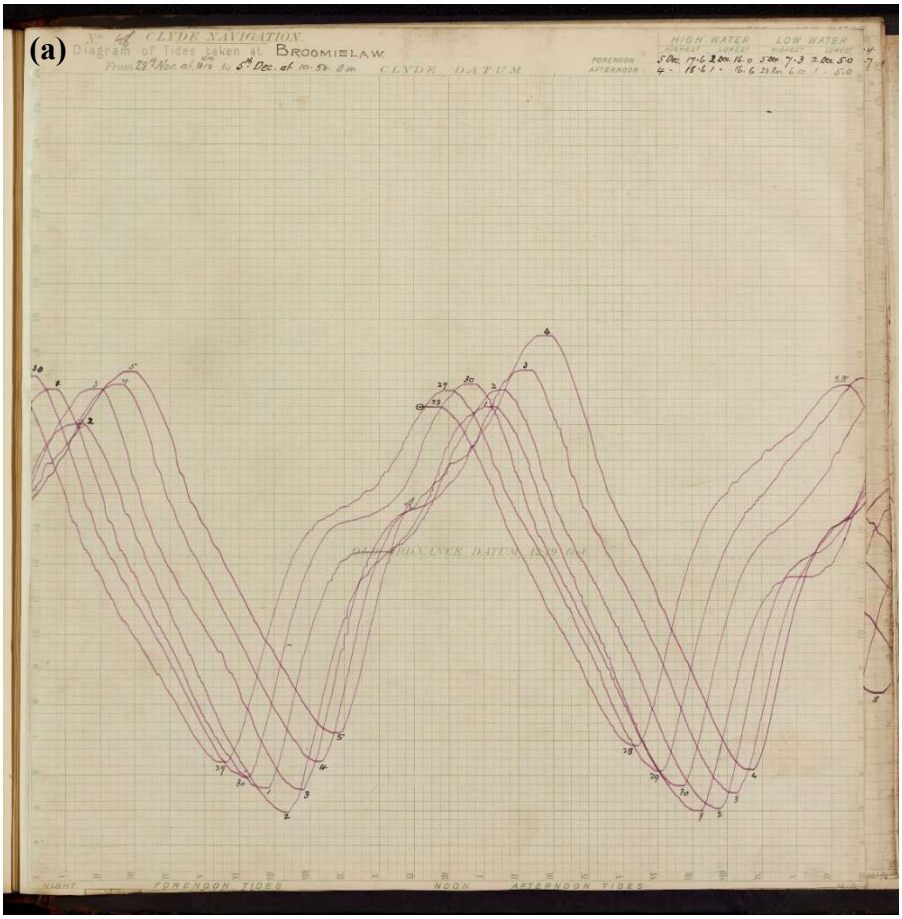

Figure 1. (a) *The test marigram used to test the digitisation tools. Water level is on the y-axis in feet, and the time of day is on the x-axis. Each red line depicts a separate day. This record shows the water level at Broomielaw, Scotland, between 28th November to 5th December 1910.* **(b)** *A map of the Clyde Estuary, Scotland, depicting the locations of Broomielaw (red) and the other nine historical tidal records (yellow).*

82

2. Materials and Methods

2.1. Study Site and Test Data

In this study we focus on a marigram from the Broomielaw tide gauge in Scotland. Broomielaw is one of ten undigitized long records in the Clyde Estuary (Figure 1b), recording near-continuous water levels from 1890 to 1980. We specifically chose a site within the estuary, rather than at the estuary mouth. The complex tidal patterns near the tidal limit of the estuary, reflecting the combined effects of River Clyde discharge and tidal amplification within the funnel-shaped estuary (Sabatino et al., 2016), provide a challenging but manageable test of the tools' digitising capabilities. A test marigram (Figure 1a) was selected from this record, displaying water-level observations, in feet, from 28th November to 5th December 1910 on a single sheet, with each overlapping trace representing one day. Native units (feet) are used for the duration of this paper.

This specific marigram was selected because the physical record exhibits a small degree of warping while the overlapping tidal traces introduce additional complexity. At the same time, the axes are clearly presented and legible and the record shows minimal deterioration. Together, these characteristics provide a challenging but fair test for evaluating the digitisation tools.

2.2. Methodology

Using a multi-user testing framework involving nine researchers (the lead author and eight co-authors), 11 of the 20 digitisation tools identified by Kenwright et al. (In review) were evaluated (Table 1). The remaining nine tools and the reasons as to why they were not tested in this study are listed in Table 2.

Based on their extraction methods, tools can be divided into two categories: trace tools and point-and-click tools. Trace tools extracted data using a colour recognition algorithm on an area manually traced by the user, whereas point-and-click tools infer values directly from user-selected points.

Each participant independently digitised the test marigram using each of the 11 selected tools and manually digitised the data from the marigram by eye. All participants received standardised instructions describing the marigram, digitisation procedure, and operation of each tool. Prior to testing, participants were required to familiarise themselves with both the

marigram and the digitisation software. In addition, testers were also instructed to take breaks of at least one hour between tests of each digitisation tool to reduce the potential effects of fatigue on performance.

Two LLMs, ChatGPT5.5 and Gemini 3.1Pro (henceforth referred to as ChatGPT and Gemini, respectively), were tested separately using a common prompt containing all task instructions and relevant contextual information. The prompt was co-developed with both LLMs before either model was used for the digitisation task (Supplementary Material 1). Each LLM was run nine times with the same prompt, as nine independent sessions, to reflect the stochastic nature of modern LLM outputs.

Data was digitised at a 15-minute temporal resolution and output as an Excel spreadsheet for both human participants and LLMS. This frequency was chosen due to the gridlines being spaced at 15-minute intervals, to record at finer resolutions would introduce a greater risk of error, creating an unreliable basis to compare digitisation methods.

The performance of each method is evaluated using two metrics, accuracy and speed of digitisation. Accuracy is quantified using root mean squared errors (RMSE), with all outputs compared against a reference dataset manually digitised by the lead author. Here we define an accurate approach as one where the mean RMSE is below 0.14ft (1% of the entire marigram range; ~0.043m).

Speed of digitisation is assessed as the time from starting the actual digitisation task to export the final Excel spreadsheet. Preparation time, including human familiarisation with the software and LLM prompt development, was not included in the time.

Some tools automatically sample the plotted trace at a higher temporal resolution, i.e., NUNIEAU, WebPlotDigitizer, Plot Digitizer, and DigitGraph (Table 1). To compare them with the 15-minute resolution reference series, their outputs were linearly interpolated to a common 15-minute interval as a post-processing step using MATLAB. Linear interpolation is not expected to qualitatively alter the accuracy metric, because interpolation passes through the original higher-resolution digitised points.

Table 1. *A list of the digitisation tools tested in this study with their category of method along with the resolution of output and the means of accuracy and speed.*

Digitisation tool	Category	Original output resolution	Mean accuracy ± standard deviation (feet)	Mean speed ± standard deviation (hours)
NUNIEAU	Trace	5-minute	0.16±0.03	1.16±0.12
WebPlotDigitizer	Trace	<1-minute	0.11±0.03	0.91±0.62
Plot Digitizer	Trace	<1-minute	0.24±0.16	0.64±0.18
DigitGraph	Trace	<1-minute	0.24±0.12	0.48±0.24
Engauge Digitizer	Point-and-click	15-minute	0.08±0.02	0.78±0.32
Dagra	Point-and-click	15-minute	0.09±0.04	0.60±0.15
DigitizeIt	Point-and-click	15-minute	0.11±0.09	0.70±0.15
GRABIT	Point-and-click	15-minute	0.09±0.02	0.76±0.13
GetData Graph Digitizer	Point-and-click	15-minute	0.12±0.13	0.71±0.29
Graph Grabber	Point-and-click	15-minute	0.09±0.04	0.66±0.23
OriginLab	Point-and-click	15-minute	0.11±0.08	0.69±0.14
ChatGPT5.5	LLM	15-minute	1.58±0.37	0.20±0.04
Gemini 3.1Pro	LLM	15-minute	5.67±1.02	0.06±0.01
Manual	Manual	15-minute	0.07±0.03	2.44±0.65

Table 2. *A list of the digitisation tools not tested in this study with their category of method along with the resolution of output and the reason for not testing them.*

Digitisation tool	Category	Original output resolution	Reason for not testing
UN-SCAN-IT	Point-and-click	15-minute	High price (£315) limits widespread usage
graphreader	Point-and-click or trace	15-minute and <1-minute	Unable to load image file size, unusable for this record
DataThief III	Trace	<1-minute	Colour detection failed to detect lines, unusable for this record
GraphClick	Point-and-click or trace	15-minute and <1-minute	Unavailable on Windows 11, limits widespread usage
Im2graph	Trace	<1-minute	Colour detection failed to detect lines, unusable for this record
WinDig 2.5	Point-and-click or trace	15-minute and <1-minute	Unavailable on Windows 11, limits widespread usage
G3data	Point-and-click	15-minute	Unavailable on Windows 11, limits widespread usage
xyExtract Graph Digitizer	Point-and-click	15-minute	Unable to load image file size, unusable for this record
LogEvolve	Point-and-click	15-minute	Unable to calibrate graph, unusable for this record

3. Results

3.1 Accuracy

First, we consider the accuracy of digitisation across the 11 selected software tools and two LLMs. Figure 2 provides examples from three of the evaluated approaches. Figures 2a-2c show that both categories of digitisation tools closely reproduce the reference data, with only minor deviations, demonstrating high accuracy in the digitisation attempts. In contrast, the outputs generated by ChatGPT and Gemini, shown in Figures 2e and 2f, exhibit substantially larger differences from the reference dataset, indicating considerably lower digitisation accuracy.

Manual digitisation achieved a mean RMSE of 0.07ft, equivalent to 0.5% of the marigram range, making it the most accurate approach evaluated (Figure 3a). It was also among the most consistent approaches, with a SD of 0.03ft (Table 1)

The point-and-click tools produced slightly higher RMSE values overall, but their accuracy was statistically comparable to that of manual digitisation (Figure 3a). Among the point-and-click tools, Engauge Digitizer recorded the lowest mean RMSE of 0.08 ft, equivalent to 0.6% of the marigram range, whereas GetData Graph Digitizer recorded the highest mean RMSE of 0.12 ft, or 1.1% of the range. Despite this variation, after performing Friedman tests no statistically significant differences in accuracy among the point-and-click tools were identified. Several point-and-click tools were also highly consistent across attempts: GRABIT and Engauge Digitizer both recorded SDs of 0.02 ft (Table 1).

The trace tools produced higher RMSE values overall, with accuracy comparable to the point-and-click tools but significantly lower than manual digitisation. WebPlotDigitizer was the most accurate trace tool, with a mean RMSE of 0.11 ft (0.81% of the marigram range), compared with 0.24 ft (approximately 1.7%) for both DigitGraph and Plot Digitizer (Figure 3a). Friedman and post hoc Wilcoxon tests, with Bonferroni correction, confirmed that WebPlotDigitizer was significantly more accurate than both DigitGraph and Plot Digitizer. Despite these differences, NUNIEAU and WebPlotDigitizer each had an SD of 0.03 ft, matching manual digitisation (Table 1).

The LLMs produced the largest and most variable RMSE values, indicating poor accuracy and consistency (Figure 3a). Gemini performed worst, with a mean RMSE of 5.67 ft (42% of the marigram range) and an SD of 1.02 ft (Table 1). ChatGPT performed significantly better but the errors remain unacceptable to use, with a mean RMSE of 1.58 ft (11.7%) and an SD of 0.37

ft. Examples of these inaccuracies are evident in Figure 2, where Gemini mis-assigned the times and dates, while both LLMs have inaccurately smoothed the data. The nature and extent of the errors exhibited by the digitisation attempts of the LLM changes between attempts (Supplementary Material2: Figures 37-42).

In defining an accurate approach as one where the RMSE are below 1% of the marigram range (0.14ft or ~0.043m), manual digitisation, WebPlotDigitizer and all point-and-click tools can be classified as accurate. The remaining trace tools and both LLMs exceeded the acceptable error threshold and were therefore not considered capable of accurate digitisation under the conditions evaluated. Comparisons between all digitised outputs and the reference dataset are provided in Supplementary Material 2.

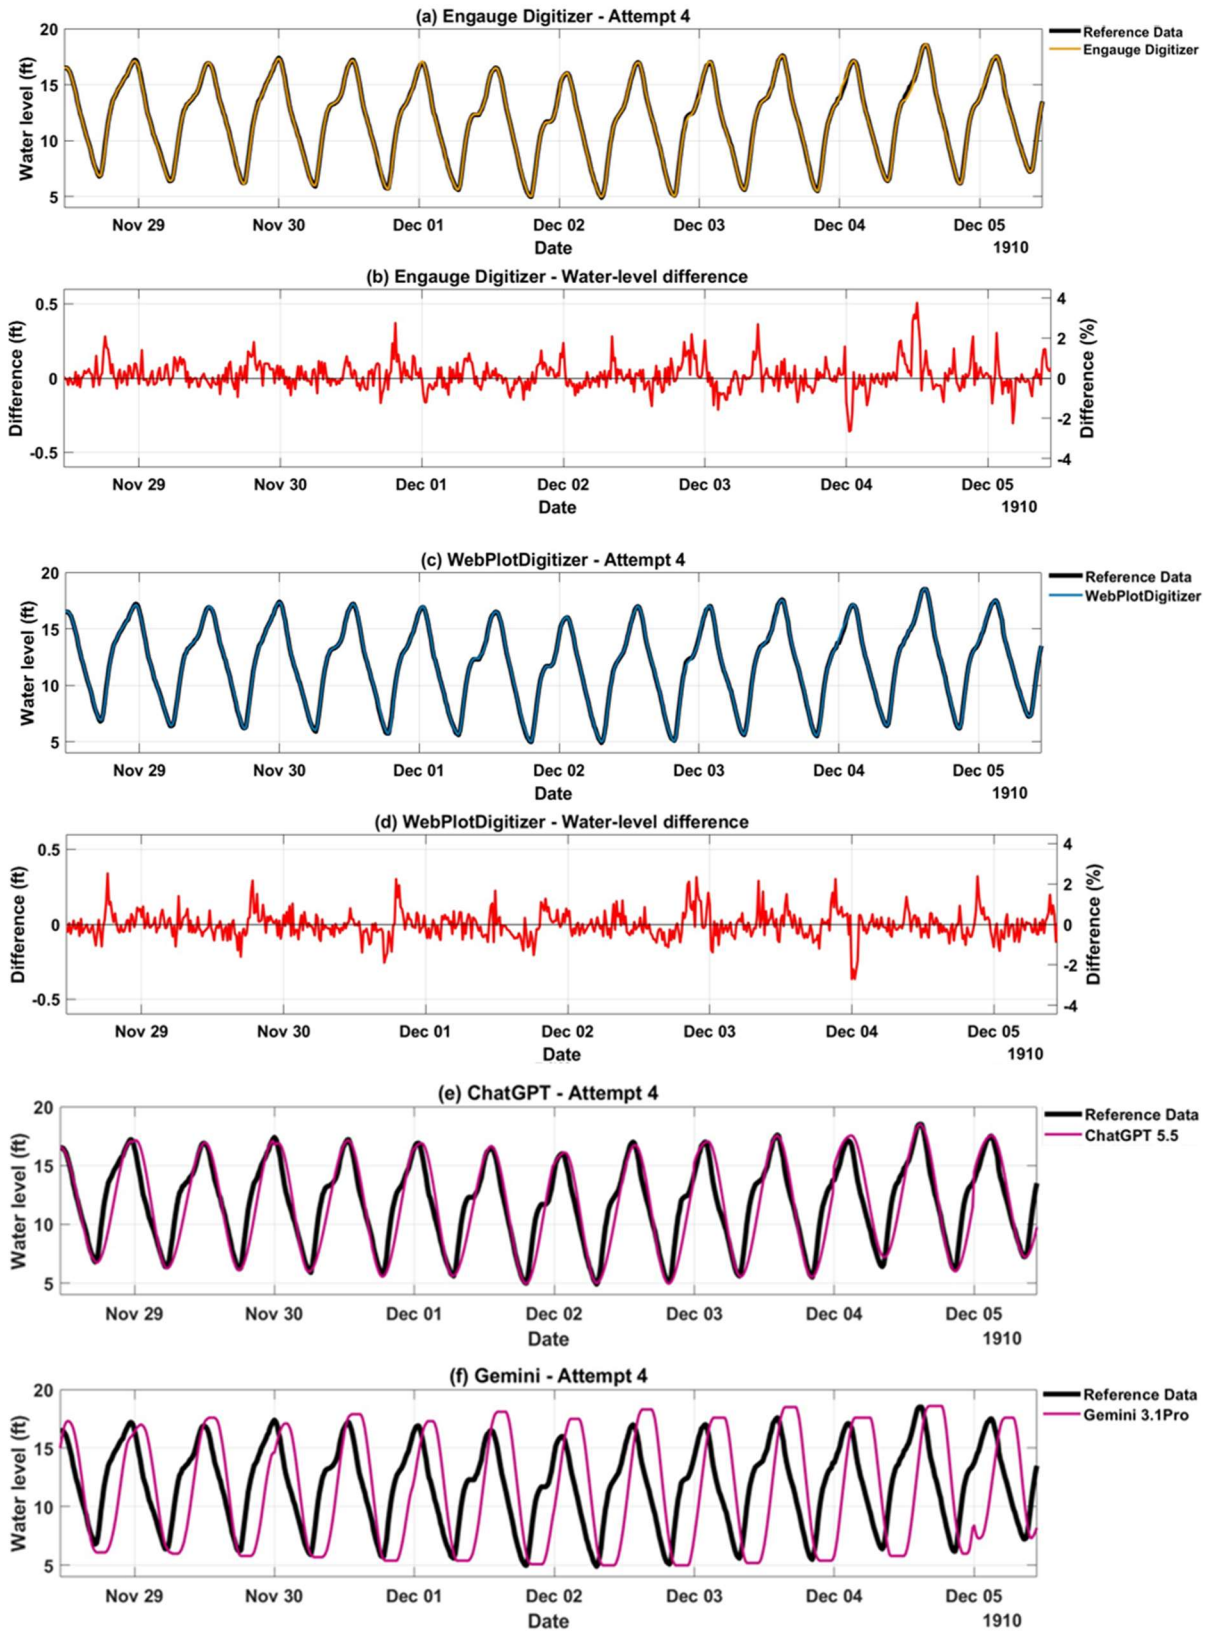

Figure 2. *The digitised output from the fourth test of* **(a)** *Engauge Digitizer and* **(b)** *WebPlotDigitizer overlaid on top of the reference dataset with* **(c)** *and* **(d)** *showing the differences between the output of the tool and the reference dataset. The digitised output of* **(e)** *ChatGPT and* **(f)** *Gemini plotted against the reference dataset.*

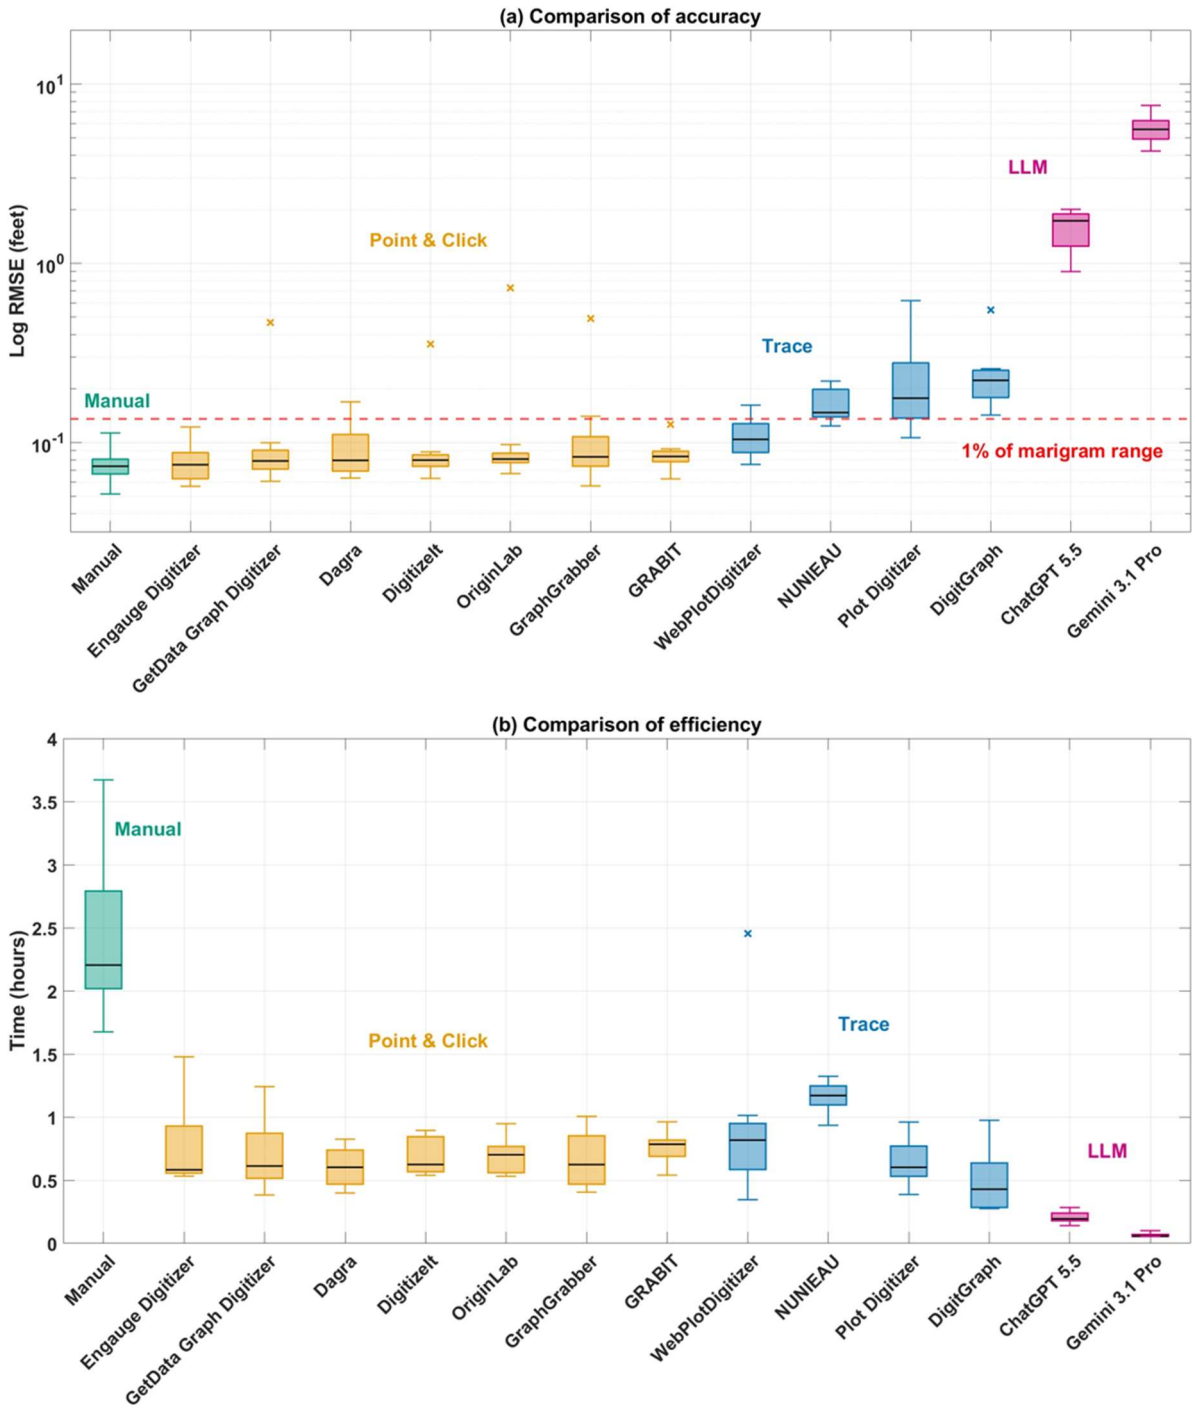

Figure 3. (a) *Boxplots of RMSE values for all digitisation approaches. The colour of the boxplot refers to the methodological categories in which it belongs. The boxplots are ordered from the lowest median RMSE value (Manual), to the highest (Gemini). The y-axis is presented on a logarithmic scale.* **(b)** *Boxplots of the time taken (in minutes) to digitise the test marigram to a 15-minute resolution by each digitisation tool. The colour of the boxplot refers to the methodological categories in which it belongs. The boxplots are ordered from the lowest median RMSE value (Manual), to the highest (Gemini).*

3.2 Speed of Digitisation

Second, we consider the time taken to digitise the selected marigram, across the 11 selected software tools and two LLMs, and the 9 researchers.

Manual digitisation was the slowest approach, requiring a mean of 2.44hrs to digitise a week of data (Figure 3b). At this rate, a continuous 100-year record would require 12,688hrs or ~6.10 years of full-time work. It also showed the greatest variability, with a SD of 0.65hrs (Table 1).

All point-and-click tools met the predefined accuracy threshold while substantially reducing digitisation time relative to the manual approach. Mean completion times ranged from 0.60hrs for Dagra to 0.78hrs for Engauge Digitizer (Figure 3b), equivalent to approximately 1.50–1.95 years of full-time work for a continuous 100-year record. Completion-time variability was moderate, with SDs ranging from 0.13hrs for GRABIT to 0.32hrs for Engauge Digitizer (Table 1). No significant differences in completion time were found among the point-and-click tools.

WebPlotDigitizer, the only trace tool to meet the predefined accuracy threshold, required a mean of 0.91hrs to digitise one week of data (Figure 3b). This corresponds to approximately 4,732hrs, or 2.3 years of full-time work, for a continuous 100-year record. It was significantly faster than manual digitisation but slower than the point-and-click tools and showed high variability in completion time (SD = 0.62hrs; Table 1).

The remaining trace tools exceeded the predefined error threshold but were significantly faster than manual digitisation (Figure 3b). Mean completion times ranged from 0.48hrs per week for DigitGraph to 1.16hrs for NUNIEAU, representing reductions of 80.3% and 52.5%, respectively. For a continuous 100-year record, these rates correspond to approximately 1.20–2.90 years of full-time work. Completion times were relatively consistent, with SDs ranging from 0.12hrs for NUNIEAU to 0.24hrs for DigitGraph (Table 1).

The LLMs were substantially faster than all other approaches, requiring a mean of 0.06hrs per week for Gemini and 0.20hrs for ChatGPT (Figure 3b). However, their large errors rendered their speed of little practical value.

3.3 Relationship of Accuracy and Speed

Figure 4 indicates an overall negative relationship between digitisation time and RMSE, suggesting that slower approaches generally produced more accurate results. Across all approaches, Spearman's rank correlation showed a significant negative association ($\rho = -0.47$, $p = 3.38 \times 10^{-8}$). However, after excluding the LLMs, which were clear outliers, the relationship, while still negative, was non-significant ($\rho = -0.16$, $p = 0.11$). Full statistical results are provided in Supplementary Material 3.

The non-significant relationship between time and accuracy is unsurprising as the digitisation tools are intended to increase digitisation speed while maintaining accuracy, although their effectiveness varied. While greater care may improve accuracy within a given workflow, additional time cannot overcome set limitations within the design of the tools such as calibration set-up, point extraction, or interface design. Completion time is therefore not the sole determinant of digitisation accuracy, which was also influenced by tool design and usability.

Figure 4 also reveals distinct groupings among the approaches. Manual digitisation is separated by its substantially longer completion time; however, it shows similar accuracy to the point-and-click tools. The point-and-click and trace tools occupy broadly similar ranges of digitisation time and RMSE, although the trace tools show greater variability in both accuracy and speed. The LLMs are distinct from all other approaches and from each other, combining very short completion times with substantially higher, but markedly different, RMSE values.

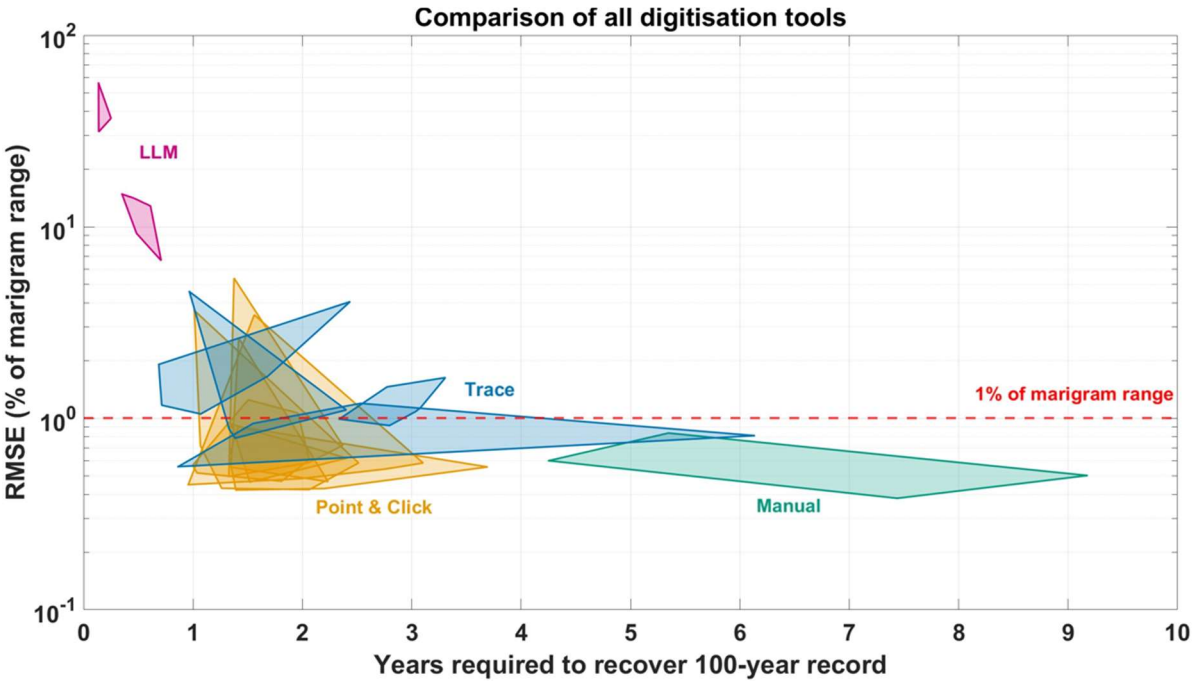

Figure 4. A comparison of all tools by showing the *RMSE as a percentage of the marigram range (shown on a log scale) plotted against number of years required to digitise a 100-year record to a 15-minute resolution. This plot uses working years and assumes that a working year consists of 260 days of 8 hours work. Each polygon represents all attempts for an individual tool, with polygon colours corresponding to the methodological categories.*

4. Discussion

This study evaluated the accuracy and speed of a range of digitisation tools, alongside the current capabilities of some LLMs, for the digitisation of marigrams.

4.1 Manual Digitisation

Manual digitisation was the most accurate approach evaluated but required substantially longer completion times because users had to visually interpret the marigram and transcribe each reading. This finding reinforces and quantifies the time-intensive nature of historical data rescue reported by McLoughlin et al. (2024) (digitising a century of hourly data using WebPlotDigitizer requires ~2.31 working years to complete) and demonstrates the limited scalability of fully manual methods for extensive archival collections.

Despite its high accuracy, manual digitisation remained vulnerable to human error. Participants reported transcription mistakes that required sections to be digitised again, further increasing completion time. Prolonged work may also cause fatigue, reducing both speed and accuracy because the process relies entirely on the digitiser's sustained concentration and judgement

4.2 Point-and-Click Tools

The point-and-click tools digitised marigrams substantially faster than manual digitisation while achieving statistically comparable accuracy. While their mean RMSE values were slightly higher, all met the predefined accuracy threshold of 1% of the marigram range. However, these tools still required considerable user input and sustained attention, which may limit their practicality for long or high-resolution records.

The slightly poorer performance of the point-and-click tools reflects their inability to accommodate image warping. Unlike human digitisers, who can adjust for local distortions, these tools assume straight, evenly spaced axes. Deformation caused by record storage therefore introduces systematic coordinate errors, potentially reducing accuracy below the acceptable threshold for substantially warped marigrams.

Although a range of point-and-click tools were tested, no statistically significant differences in either accuracy or speed were detected among the individual point-and-click tools, which reflects the similarities in their workflows and user interactions.

4.3 Trace Tools

The trace tools reduced digitisation time relative to manual recovery but still required substantial user input. Overall, they were less accurate than the point-and-click tools and significantly less accurate than manual digitisation, partly because they could not accommodate image warping. WebPlotDigitizer was the only trace tool to meet the predefined accuracy threshold and matched the consistency of manual digitisation, whereas Plot Digitizer and DigitGraph produced errors nearly twice the acceptable threshold.

Performance varied substantially among the trace tools, reflecting differences in workflow and interface design. Despite using broadly similar approaches, WebPlotDigitizer was significantly more accurate than Plot Digitizer. It extracted more data points, producing a more detailed trace but requiring additional post-processing to remove false positives, which increased completion time. Plot Digitizer extracted fewer points and was therefore faster, but participants found false positives more difficult to identify where daily traces overlapped because the data markers did not scale with zoom. This interface limitation contributed to its lower accuracy.

NUNIEAU also produced comparatively high RMSE values. Its workflow required users to recalibrate the graph after each trace, increasing both user interaction and completion time while introducing calibration inconsistencies between digitised days.

DigitGraph similarly produced high RMSE values because it relied entirely on user-defined tracing rather than colour-recognition methods. Its performance therefore depended entirely on user precision and was further constrained by the need to draw each trace in a single continuous action, preventing users from zooming in during tracing, and by the inability to edit completed traces.

Although the trace tools were generally less accurate than manual digitisation and the point-and-click tools, they could extract high-frequency data without a proportional increase in completion time. By contrast, increasing the temporal resolution of manual digitisation from 15-minute to 1-minute intervals would require 15 times more readings, potentially increasing the time needed to digitise one week from 2.44hrs to approximately 36.6hrs. Accuracy could

also decline because values would need to be estimated at a finer resolution than the gridline spacing, while the additional time required could increase fatigue and the risk of human error.

Point-and-click tools would likewise require substantially more time at 1-minute resolution because each additional observation would need to be selected manually, although they would likely remain faster than fully manual digitisation. Unlike manual digitisation, however, their accuracy would be less affected by increased temporal resolution, provided that the image is not substantially warped.

4.4 LLMs

Although both LLMs achieved very short digitisation times, their accuracy was substantially lower than that of the manual and software-based approaches. Gemini performed particularly poorly, with a mean RMSE more than an order of magnitude greater than those of the other tools. Because the internal decision-making processes of LLMs are not directly observable, it is difficult to identify which stages of digitisation caused the errors or how their outputs could be improved through further prompting (Zhao et al., 2024).

These findings indicate that the LLMs evaluated are not currently suitable for accurate and reliable marigram digitisation. Although their capabilities may improve as LLMs continue to develop rapidly (Naveed et al., 2025), substantial further development and evaluation will be needed before they can accurately digitise marigrams of varying formats and quality with minimal user intervention.

4.5 Study Limitations

A key limitation of this study is that all evaluations were conducted using a single marigram. Consequently, the results may not fully represent tool performance across records of varying quality, complexity, and format. Additionally, the relatively small number of participants also means that individual differences may have disproportionately influenced the observed variability. Furthermore, restricting the evaluation to data extracted at 15-minute intervals may have understated the relative advantages of trace tools, which can generate observations at sub-minute resolution without a proportional increase in completion time.

5. Conclusion

The results of this study demonstrate that digitisation tools can substantially reduce the time required to digitise historical marigram records while maintaining acceptable error levels. Point-and-click tools provided the most favourable balance between accuracy and speed when extracting observations at clearly identifiable, fixed temporal intervals, achieving accuracy comparable to manual digitisation in substantially less time. However, trace tools, particularly WebPlotDigitizer, may be more appropriate when higher-frequency data are required. Although LLMs are now capable of producing data from simple graphs they are not at the time of this study a suitable method for such complex digitisations. Nevertheless, all sufficiently accurate approaches still require substantial user input, meaning that the digitisation of extensive archival datasets remains time intensive. This highlights the continued need for more efficient and reliable automated methods for large-scale marigram digitisation.

Author Contributions

Christian Kenwright: Conceptualization; Investigation; Testing; Writing; Visualization; Resources; Data Curation

Ivan D. Haigh: Conceptualization; Writing - Review & Editing; Visualization; Resources; Obtained funding

Joanne Williams: Conceptualization; Writing - Review & Editing; Visualization; Resources

Robert Nicholls: Conceptualization; Writing - Review & Editing; Visualization; Resources

Steve McFarland: Resources

Marc Becker: Writing – Review & Editing; Resources

Nurul Tazaroh: Testing; Writing – Review & Editing

Seyedeh Fardis Pourreza Ahmadi: Testing; Writing – Review & Editing

Callum Slade: Testing; Writing – Review & Editing

Molly A. Phillips: Testing; Writing – Review & Editing

Sam E. Pearson-Smith: Testing; Writing – Review & Editing

Minnie J. Darby: Testing; Writing – Review & Editing

Patrick Sharpe: Testing; Writing – Review & Editing

Duo Chan: Conceptualization; Writing - Review & Editing; Visualization; Resources

Acknowledgements

This research was funded by the NERC Flood-CDT.

We would like to thank Frederic Pons and Patrick McLoughlin for their guidance and support in the use of some digitisation tools.

Conflict of Interest

The authors declare no conflicts of interest regarding this paper.

Availability Statement

The dataset generated from testing the digitisation tools and LLMs in this paper is stored on GitHub at: ckenw77/Digitisation-Tools-Test: MATLAB scripts for loading, analysing and plotting testing data from digitisation tool tests

The MATLAB scripts used to analyse the dataset and generate Figures 2-4 as well as those in Supplementary Materials 2 and 3 is available on GitHub at: ckenw77/Digitisation-Tools-Test: MATLAB scripts for loading, analysing and plotting testing data from digitisation tool tests

Links to the digitisation tools discussed in this paper can be found at: TideVault

References

Askarov, M. et al. (2025) 'Preserving Historical Documents Using OCR and Natural Language Processing (NLP)', 2025 International Conference on Computational Innovations and Engineering Sustainability (ICCIES). 2025 International Conference on Computational Innovations and Engineering Sustainability (ICCIES), pp. 1–6. Available at: https://doi.org/10.1109/ICCIES63851.2025.11032769.

Bhaskar, M. et al. (2025) Predicting the Past: Estimating Historical Appraisals with OCR and Machine Learning, arXiv.org. Available at: https://doi.org/10.1145/3715335.3735488.

Bormann, I. (2025) 'DigitizeIt'. DigitizeIt. Available at: https://www.digitizeit.xyz/.

Dagdelen, J. et al. (2024) 'Structured information extraction from scientific text with large language models', Nature Communications, 15(1), p. 1418. Available at: https://doi.org/10.1038/s41467-024-45563-x.

Das, A.K., Tarun, M. and Mueller, K. (2026) 'Charts-of-Thought: Enhancing LLM Visualization Literacy Through Structured Data Extraction', IEEE Transactions on Visualization and Computer Graphics, 32(1), pp. 1427–1437. Available at: https://doi.org/10.1109/TVCG.2025.3634813.

Doke, J. (2016) 'GRABIT'. MATLAB Central File Exchange. Available at: https://uk.mathworks.com/matlabcentral/fileexchange/7173-grabit.

Fedorov, S. (2017) 'GetData Graph Digitizer'. GetData Software. Available at: https://getdata-graph-digitizer.com/.

Google DeepMind (2026) 'Gemini 3.1 Pro'. Google LLC. Available at: https://gemini.google.com/app.

Gupta, A. et al. (2015) 'Automatic Assessment of OCR Quality in Historical Documents', Proceedings of the AAAI Conference on Artificial Intelligence, 29(1). Available at: https://doi.org/10.1609/aaai.v29i1.9487.

Haigh, I.D. et al. (2025) 'Coastal flooding', MCCIP Annual Summary 2025: Marine Climate Change Impacts Evidence Updates. Plymouth, UK: Marine Climate Change Impacts Partnership (MCCIP). Available at: https://doi.org/10.14465/2025.anl00.sum.

Hogarth, P. (2014) 'Preliminary analysis of acceleration of sea level rise through the twentieth century using extended tide gauge data sets (August 2014)', Journal of Geophysical Research: Oceans, 119(11), pp. 7645–7659. Available at: https://doi.org/10.1002/2014JC009976.

Jenkins, D. (2020) 'DigitGraph Spreadsheet and WebPlot Digitiser'. Available at: https://newtonexcelbach.com/2020/03/03/digitgraph-spreadsheet-and-webplot-digitiser/.

Kenwright, C. (2026) 'TideVault'. Zenodo. Available at: https://doi.org/10.5281/zenodo.20510038.

Kenwright, C. et al. (In review) 'Unlocking Historical Sea Level Records through Digitisation—A Review and Inventory', Manuscript submitted for publication.

Latapy, A. et al. (2023) 'Data rescue process in the context of sea level reconstructions: An overview of the methodology, lessons learned, up-to-date best practices and recommendations', Geoscience Data Journal, 10(3), pp. 396–425. Available at: https://doi.org/10.1002/gdj3.179.

Lovón-Melgarejo, J. et al. (2025) 'Evaluating LLM Abilities to Understand Tabular Electronic Health Records: A Comprehensive Study of Patient Data Extraction and Retrieval', in C. Hauff et al. (eds) Advances in Information Retrieval. Cham: Springer Nature Switzerland, pp. 153–168. Available at: https://doi.org/10.1007/978-3-031-88711-6_10.

Martinsen, P. (2019) 'Dagra Data Digitizer'. Blue Leaf Software. Available at: https://www.datadigitization.com/ (Accessed: 30 September 2025).

McInnes, K.L. et al. (2024) 'Digitizing the Williamstown, Australia Tide-Gauge Record Back to 1872: Insights Into Changing Extremes', Journal of Geophysical Research: Oceans, 129(8), p. e2024JC020908. Available at: https://doi.org/10.1029/2024JC020908.

McLoughlin, P.J. et al. (2024) 'The accurate digitization of historical sea level records', Geoscience Data Journal, 11(4), pp. 790–805. Available at: https://doi.org/10.1002/gdj3.256.

Mitchell, M. et al. (2020) 'markummitchell/engauge-digitizer: Nonrelease'. Available at: https://doi.org/10.5281/zenodo.3941227.

Naveed, H. et al. (2025) 'A Comprehensive Overview of Large Language Models', ACM Transactions on Intelligent Systems and Technology, 16(5), p. 106:1-106:72. Available at: https://doi.org/10.1145/3744746.

OpenAI (2026) ‘ChatGPT 5.5’. OpenAI. Available at: https://chatgpt.com/?utm_source=chatgpt.com.

OriginLab Corporation (2025) ‘Origin’. OriginLab Corporation. Available at: https://www.originlab.com/index.aspx?go=Products/Origin/2026&pid=5478.

PlotDigitizer PRO App (2026). Available at: https://plotdigitizer.com/app (Accessed: 19 June 2026).

Quintessa Ltd (2020) ‘Graph Grabber’. Quintessa Ltd. Available at: https://www.quintessa.org/software/downloads-and-demos/graph-grabber-2.0.2.

Rohatgi, A. (2024) WebPlotDigitizer (Version 4.8), WebPlotDigitizer. Available at: https://apps.automeris.io/wpd4/ (Accessed: 19 June 2026).

Sabatino, A.D. et al. (2016) ‘Modelling sea level surges in the Firth of Clyde, a fjordic embayment in south-west Scotland’, Natural Hazards, 84(3), pp. 1601–1623. Available at: https://doi.org/10.1007/s11069-016-2506-7.

Singh, A., Bacchuwar, K. and Bhasin, A. (2012) ‘A Survey of OCR Applications’, International Journal of Machine Learning and Computing (IJMLC) [Preprint]. Available at: https://doi.org/10.7763/IJMLC.2012.V2.137.

Talke, S.A. and Jay, D. (2013) ‘Nineteenth Century North American and Pacific Tidal Data: Lost or Just Forgotten?’, Journal of Coastal Research, 29, pp. 118–127. Available at: https://doi.org/10.2112/JCOASTRES-D-12-00181.1.

Ullmann, A., Pons, F. and Moron, V. (2005) ‘Tool kit helps digitize tide gauge records’, Eos, Transactions American Geophysical Union, 86(38), pp. 342–342. Available at: https://doi.org/10.1029/2005EO380004.

Wöppelmann, G. et al. (2014) ‘Rescue of the historical sea level record of Marseille (France) from 1885 to 1988 and its extension back to 1849–1851’, Journal of Geodesy, 88. Available at: https://doi.org/10.1007/s00190-014-0728-6.

Zhao, H. et al. (2024) ‘Explainability for Large Language Models: A Survey’, ACM Transactions on Intelligent Systems and Technology, 15(2), p. 20. Available at: https://doi.org/10.1145/3639372.

Zhao, W.X. et al. (2026) ‘A Survey of Large Language Models’. arXiv. Available at: https://doi.org/10.48550/arXiv.2303.18223.

Supplementary Material 1

Prompt for OpenAI ChatGPT Plus (GPT-4o) and Google Gemini Advanced (Gemini 2.0 Flash)

Extract tidal height data from the provided tide chart image.

Each red curve represents one day, covering November 28 to December 5, 1910 (8 total curves). Curves are labelled with small handwritten numbers near peaks/troughs and may overlap.

Axis definitions:

• **X-axis:** time over a 24-hour period
o Major gridlines: 1 hour
o Minor gridlines: 15 minutes
o Assume the leftmost vertical gridline corresponds to 00:00

• **Y-axis:** water level in feet (range 3–26 ft)
o Major gridlines: 1 ft
o Minor gridlines: 0.25 ft
o Values are absolute (no datum correction required)

Time coverage:

• Nov 28, 1910, starts at ~11:15
• Nov 29–Dec 4, 1910, are full 24-hour days
• Dec 5, 1910, ends at ~10:45

Additional constraints:

• On the rightmost side of the chart, there may be partial curves or gridlines from the following week
• Do not include any data beyond 23:45 for any day
• Ignore any lines that extend past the valid time range for Dec 5

Sampling method:
• Sample strictly at each 15-minute interval using the vertical gridlines

For each time step:
• Locate the vertical gridline corresponding to the timestamp
• Record the water level at the point where the curve intersects that gridline.

If the curve does not pass exactly through the gridline:
• Estimate the height using the closest visible point of the curve at that time slice
• Do NOT infer values between gridlines.

If the curve has visible thickness:
• Sample the **centreline** of the stroke.

Curve tracking rules:

• Each red curve represents one full day
• When curves overlap or intersect, follow the **continuous trajectory of the same curve**
• Use nearby labelled numbers and curve continuity to maintain correct tracking
• Do not switch between curves at intersections

Data limitations:

• Do not extrapolate beyond visible portions of any curve
• Only record values where the curve is clearly visible
• For Nov 28 and Dec 5, only include times where data exists

Output requirements:

• Format: CSV
• Columns:

- date (YYYY-MM-DD)
- time (HH:MM)
- height ft (float, rounded to 2 decimal places)

• Sampling frequency: every 15 minutes
• Clearly separate data by day
• Expected output: ~650–750 rows
• Provide a **single CSV file containing all days**

Notes:

• Prioritize accuracy over speed
• Use careful visual estimation where required
• All values must be directly derived from the image at each sampling point

Supplementary Material 2

1. Manual Comparison Plots

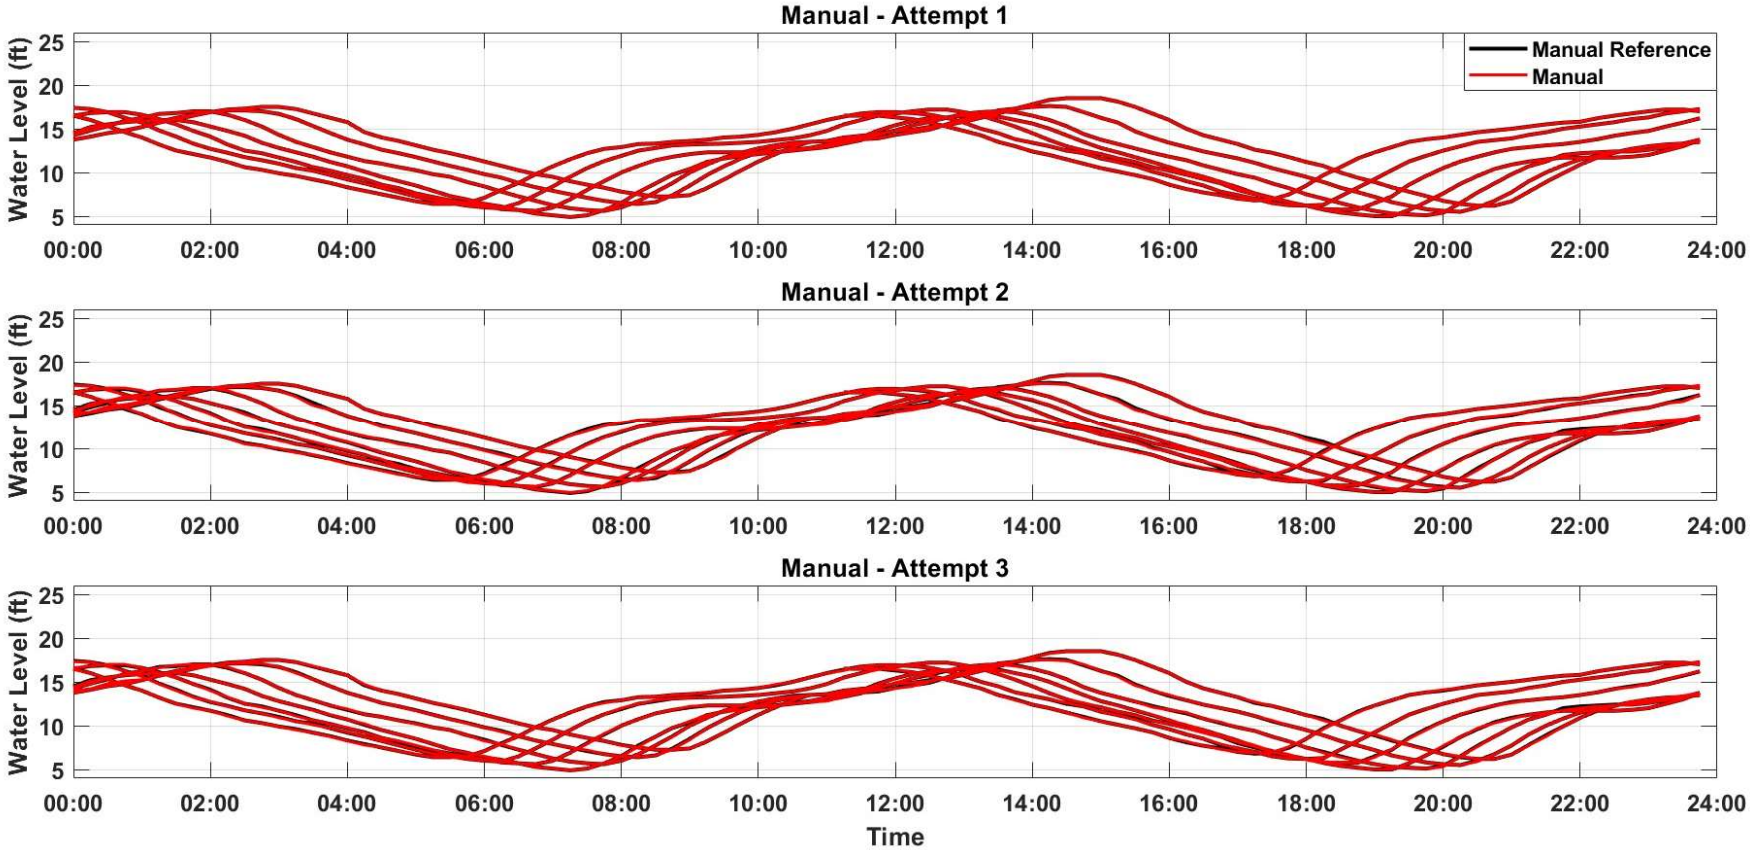

Figure 1. *Attempts 1-3 for manual digitisation of the test marigram, where the red lines show the recovered data and the black shows the reference data.*

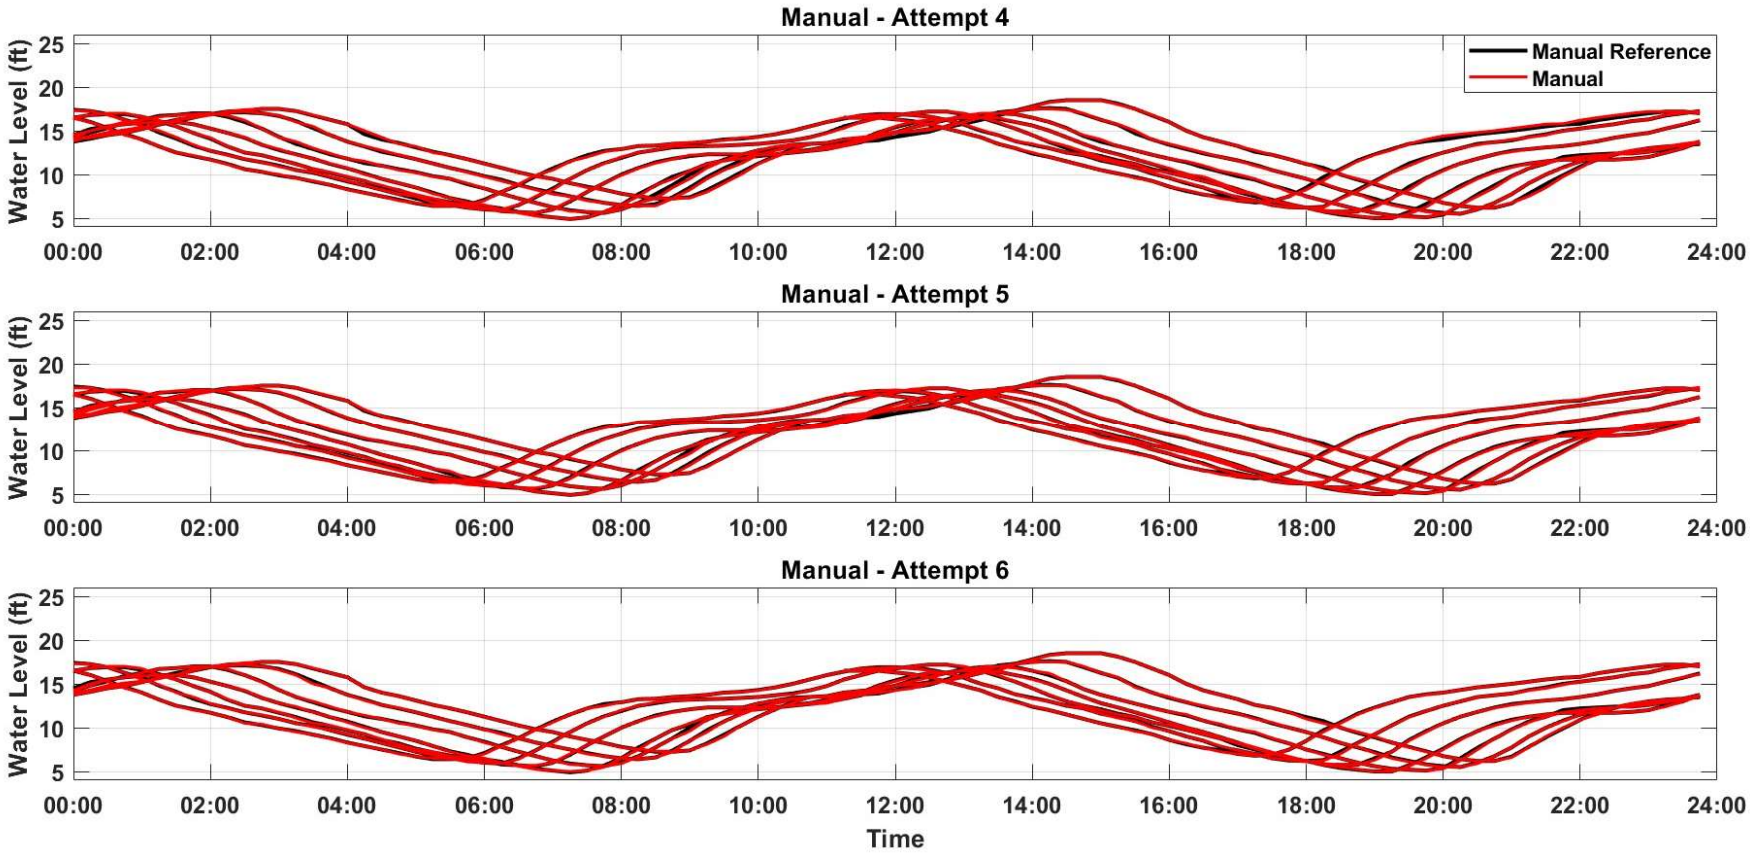

Figure 2. *Attempts 4-6 for manual digitisation of the test marigram, where the red lines show the recovered data and the black shows the reference data.*

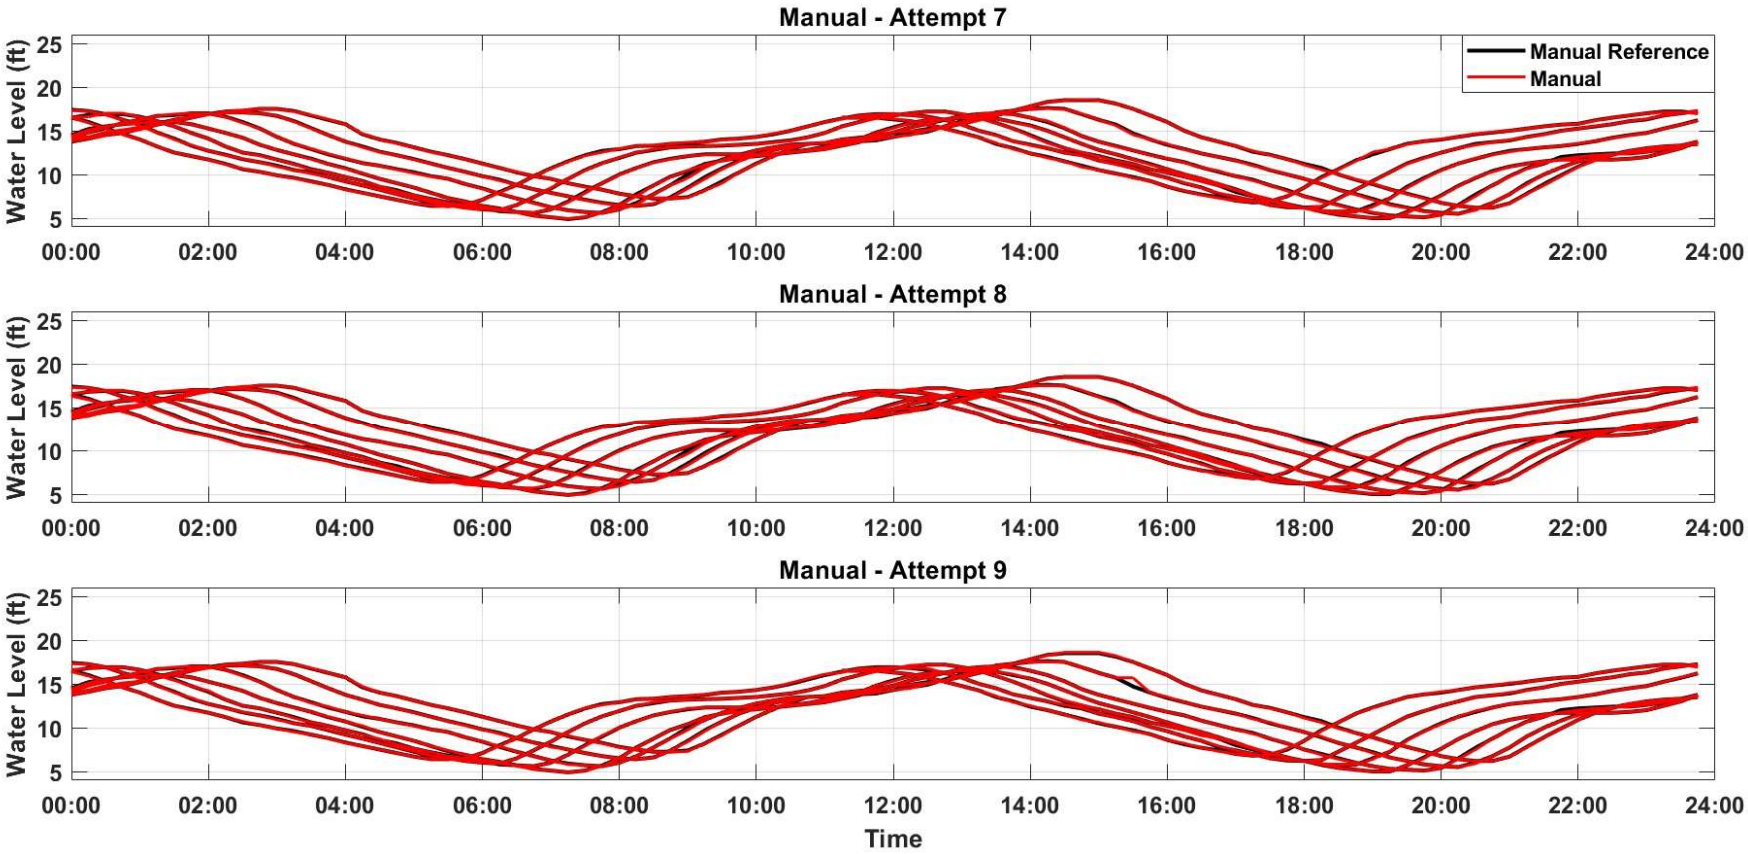

Figure 3. *Attempts 7-9 for manual digitisation of the test marigram, where the red lines show the recovered data and the black shows the reference data.*

2. Point-and-Click Comparison Plots

2.1. Dagra

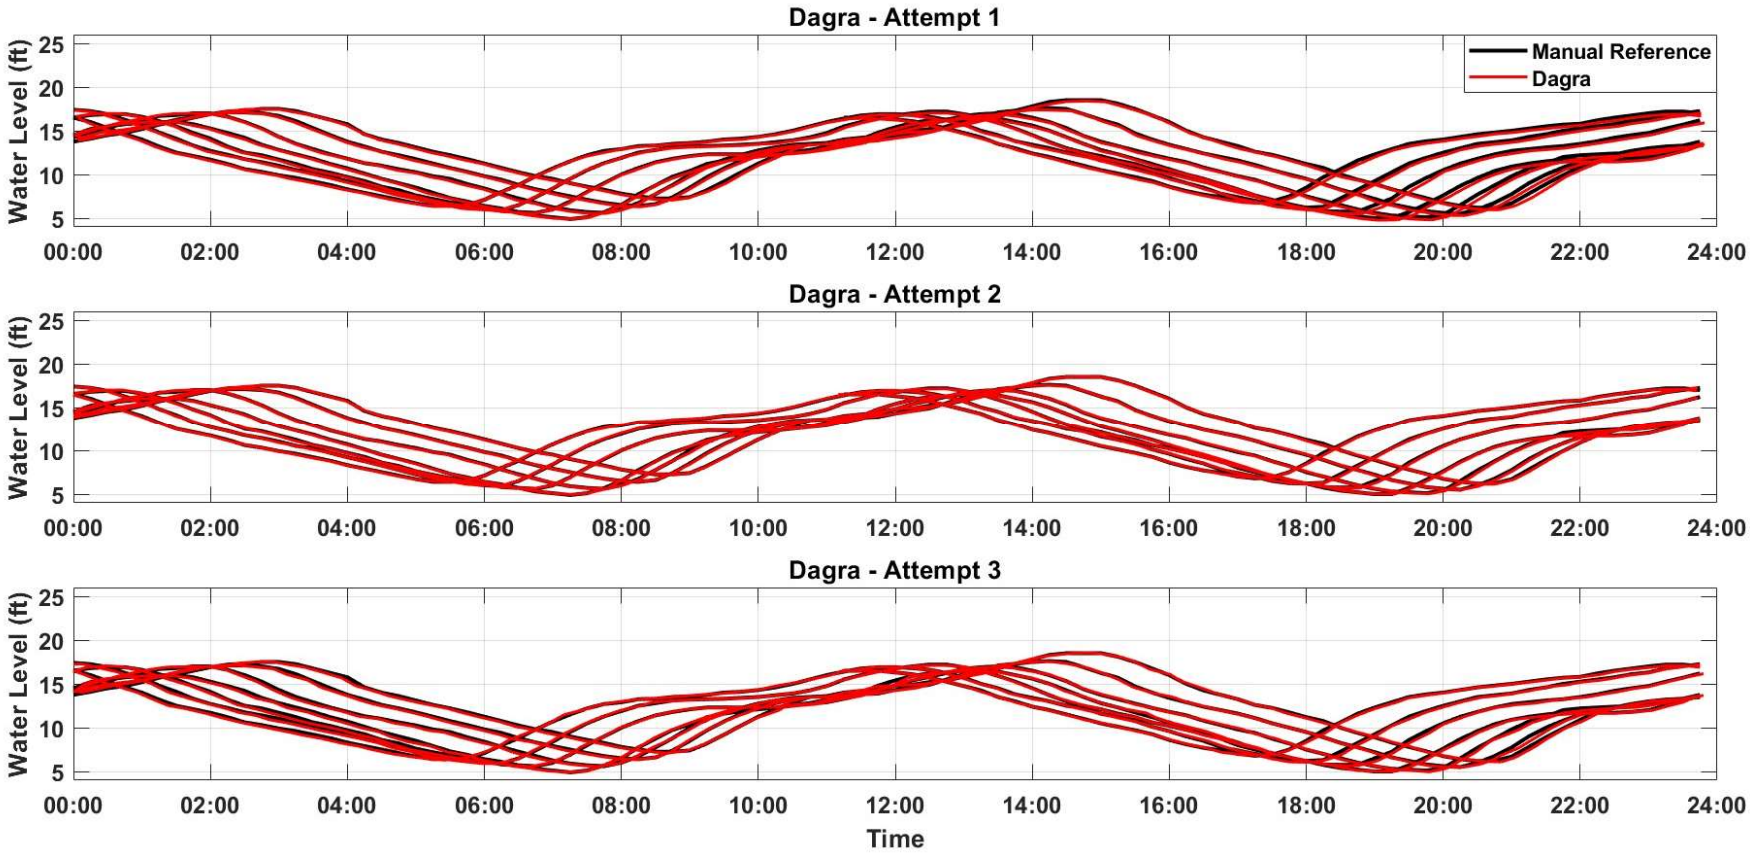

Figure 4. *Attempts 1-3 of using Dagra to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

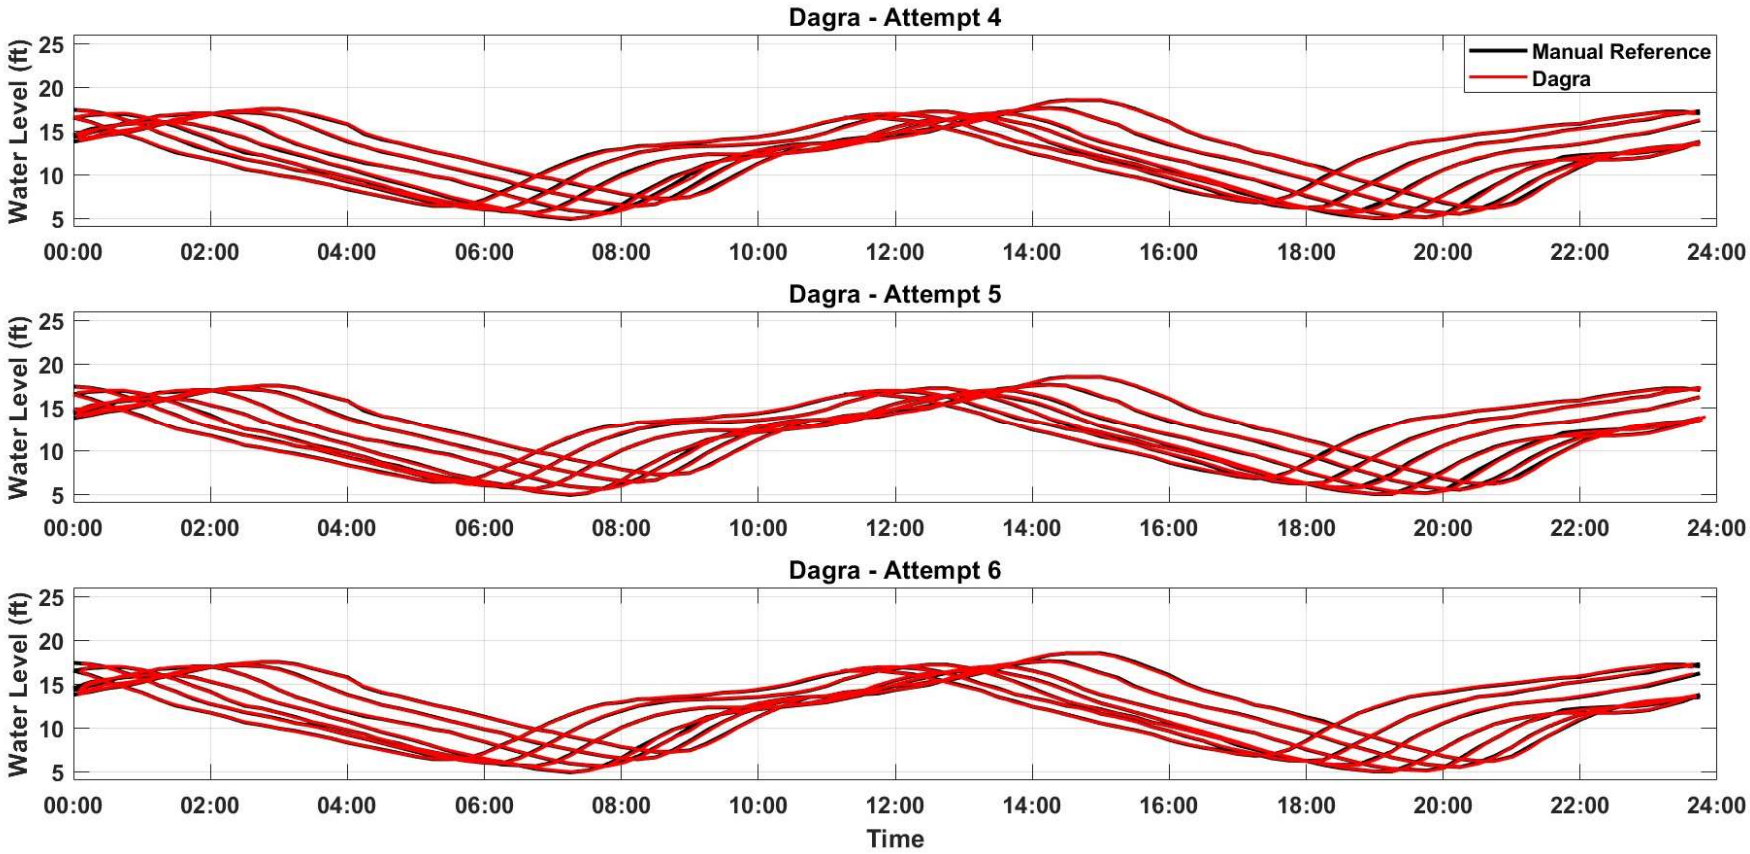

Figure 5. *Attempts 4-6 of using Dagra to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

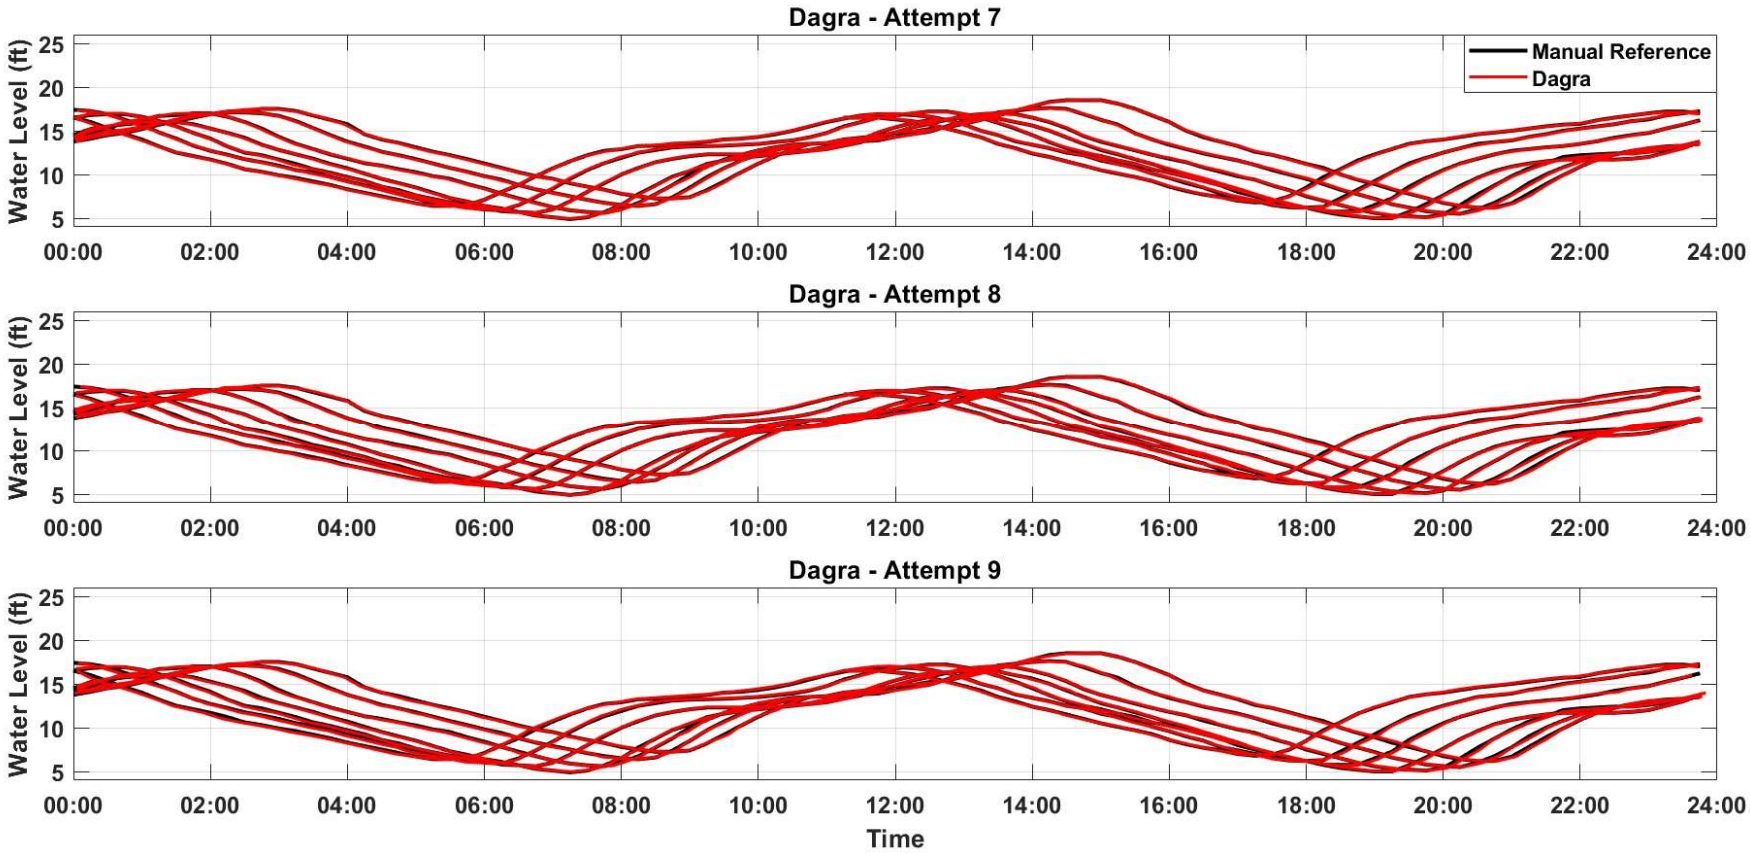

Figure 6. *Attempts 7-9 of using Dagra to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

2.2. DigitizeIt

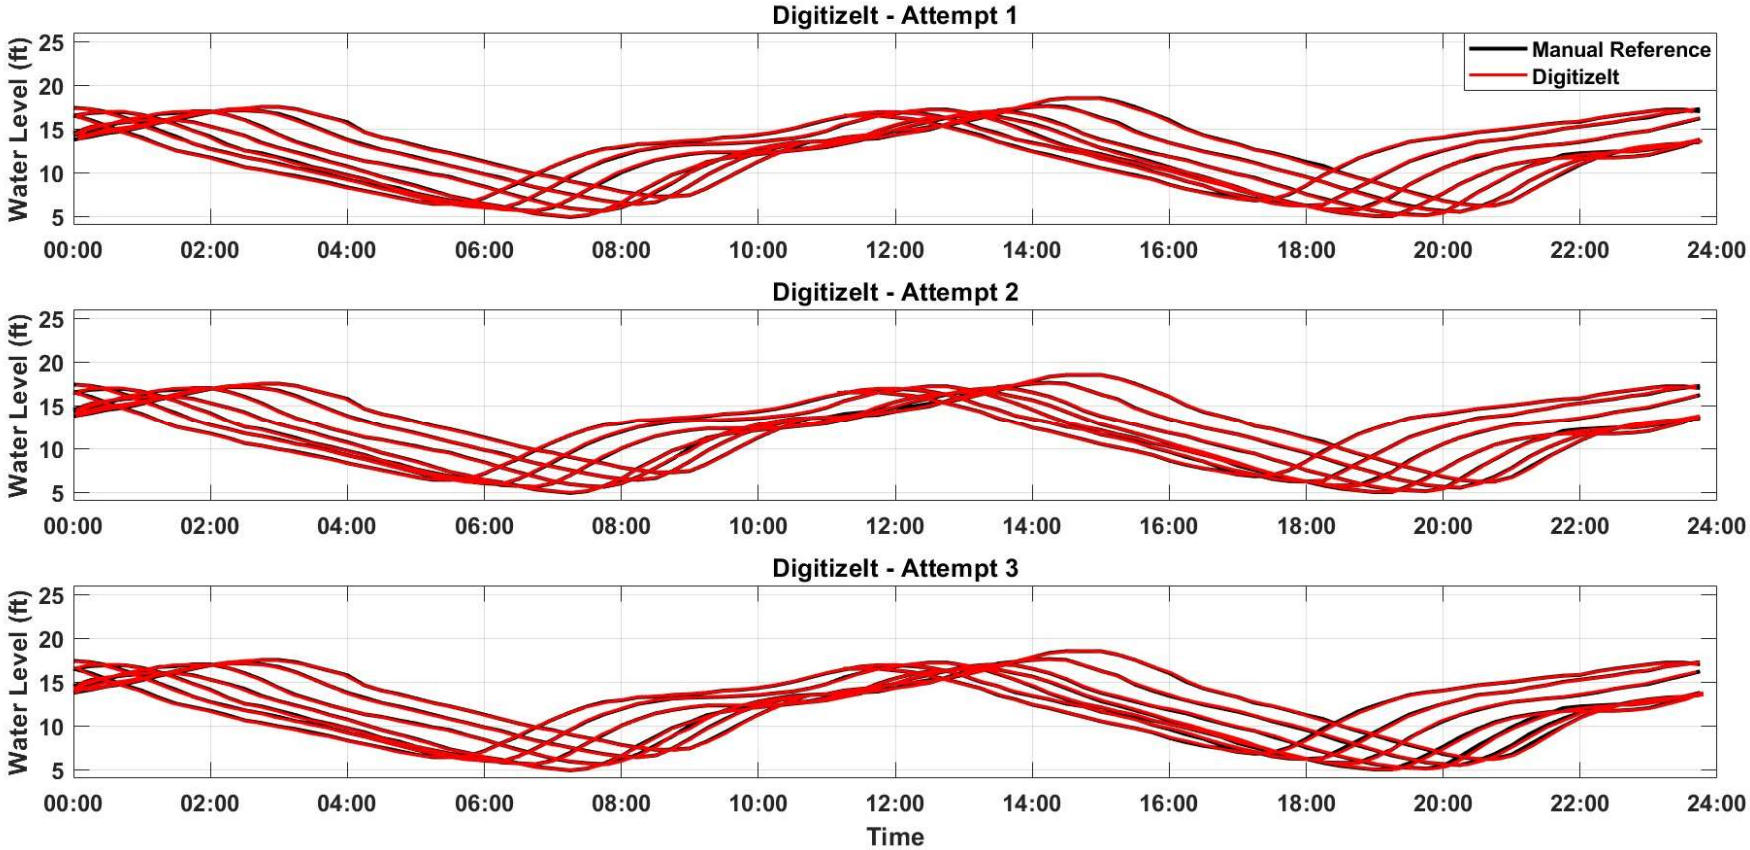

Figure 7. *Attempts 1-3 of using DigitizeIt to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

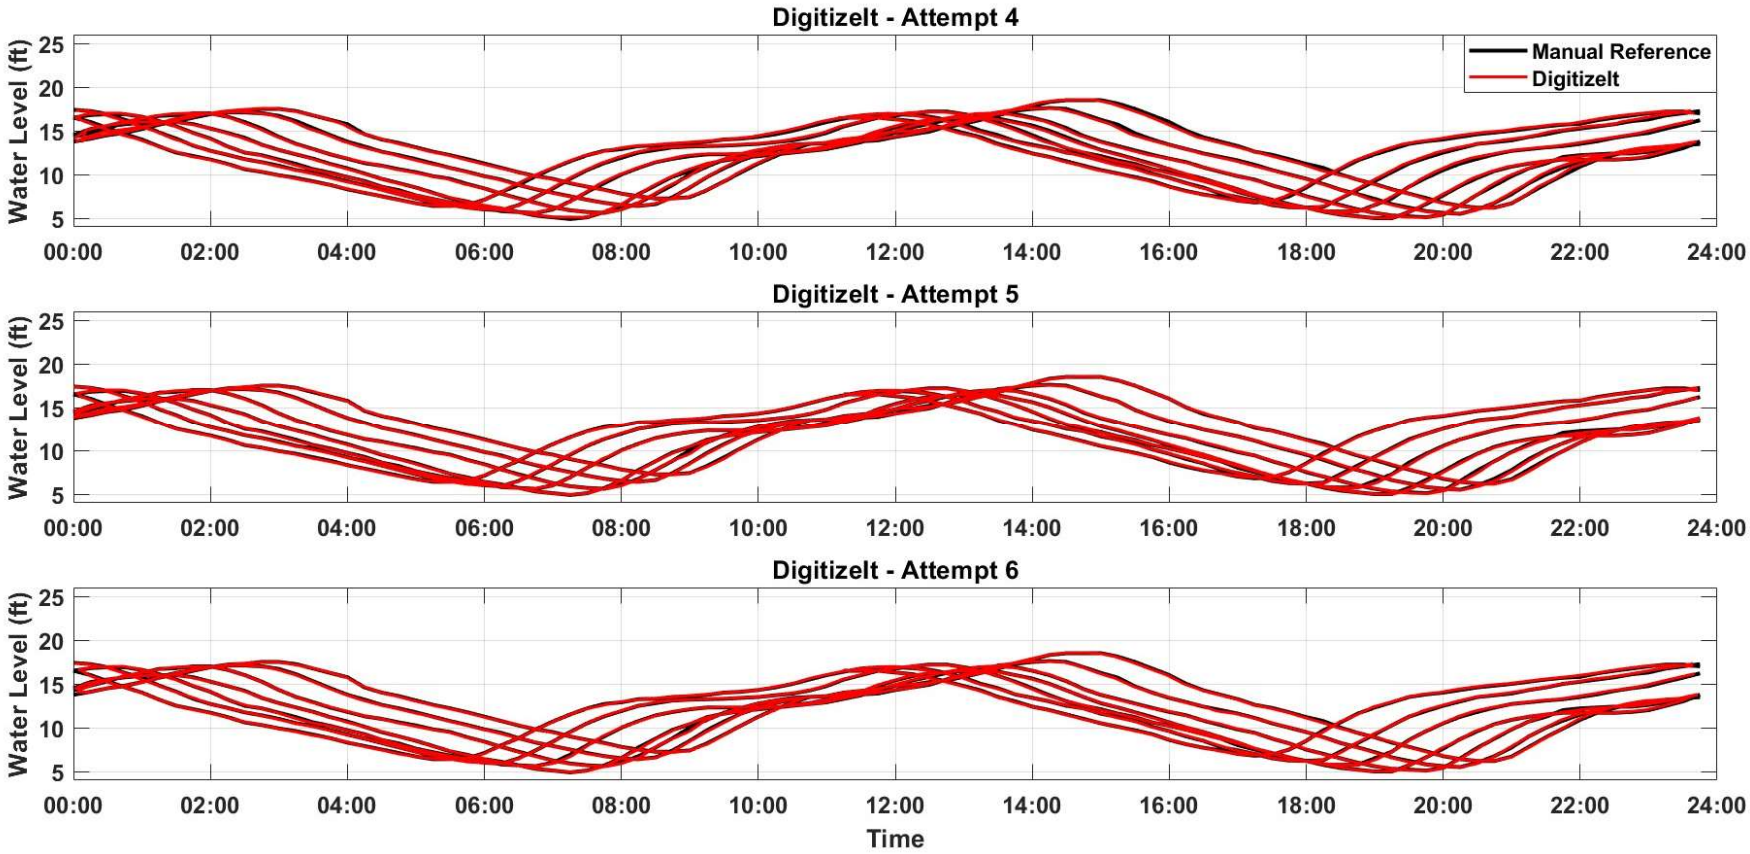

Figure 8. *Attempts 4-6 of using DigitizeIt to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

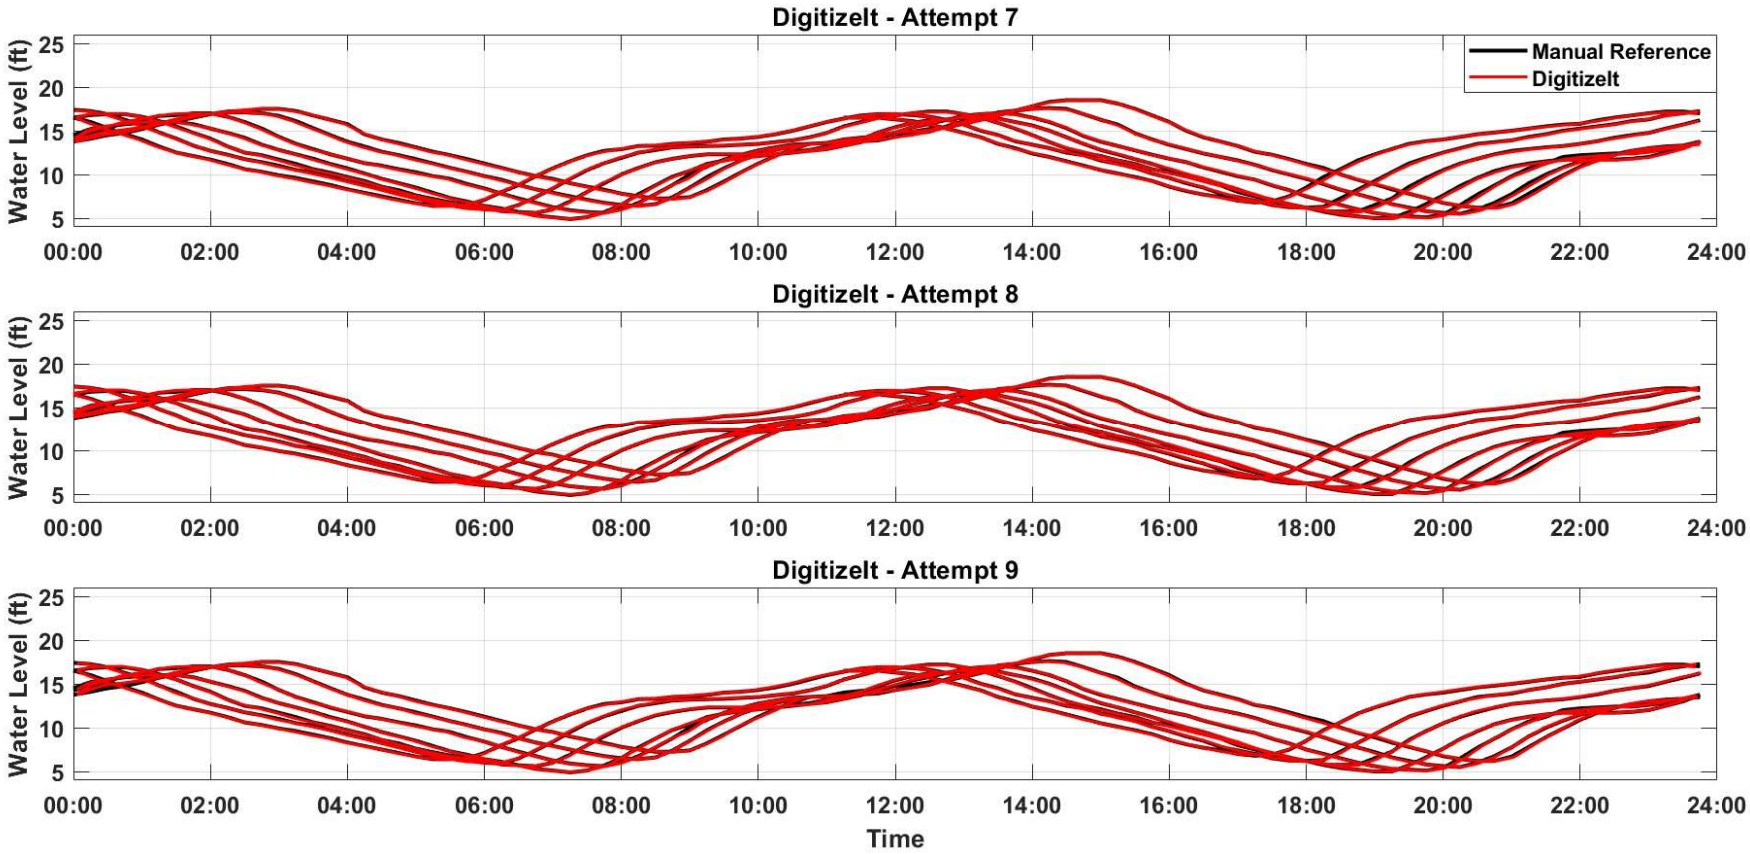

Figure 9. *Attempts 7-9 of using DigitizeIt to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

2.3. Engauge Digitizer

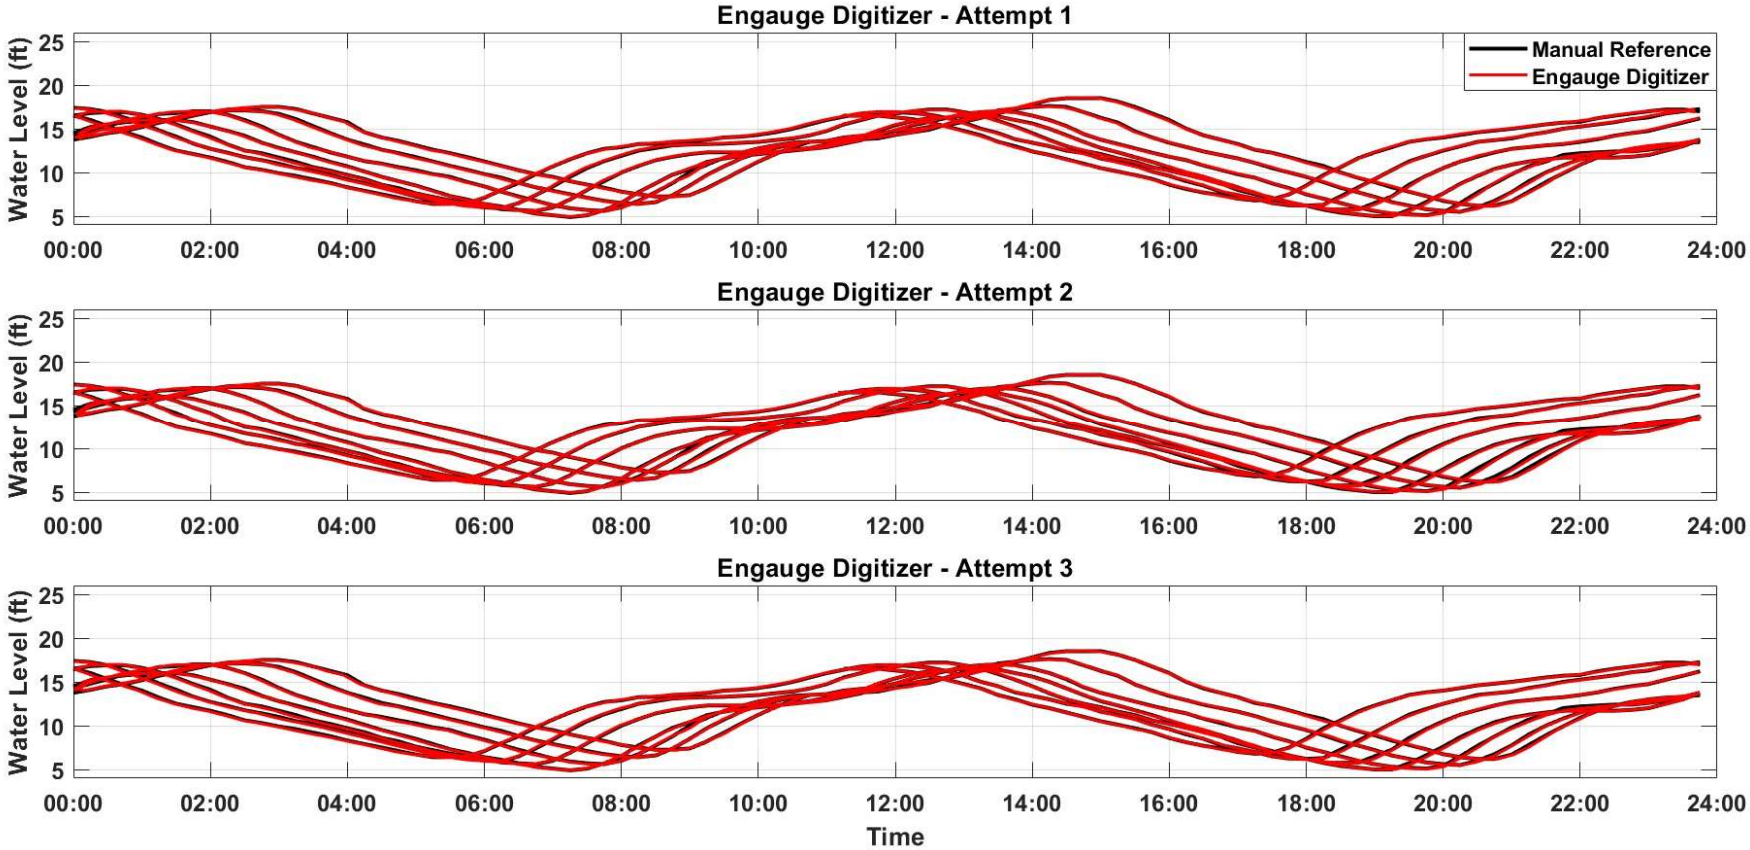

Figure 10. *Attempts 1-3 of using Engauge Digitizer to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

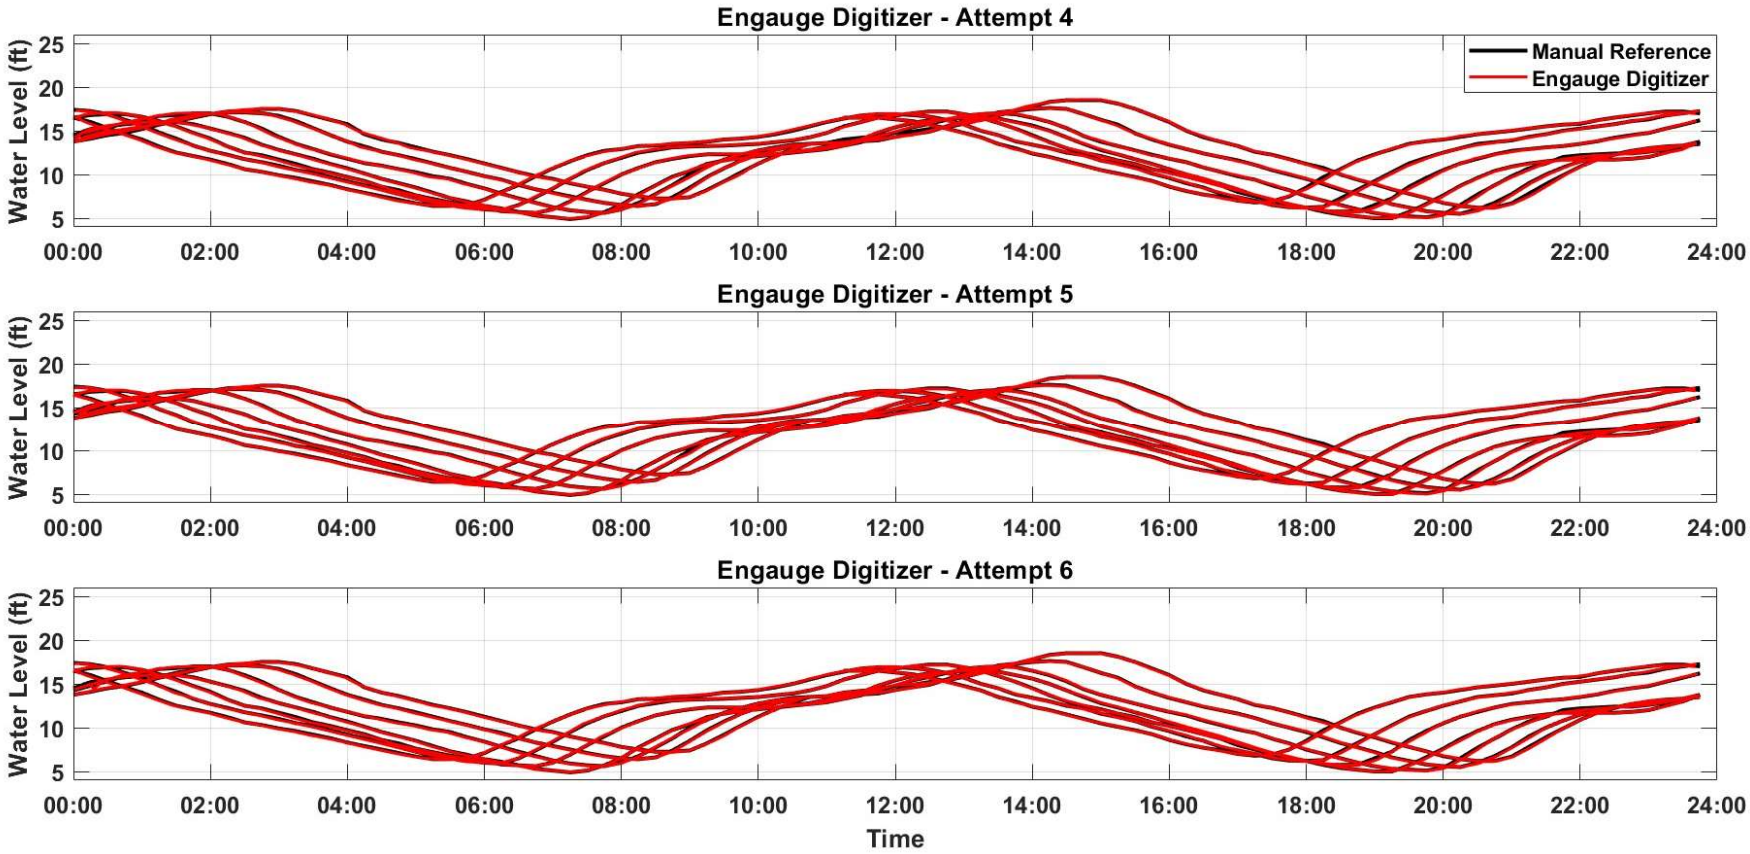

Figure 11. *Attempts 4-6 of using Engauge Digitizer to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

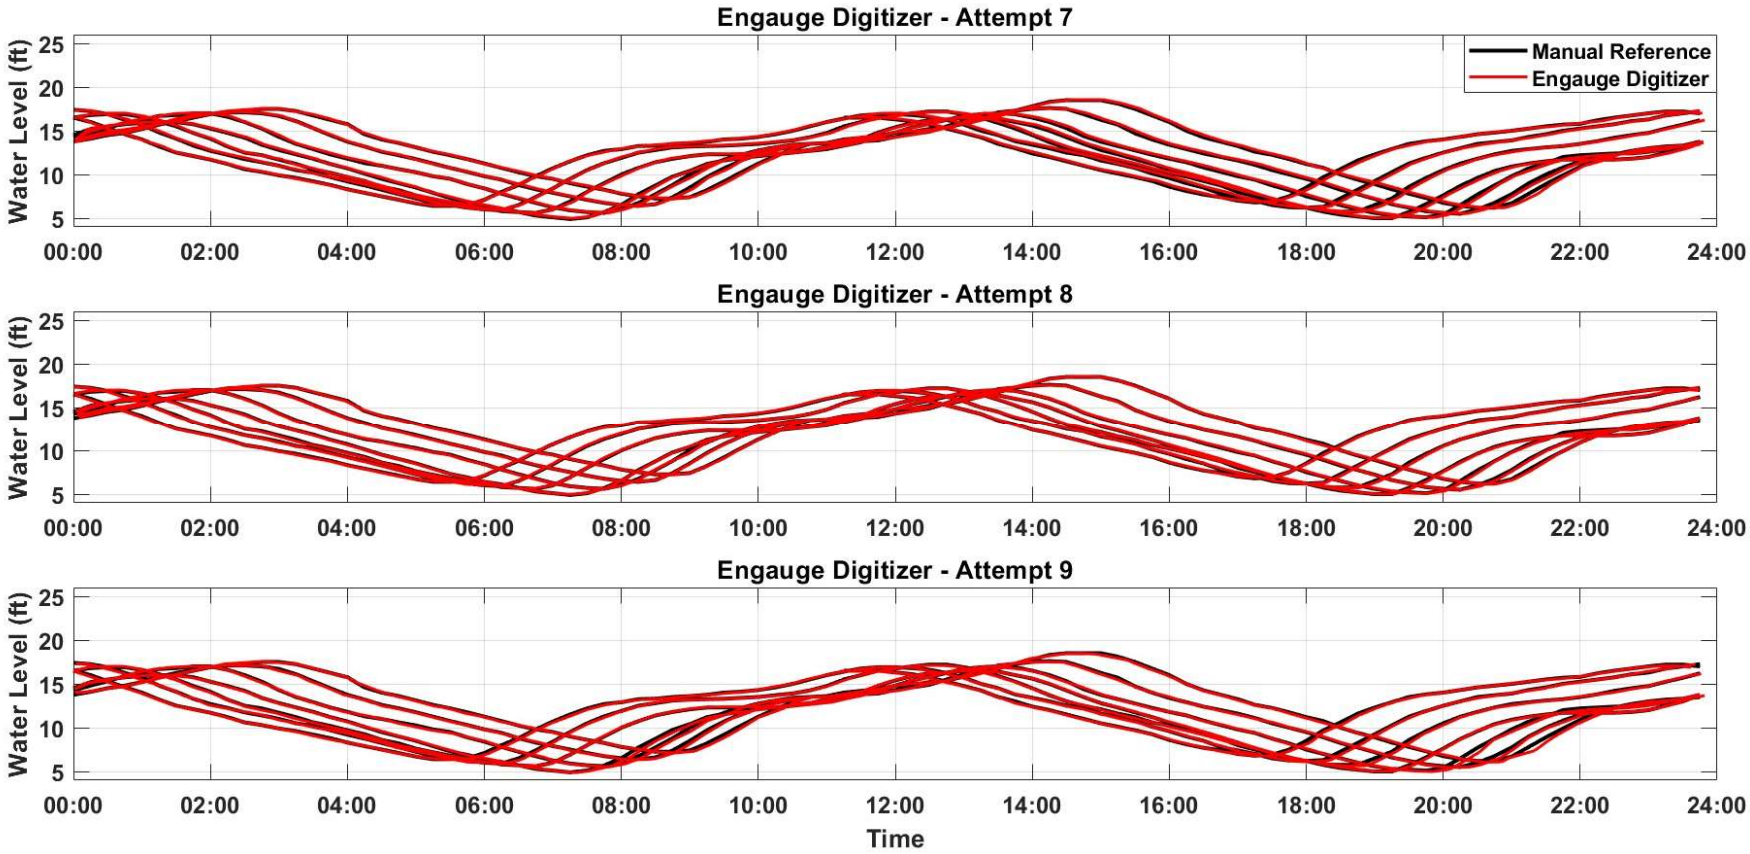

Figure 12. *Attempts 7-9 of using Engauge Digitizer to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

2.4. GetData Graph Digitizer

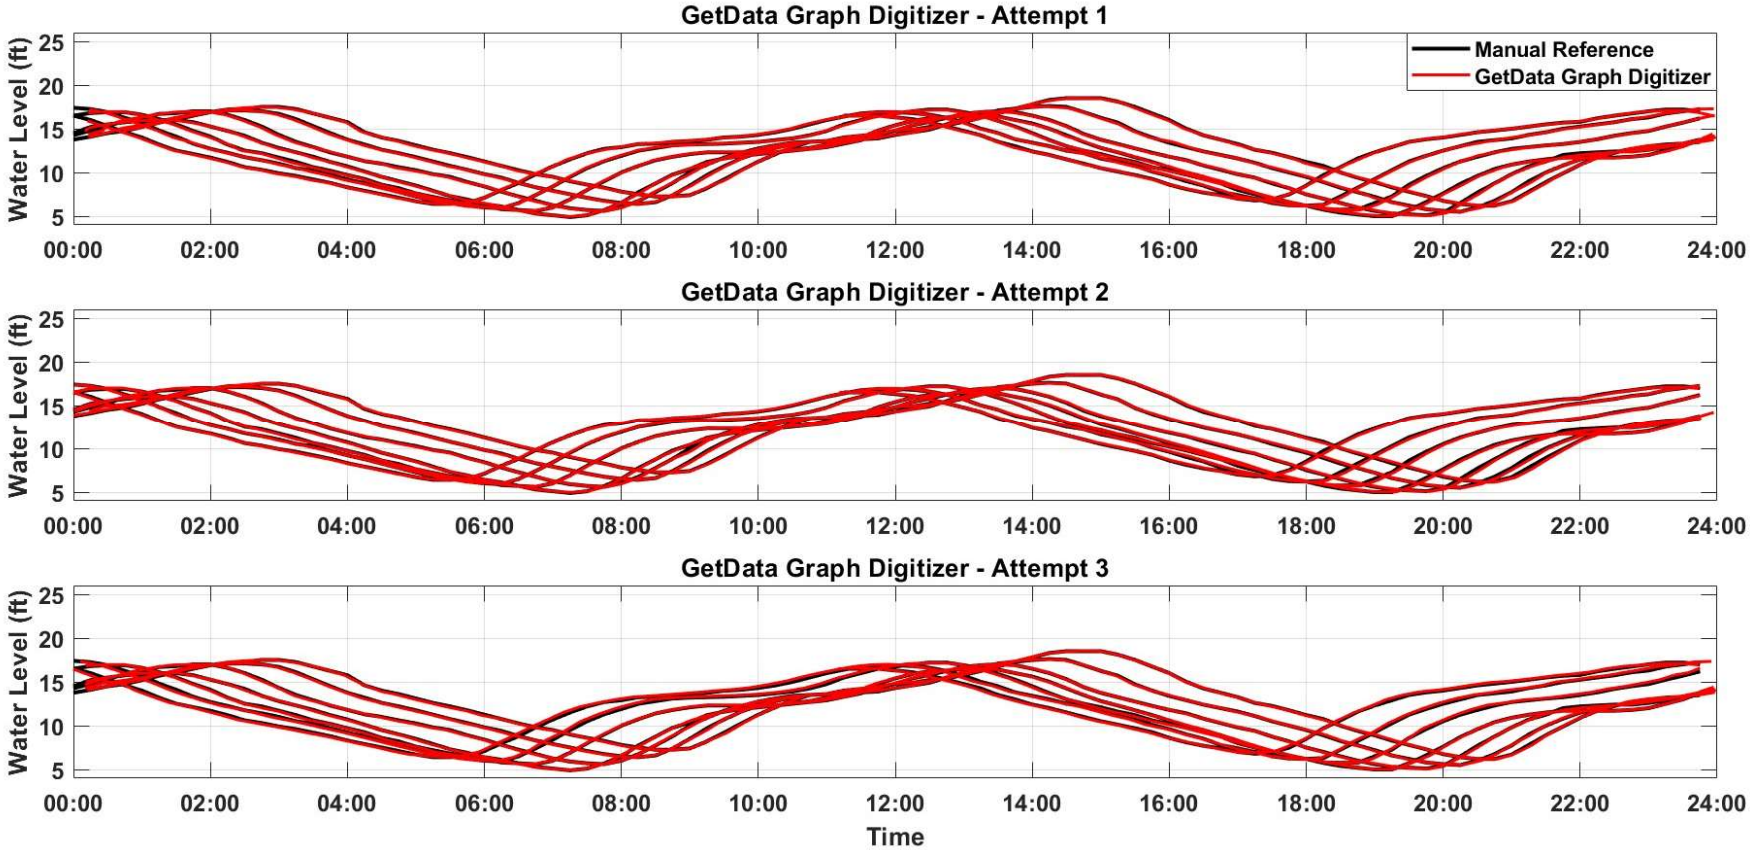

Figure 13. *Attempts 1-3 of using GetData Graph Digitizer to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

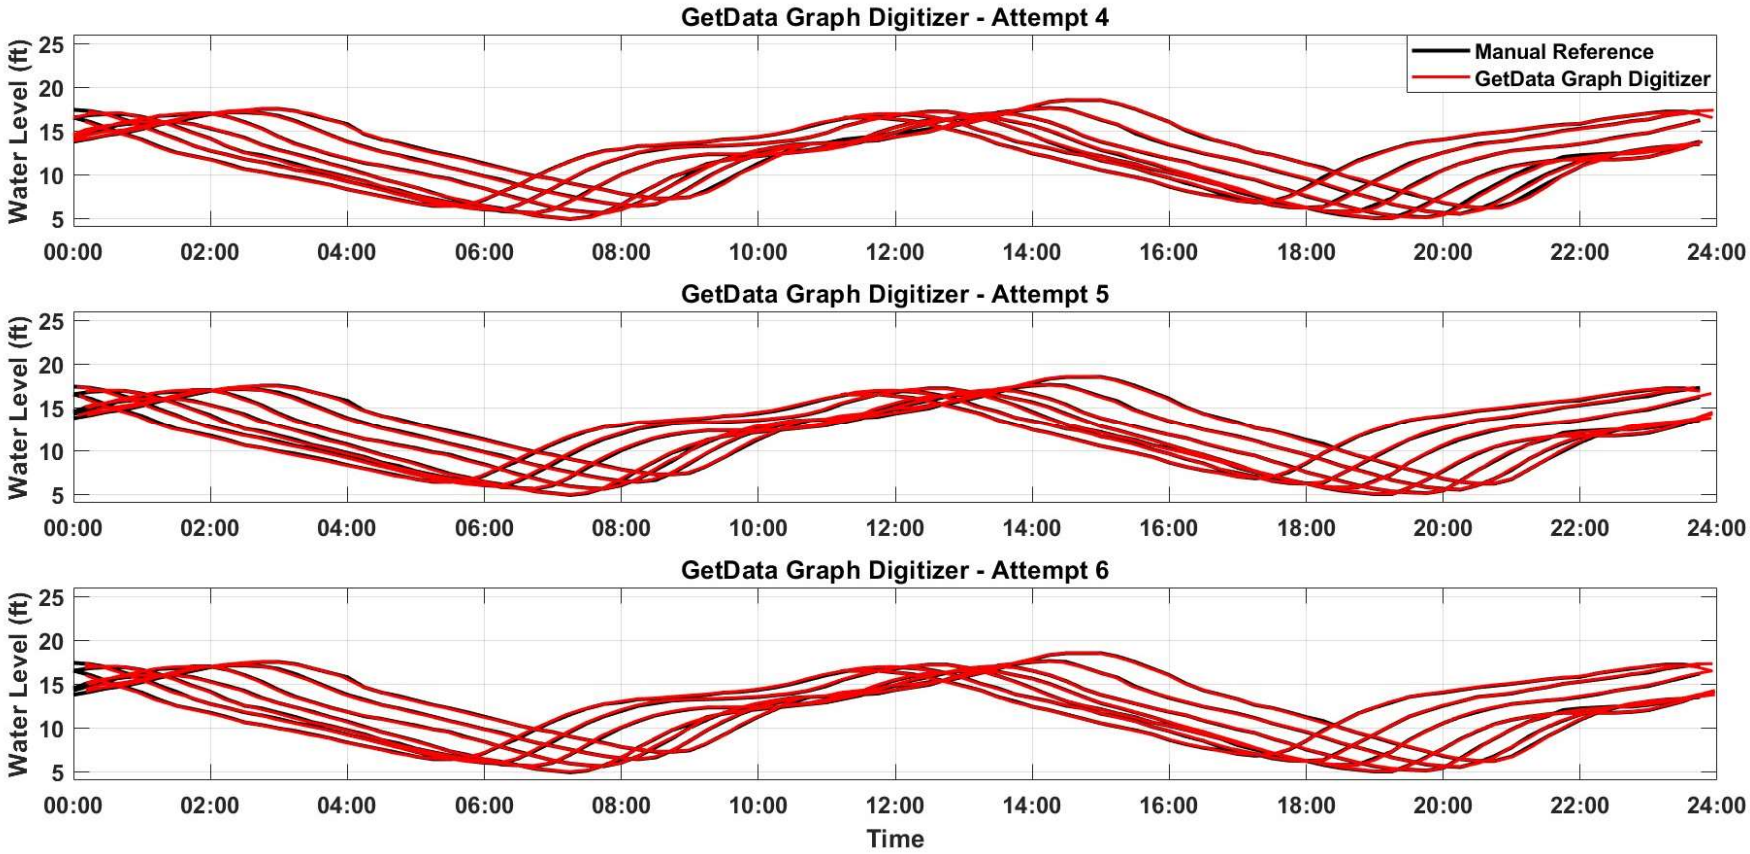

Figure 14. *Attempts 4-6 of using GetData Graph Digitizer to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

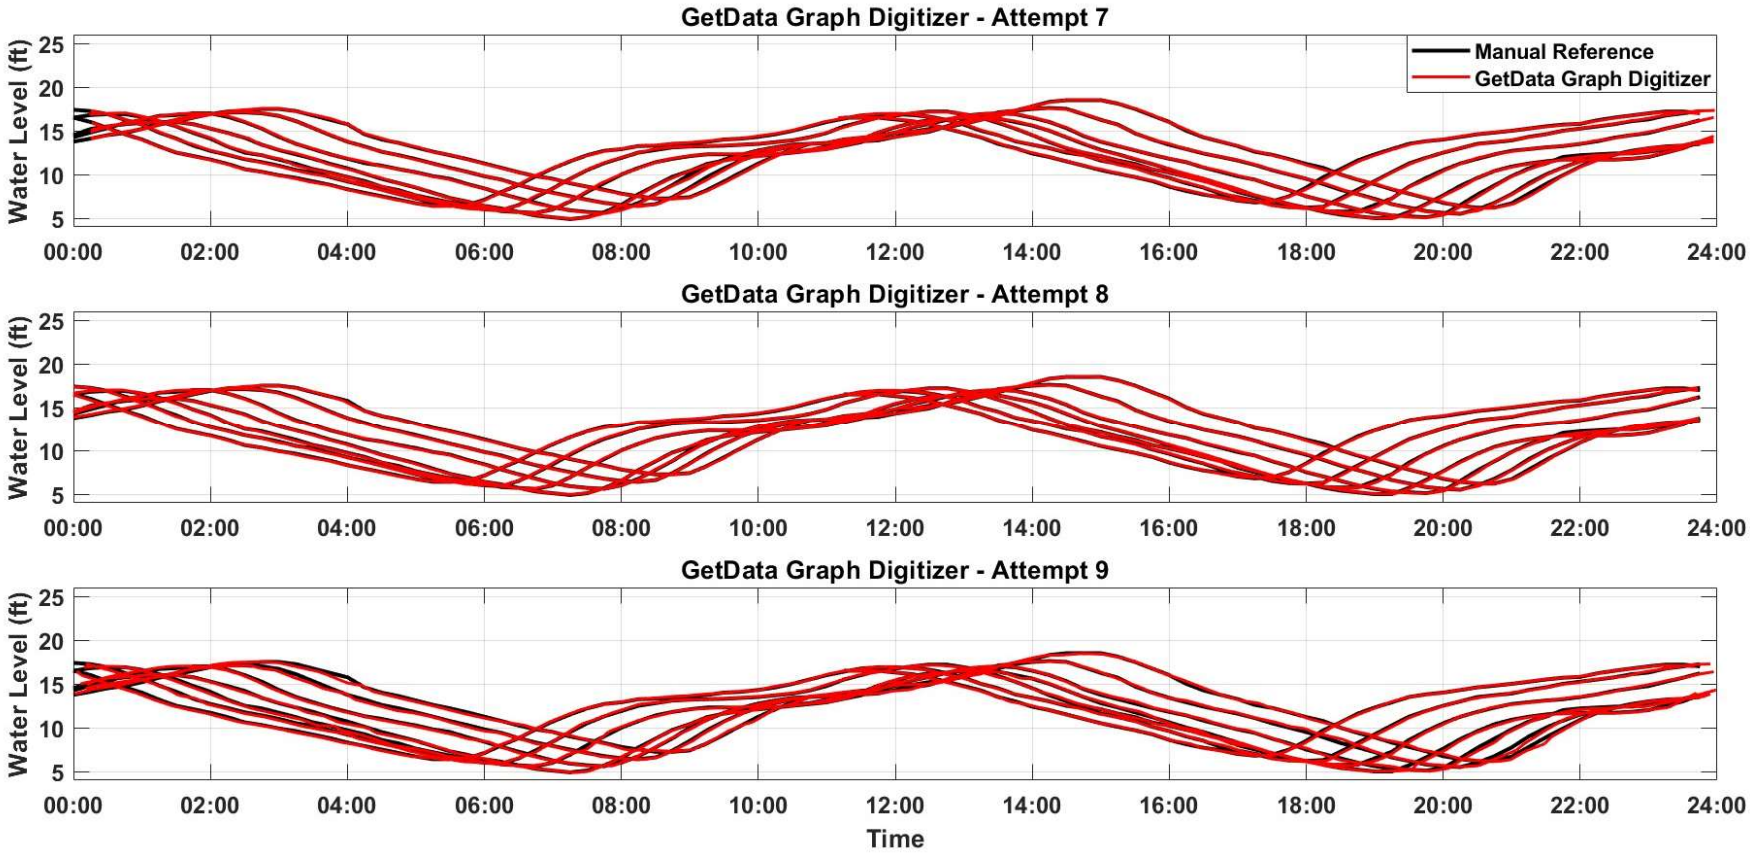

Figure 15. *Attempts 7-9 of using GetData Graph Digitizer to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

2.5. GRABIT

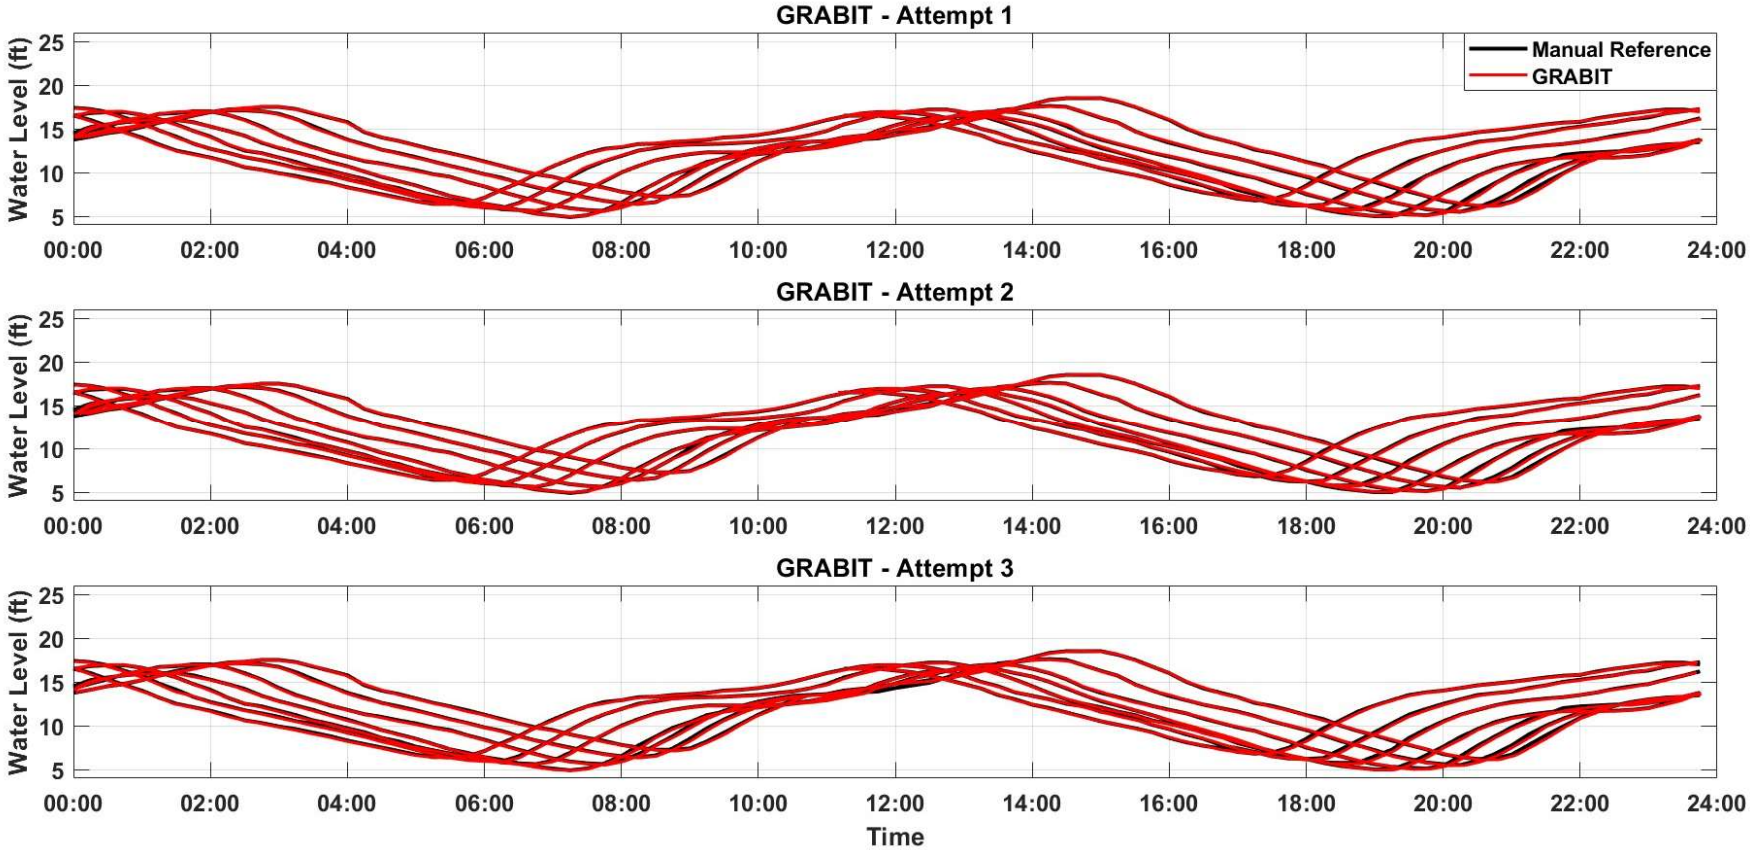

Figure 16. *Attempts 1-3 of using GRABIT to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

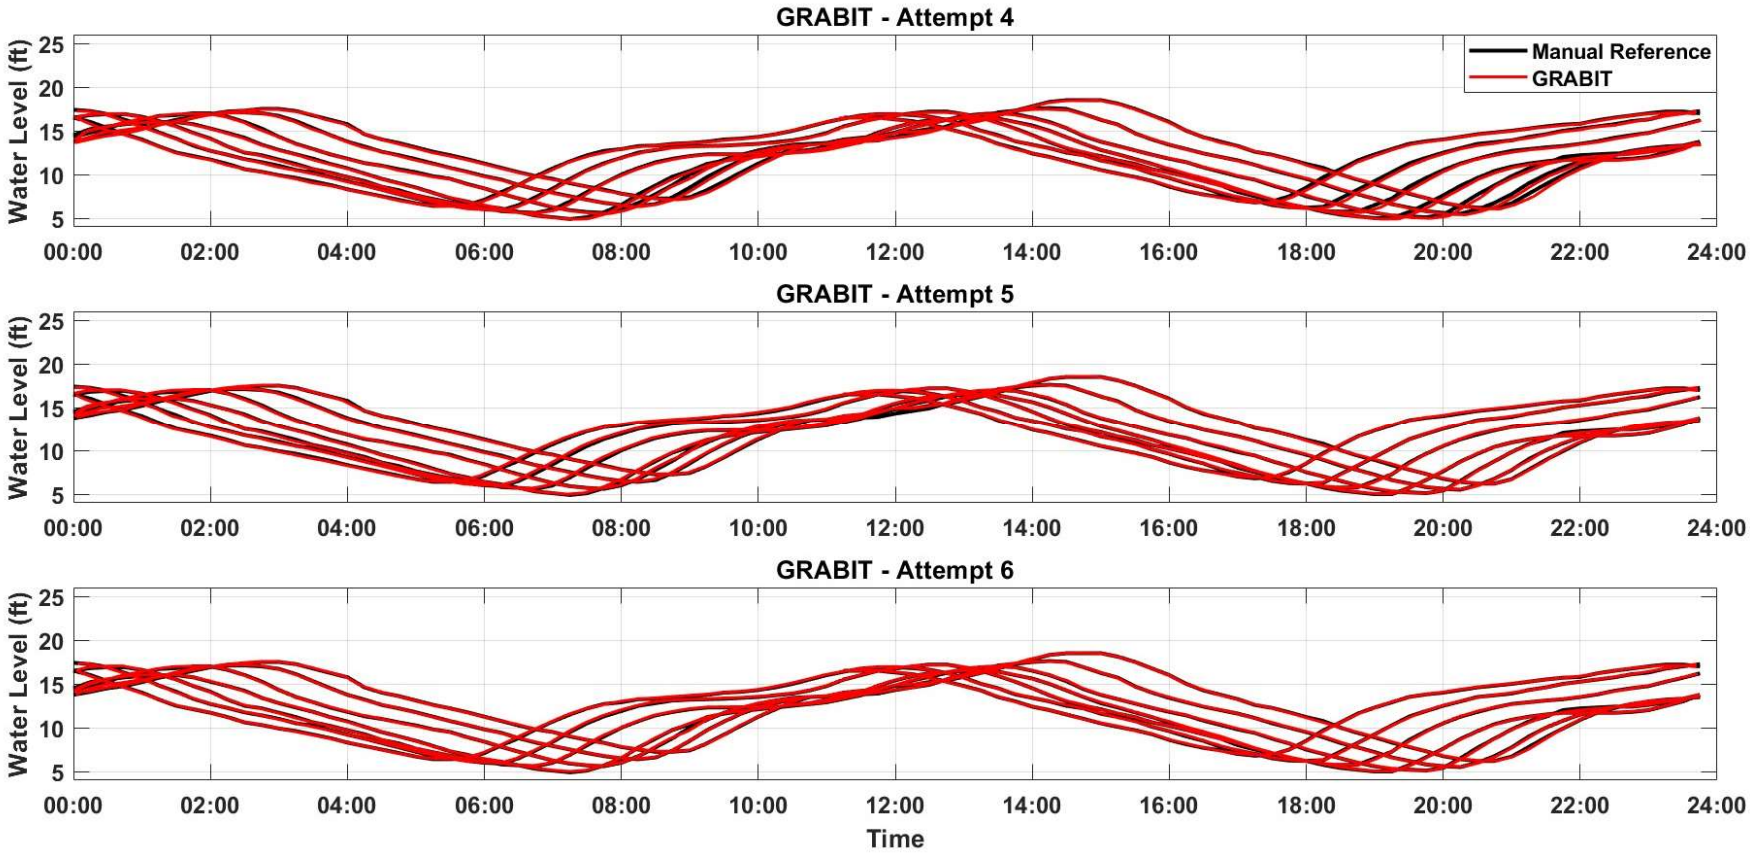

Figure 17 *Attempts 4-5 of using GRABIT to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

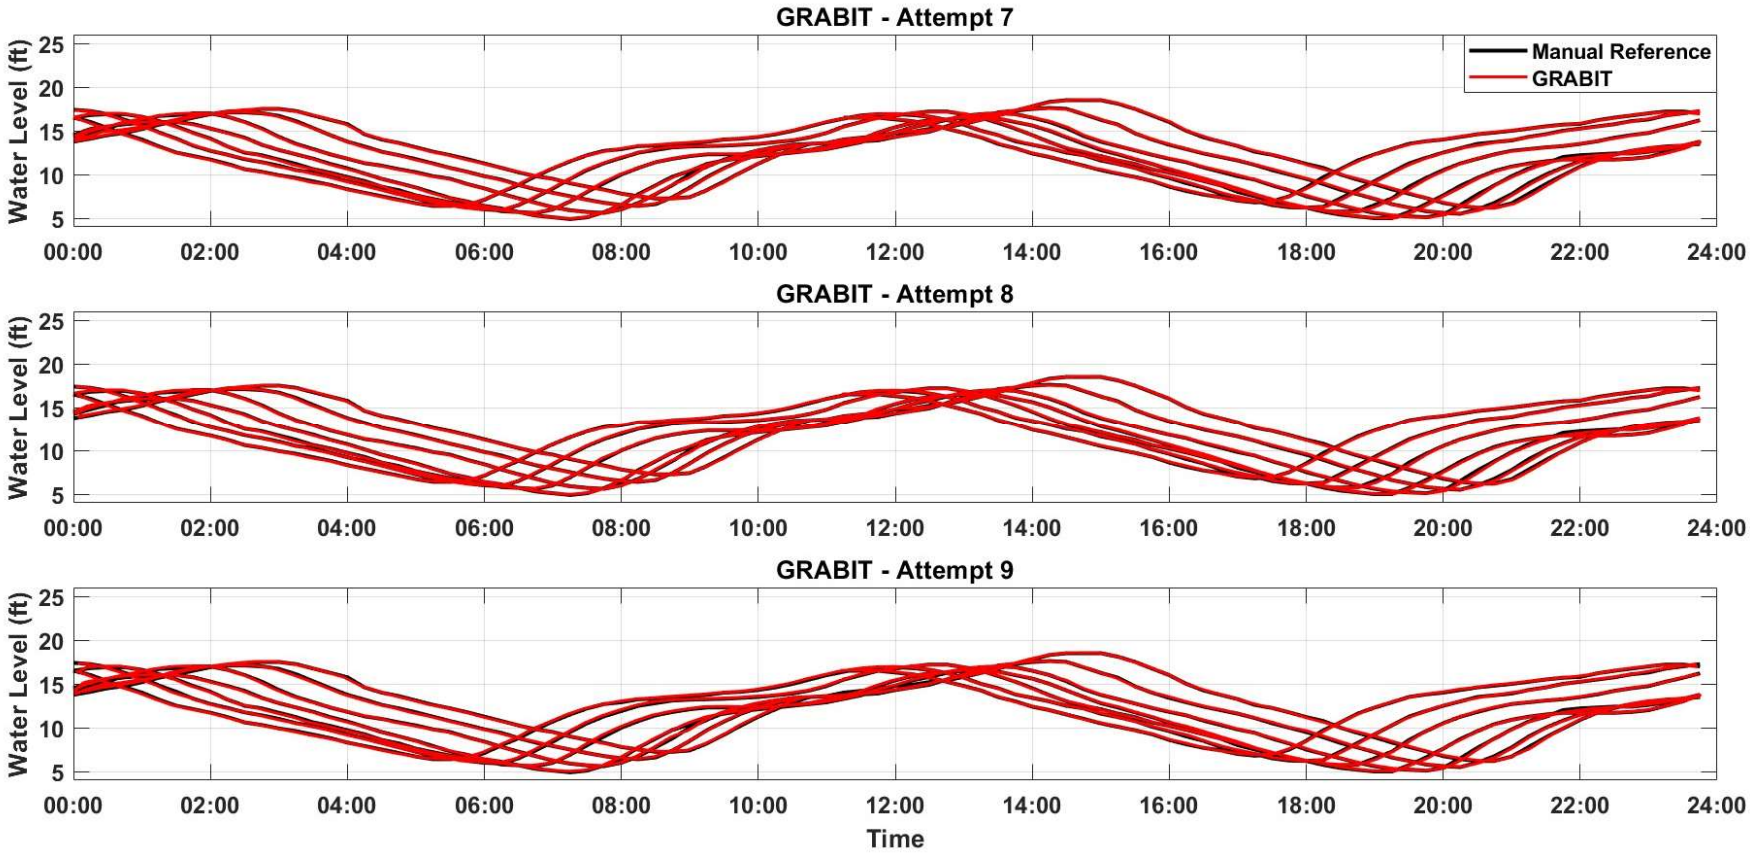

Figure 18. *Attempts 7-9 of using GRABIT to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

2.6. GraphGrabber

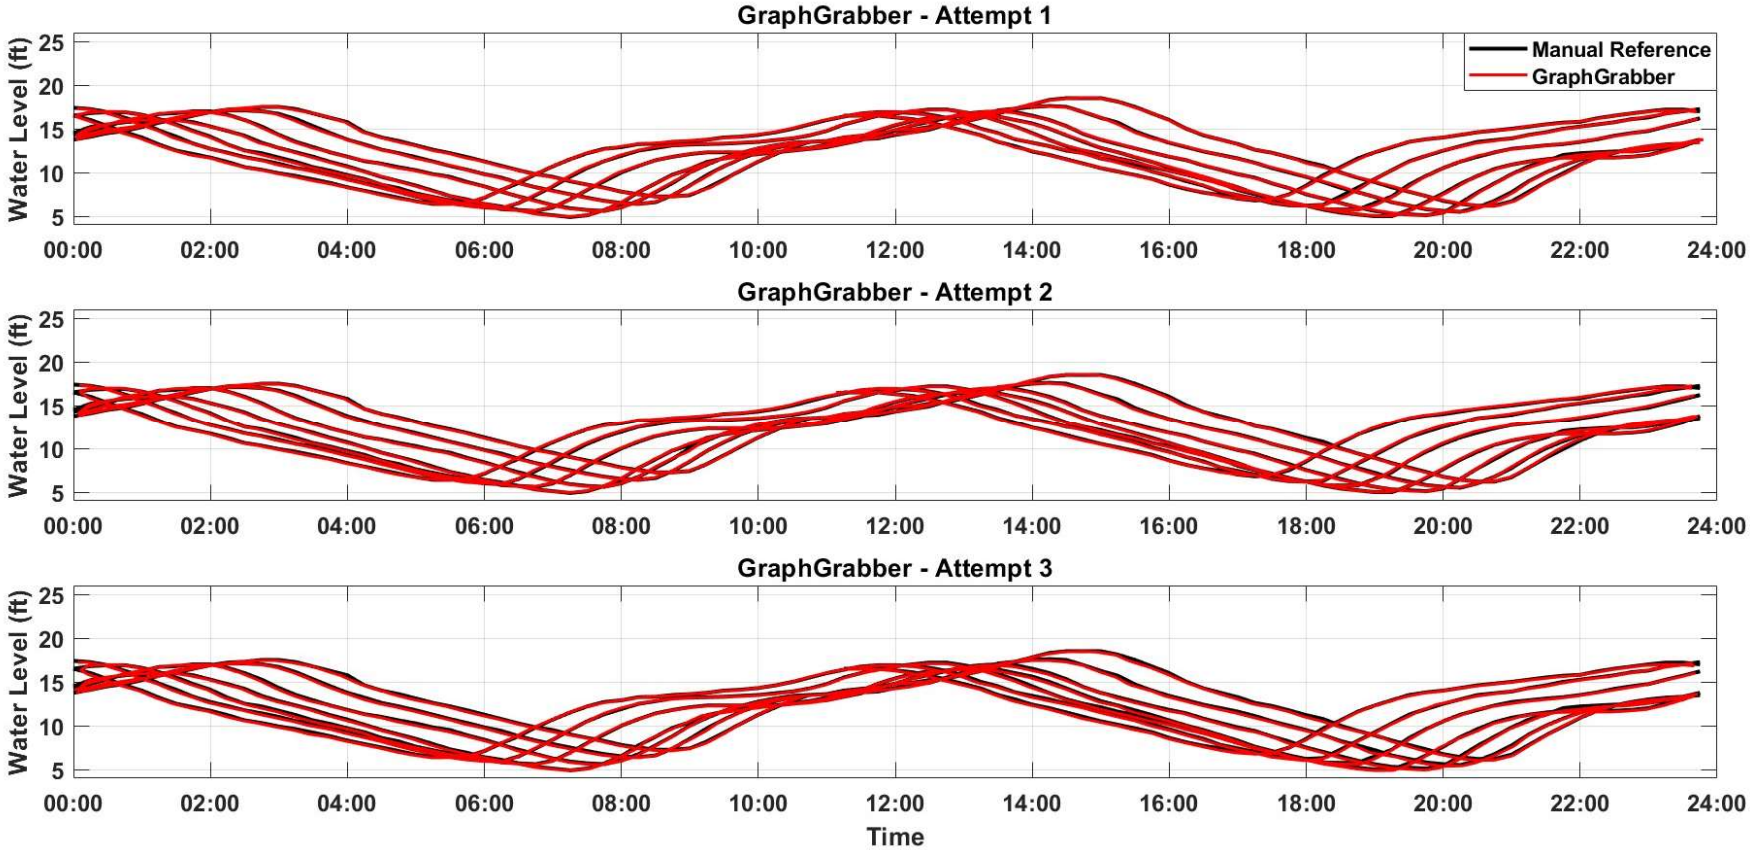

Figure 19. *Attempts 1-3 of using GraphGrabber to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

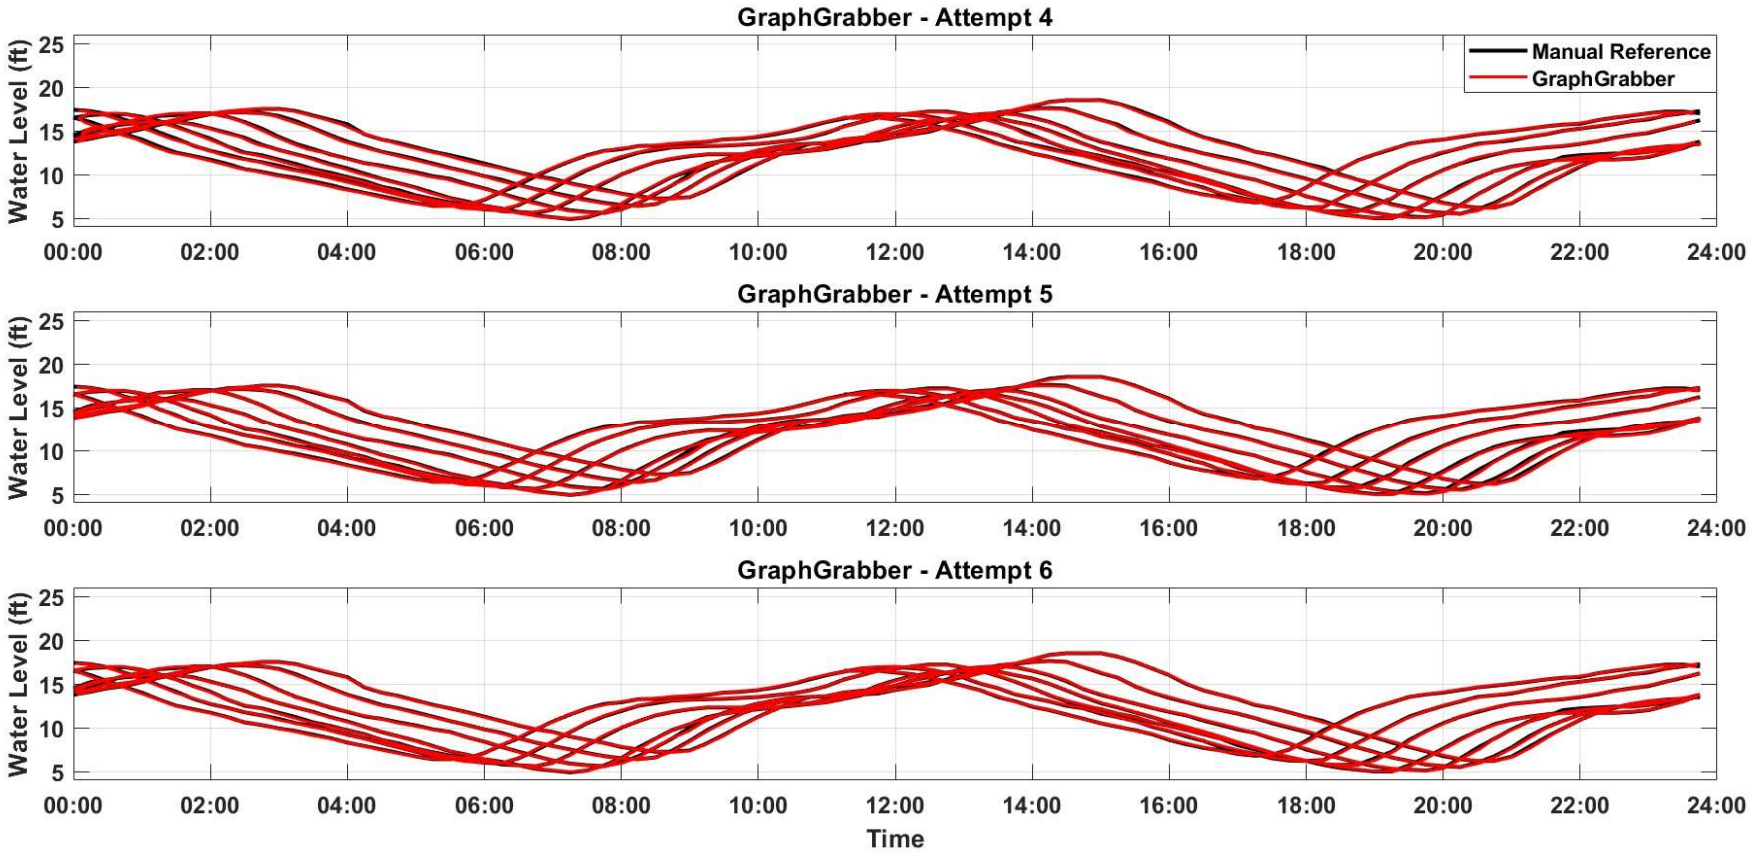

Figure 20. *Attempts 4-6 of using GraphGrabber to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

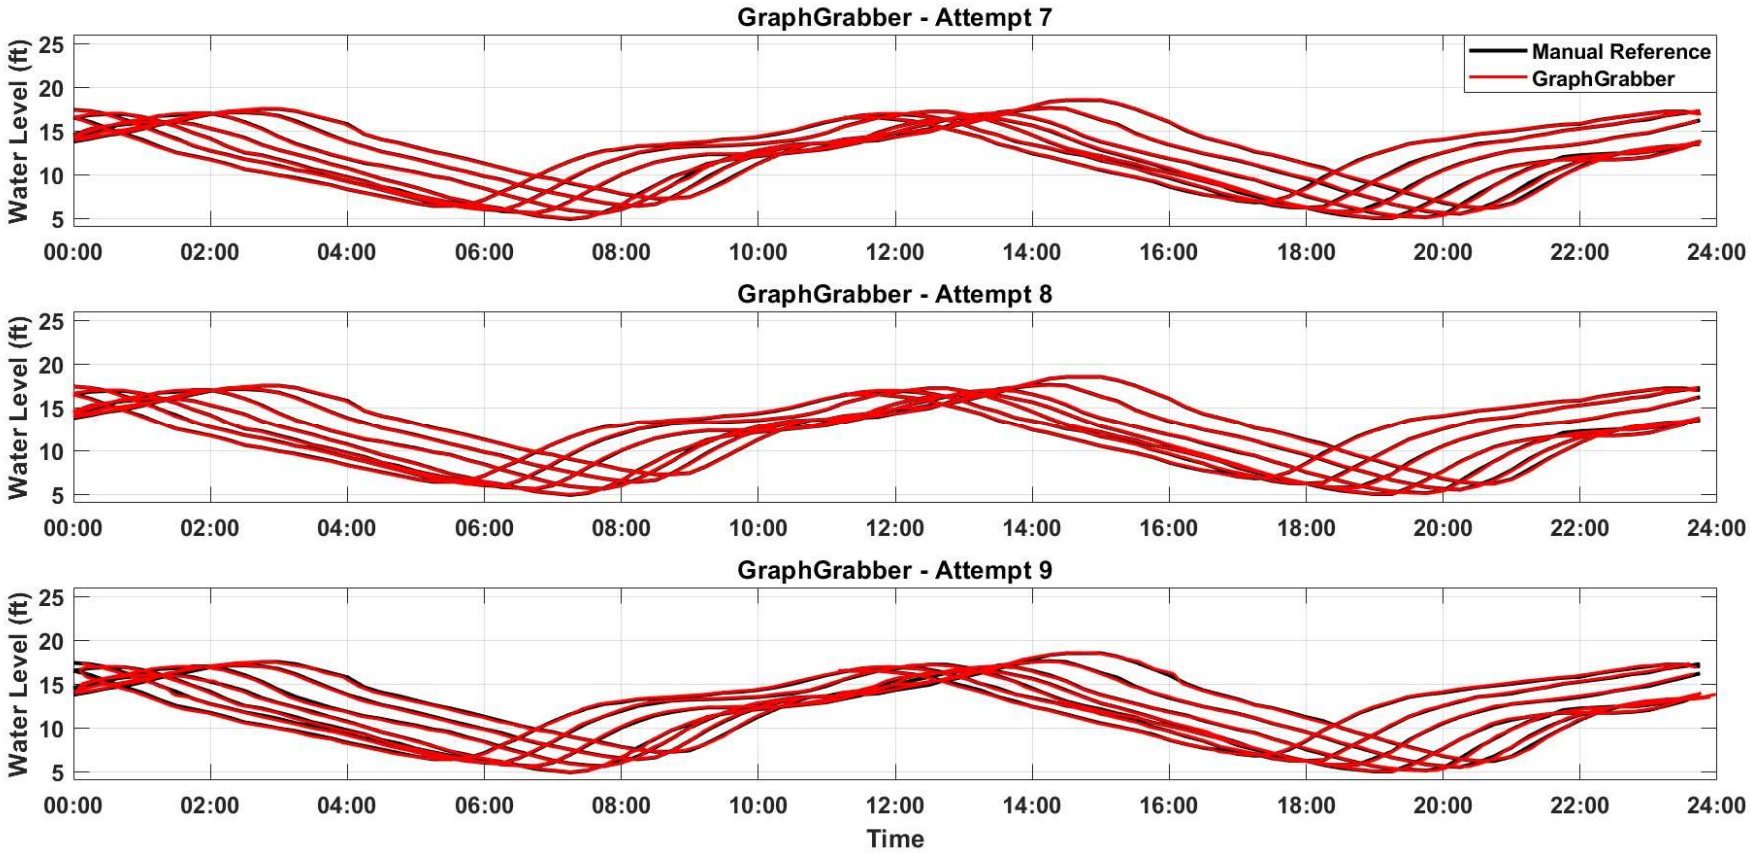

Figure 21. *Attempts 7-9 of using GraphGrabber to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

2.7. OriginLab

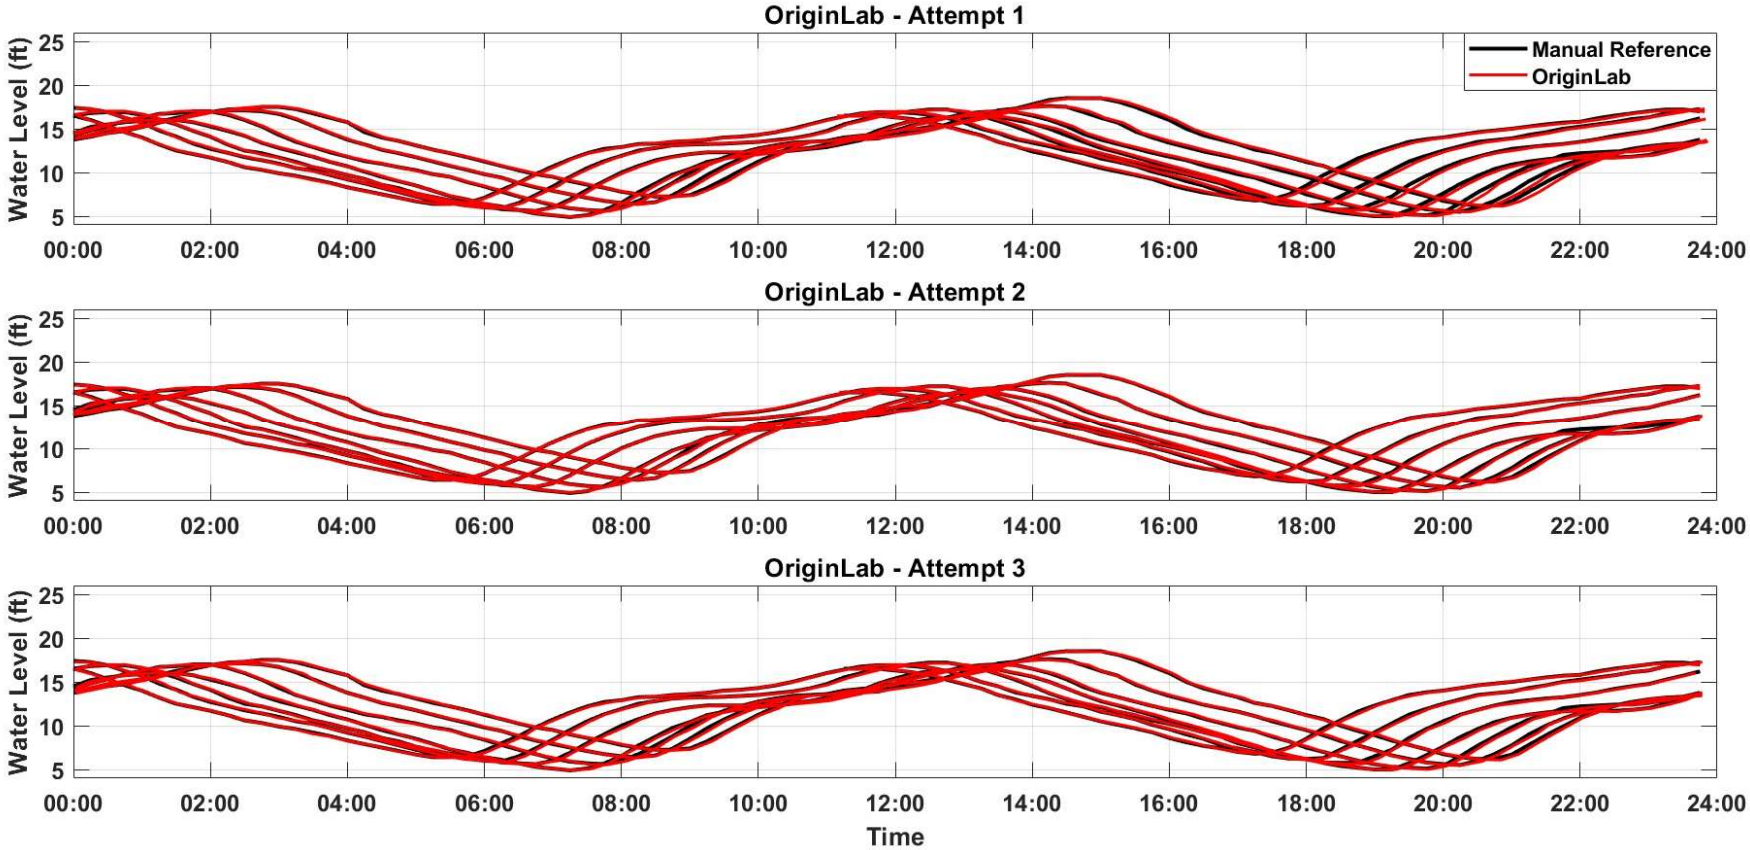

Figure 22. *Attempts 1-3 of using OriginLab to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

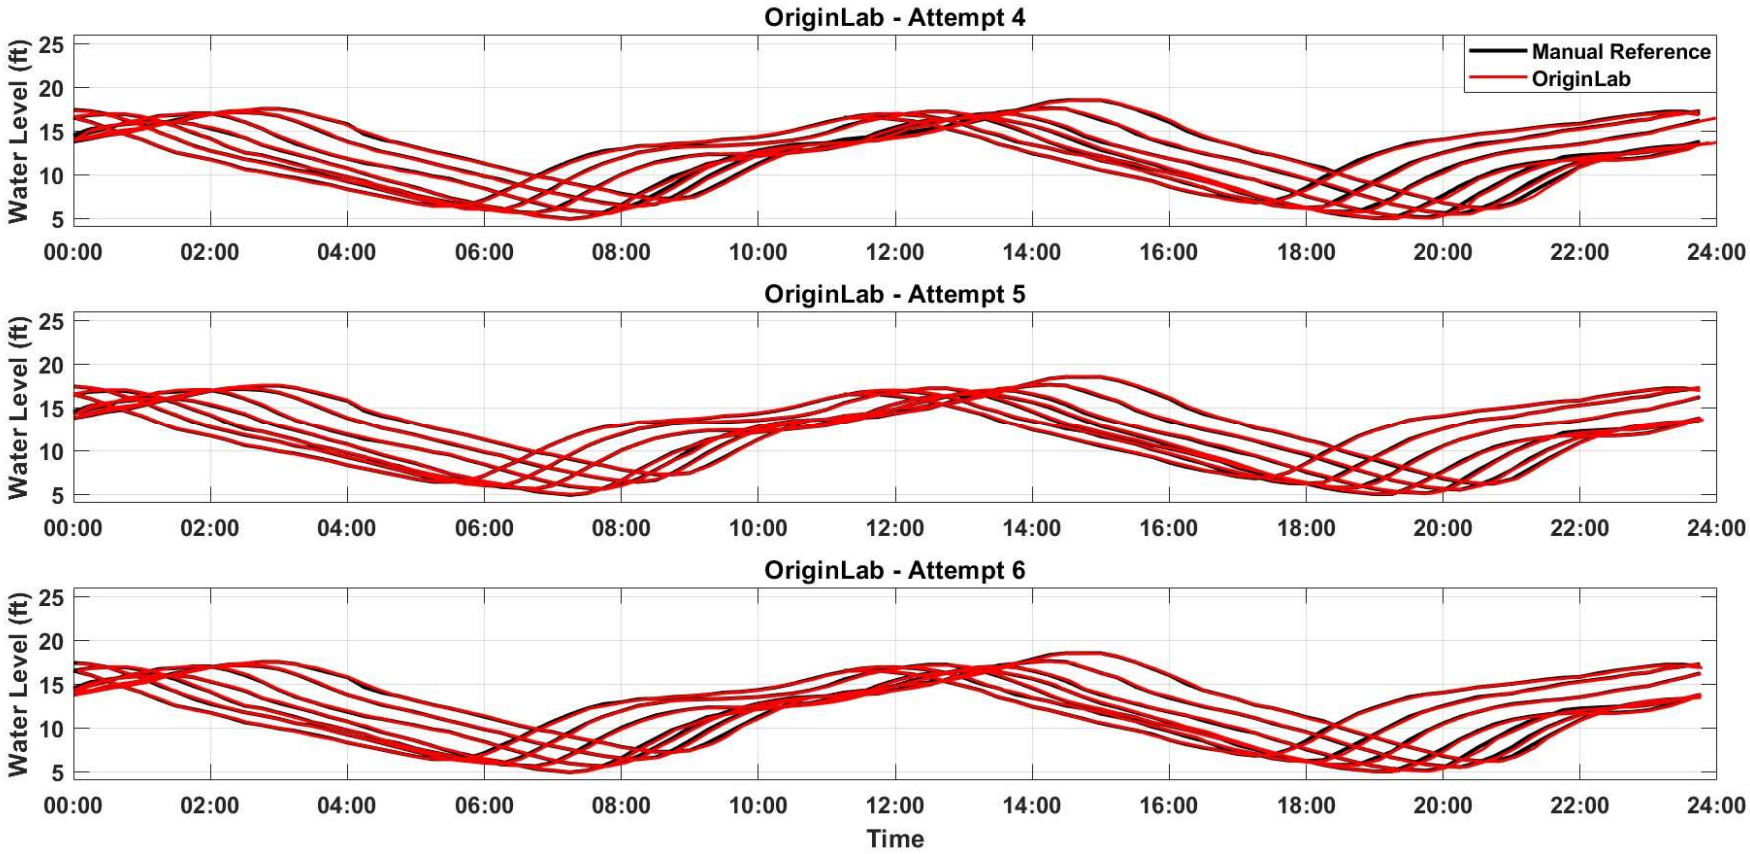

Figure 23. *Attempts 4-6 of using OriginLab to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

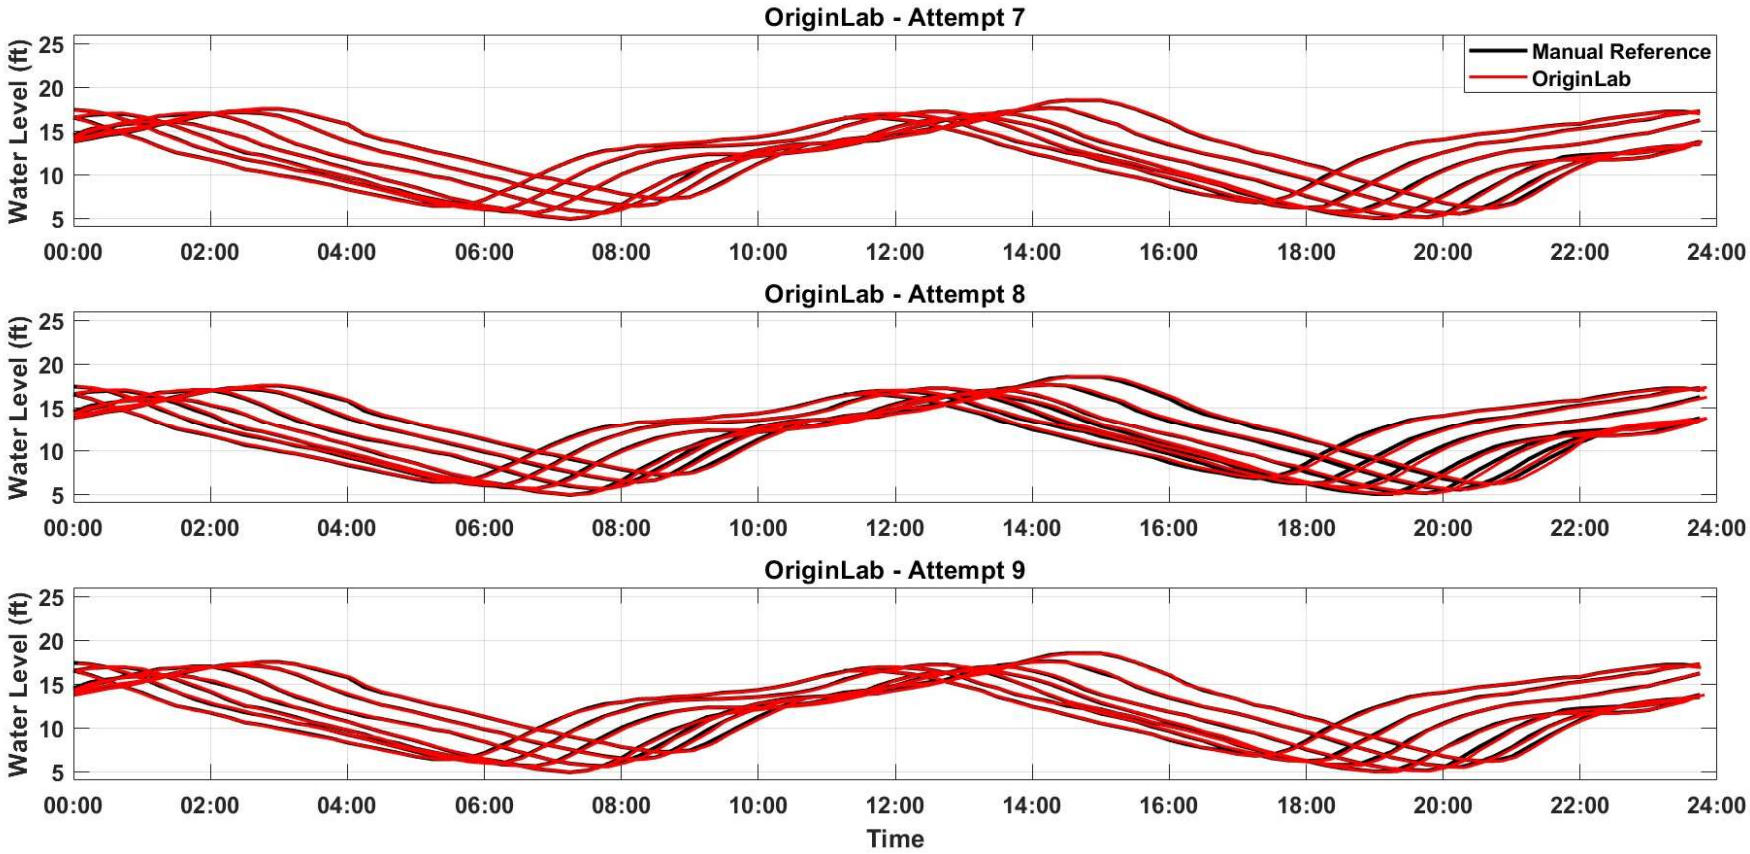

Figure 24. *Attempts 7-9 of using OriginLab to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

3. Trace Tools Comparison Plots

3.1. DigitGraph

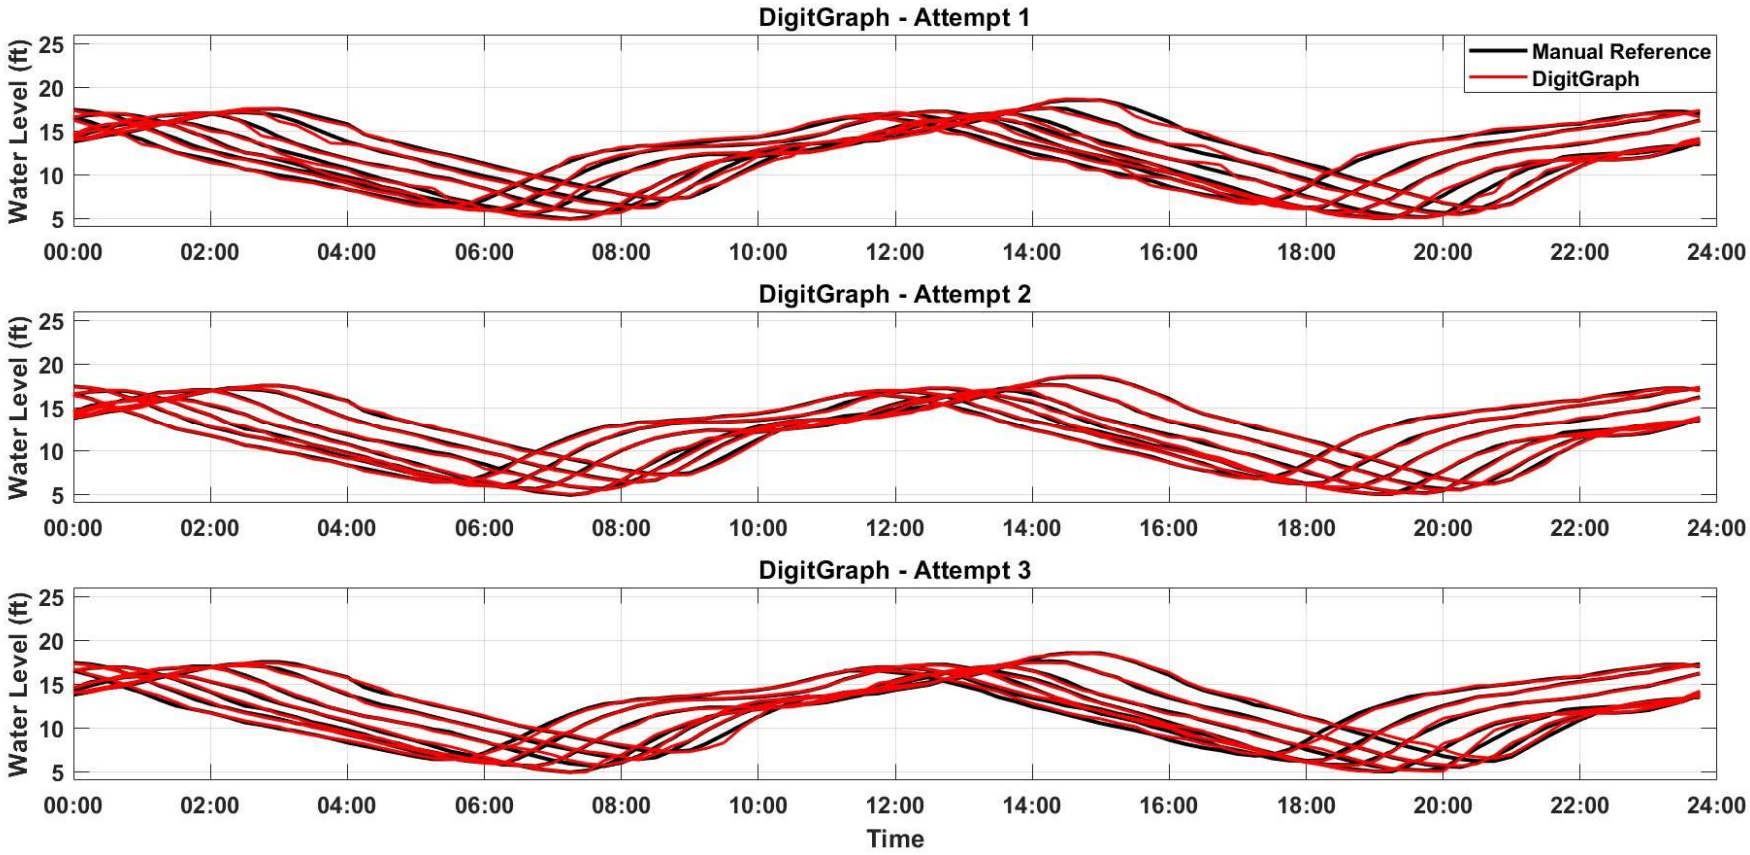

Figure 25. *Attempts 1-3 of using DigitGraph to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

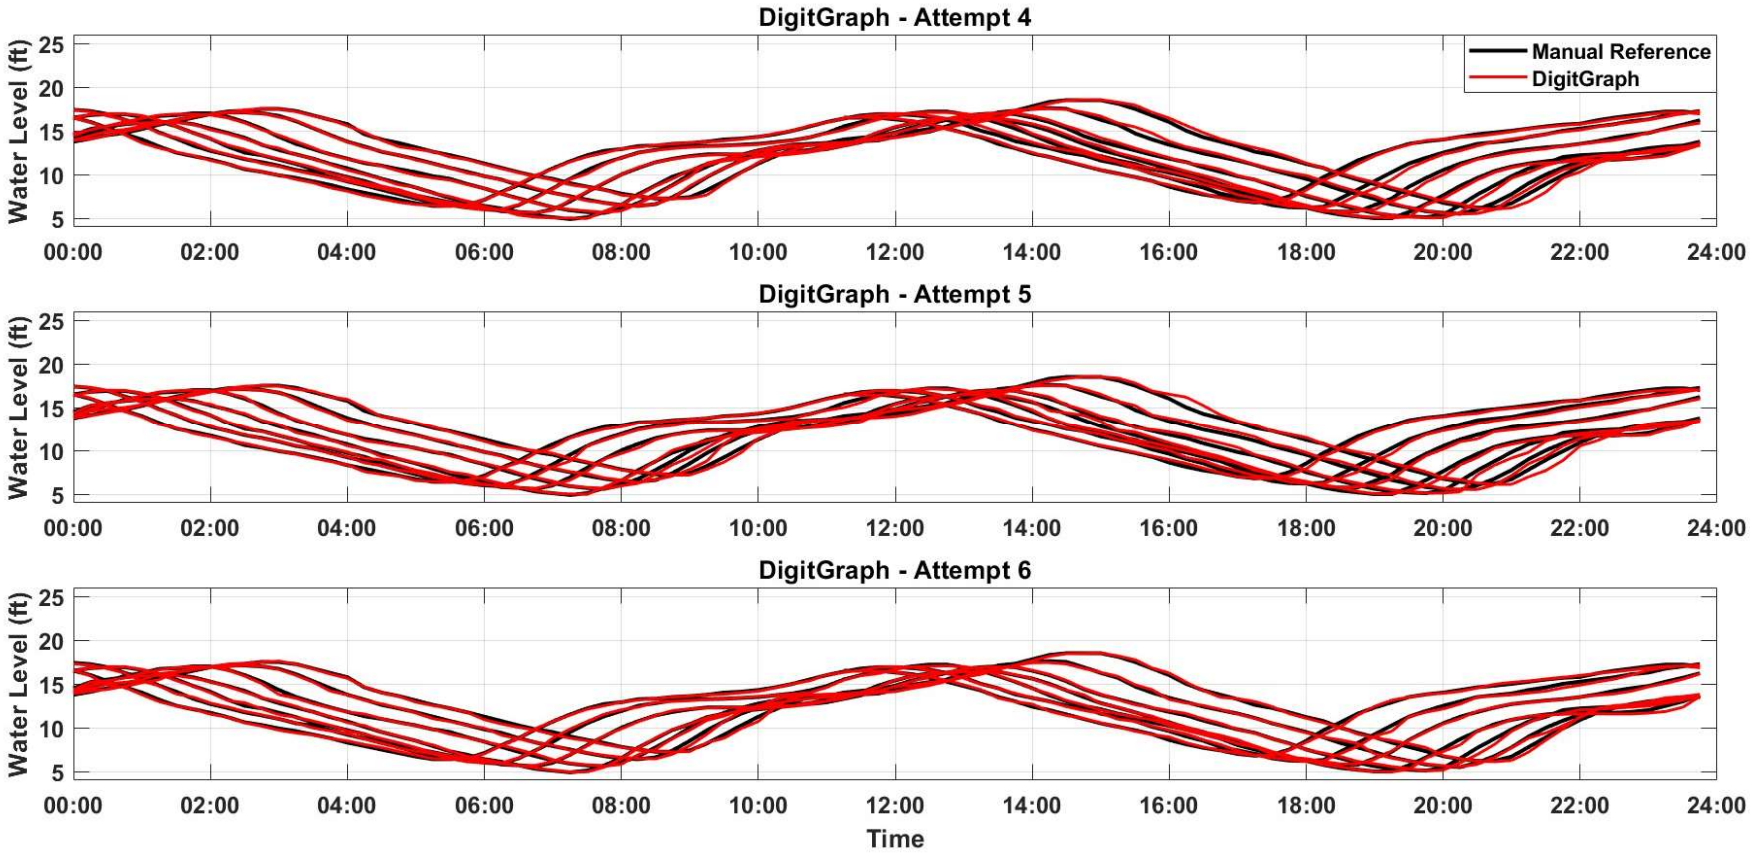

Figure 26. *Attempts 4-6 of using DigitGraph to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

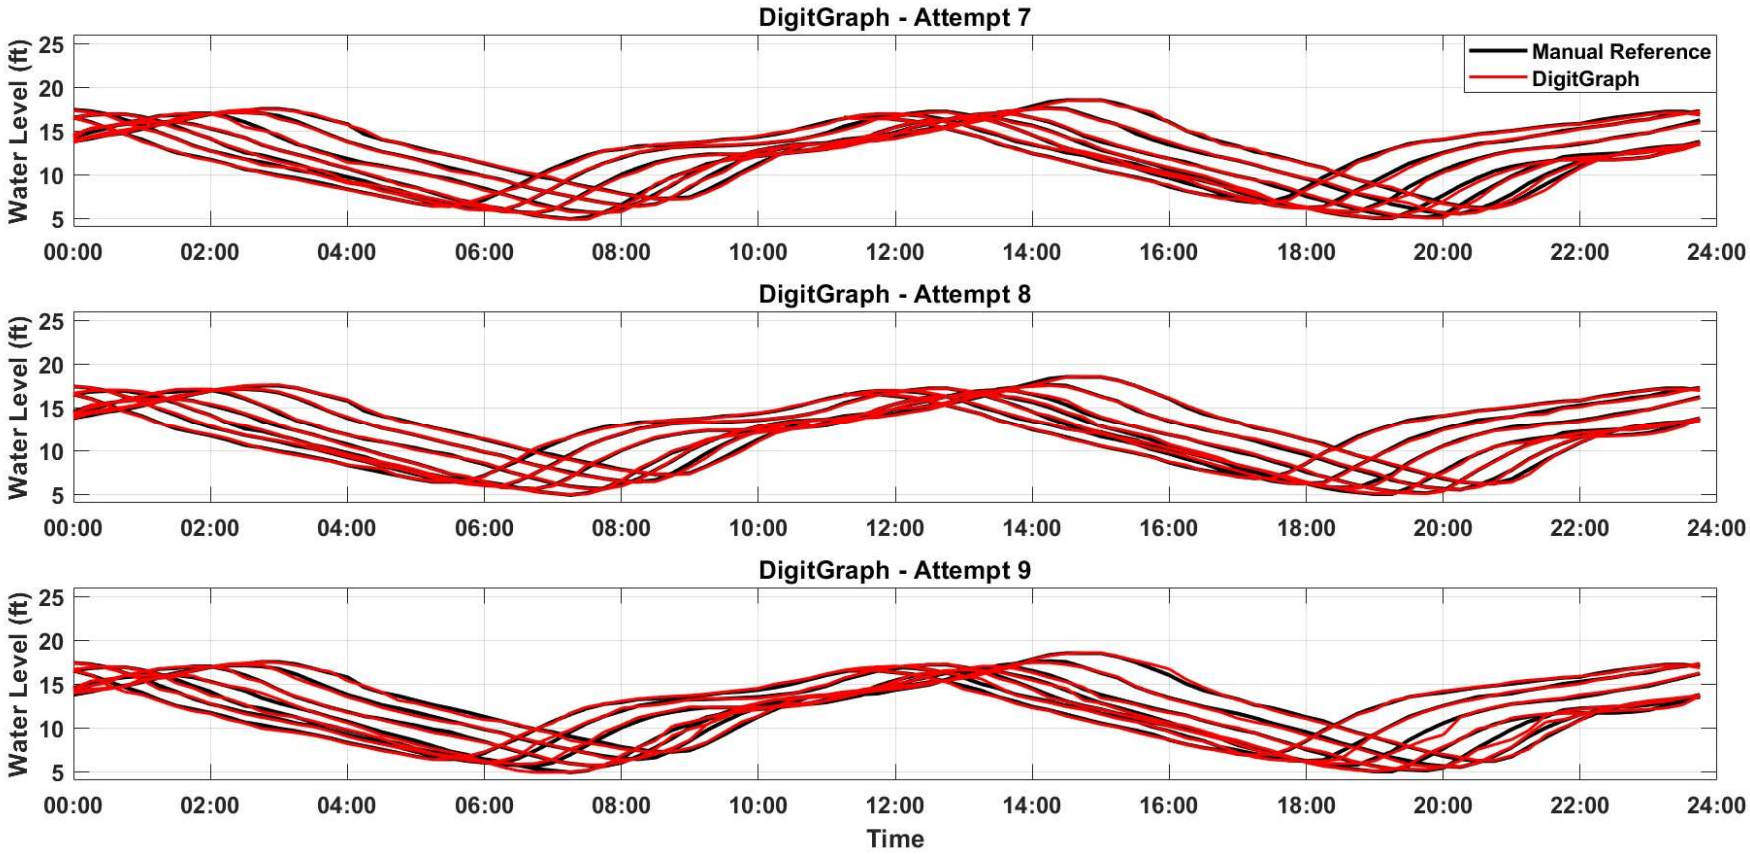

Figure 27. *Attempts 7-9 of using DigitGraph to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

3.2. NUNIEAU

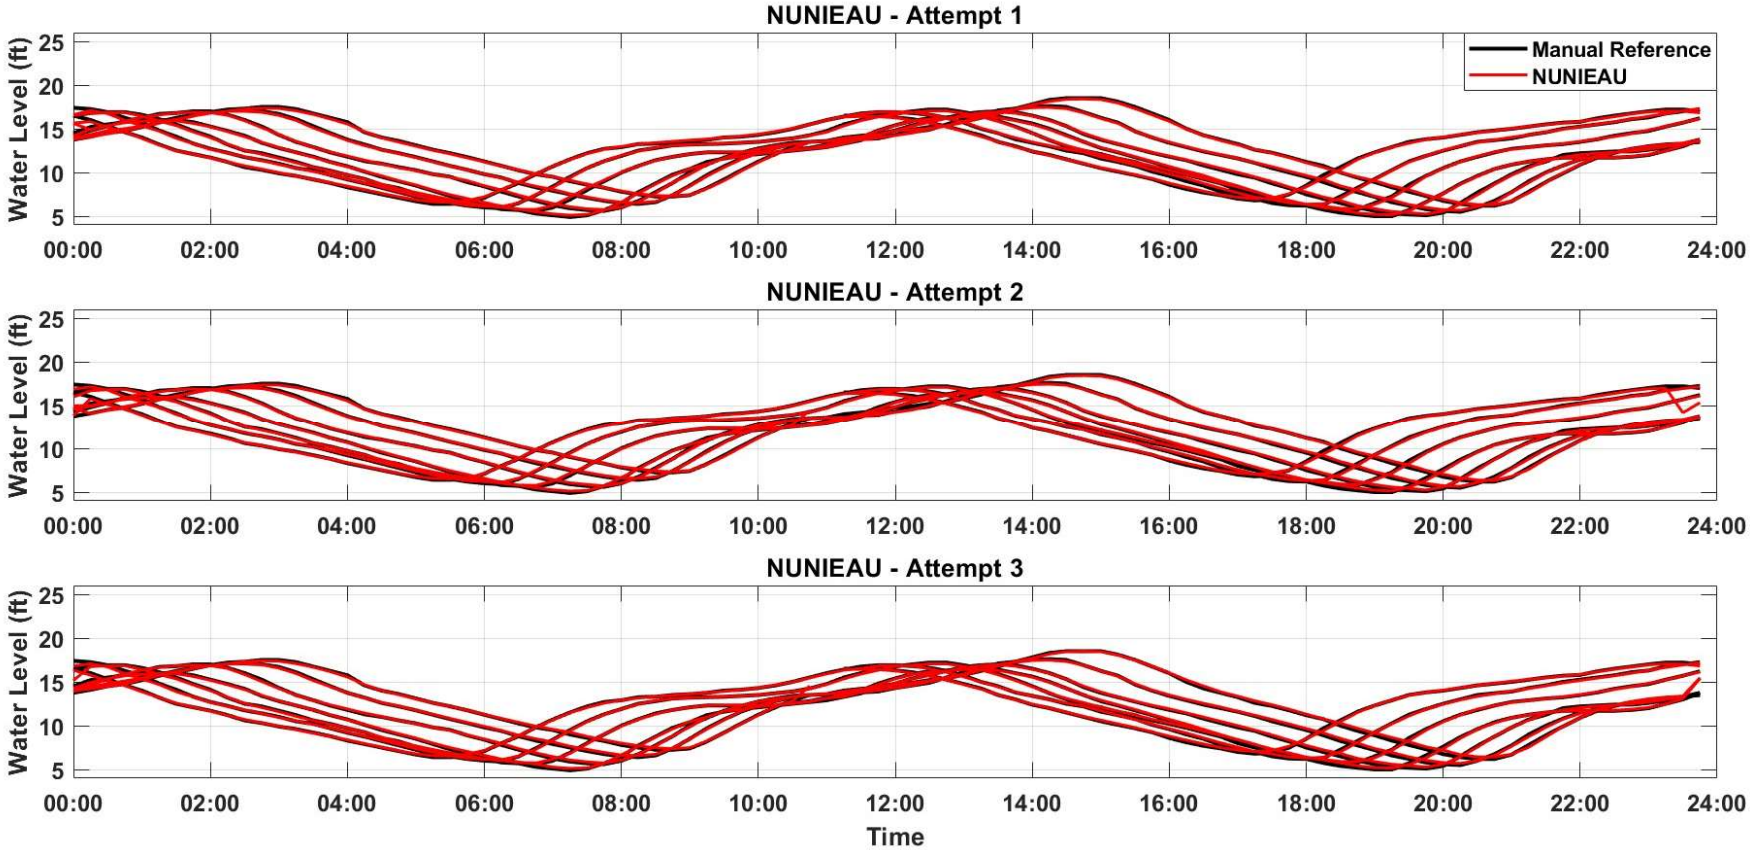

Figure 28. *Attempts 1-3 of using NUNIEAU to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

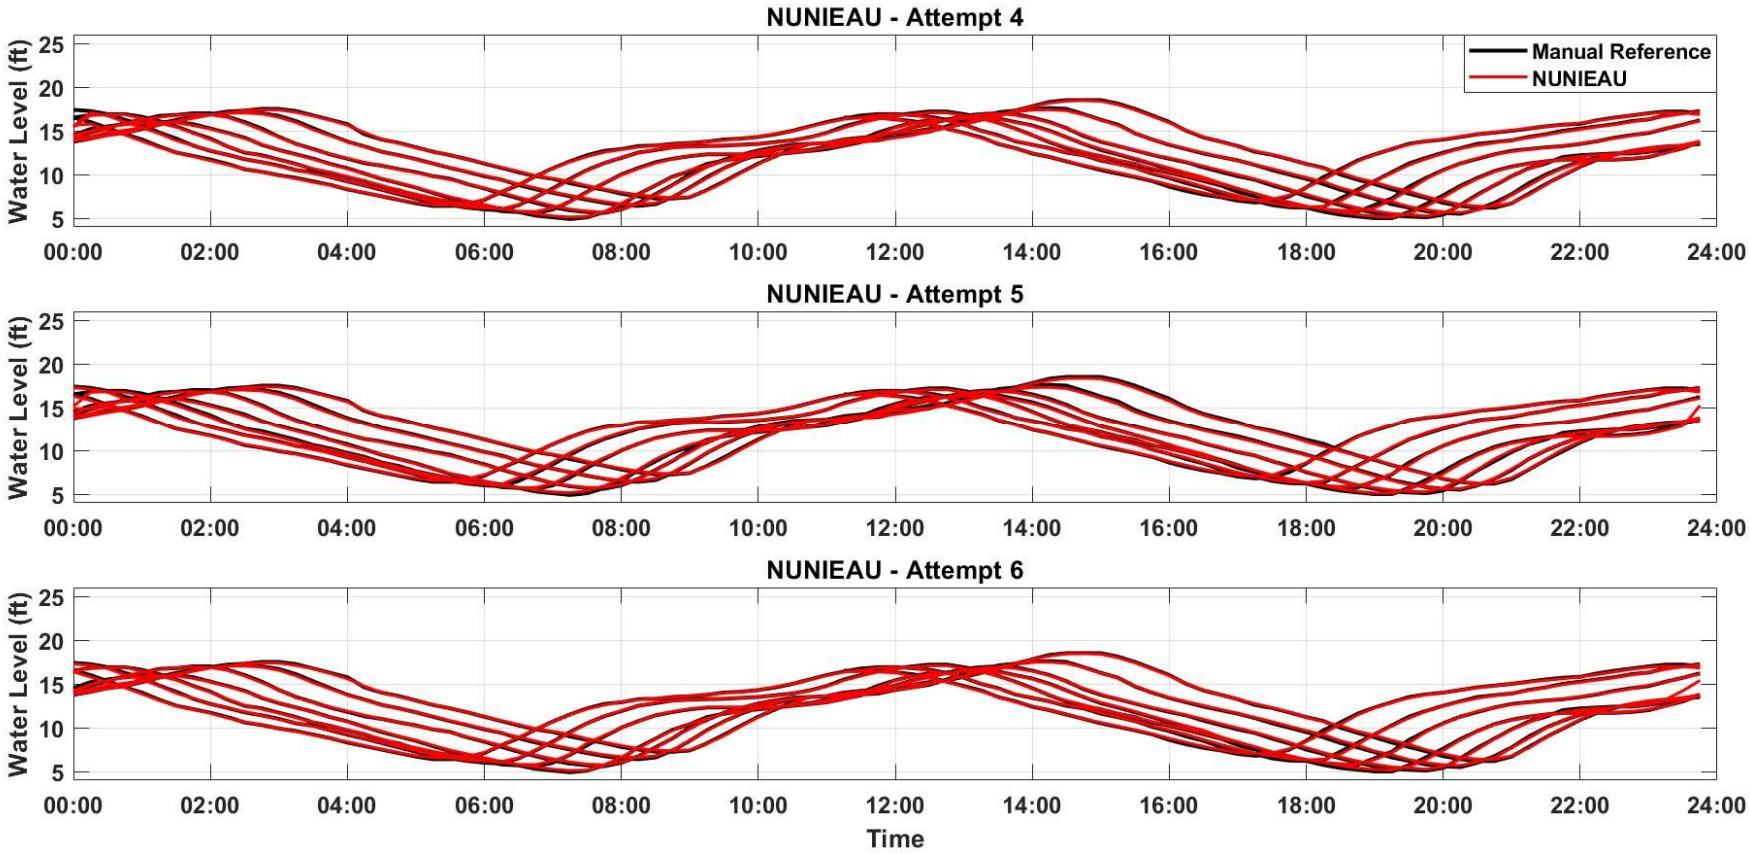

Figure 29. *Attempts 4-6 of using NUNIEAU to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

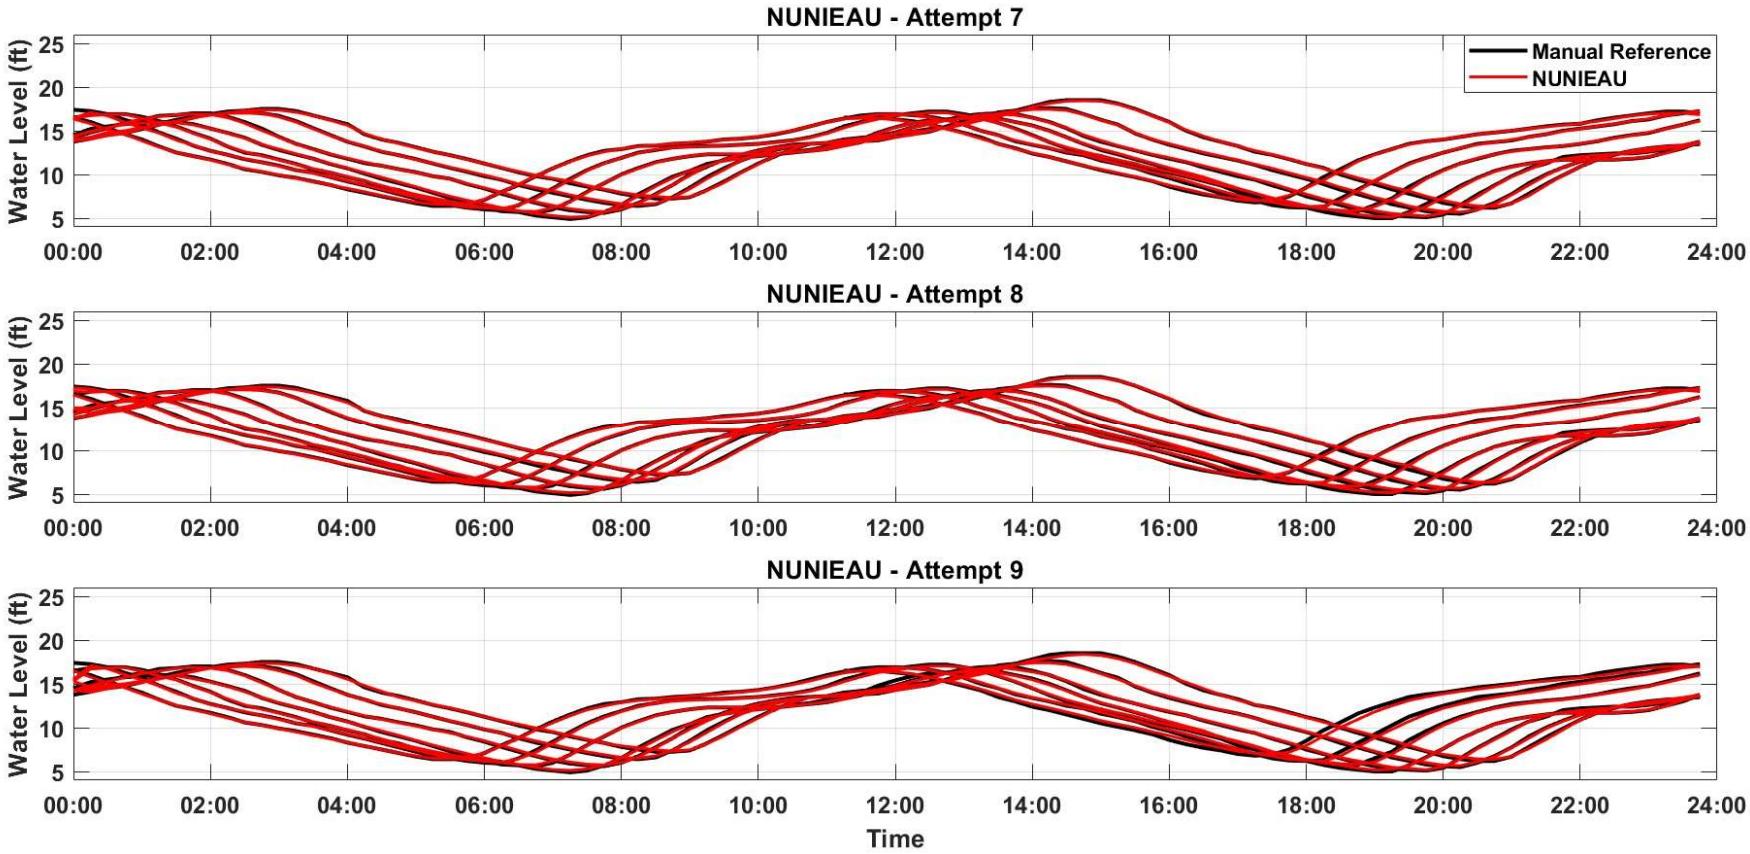

Figure 30. *Attempts 7-9 of using NUNIEAU to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

3.3. Plot Digitizer

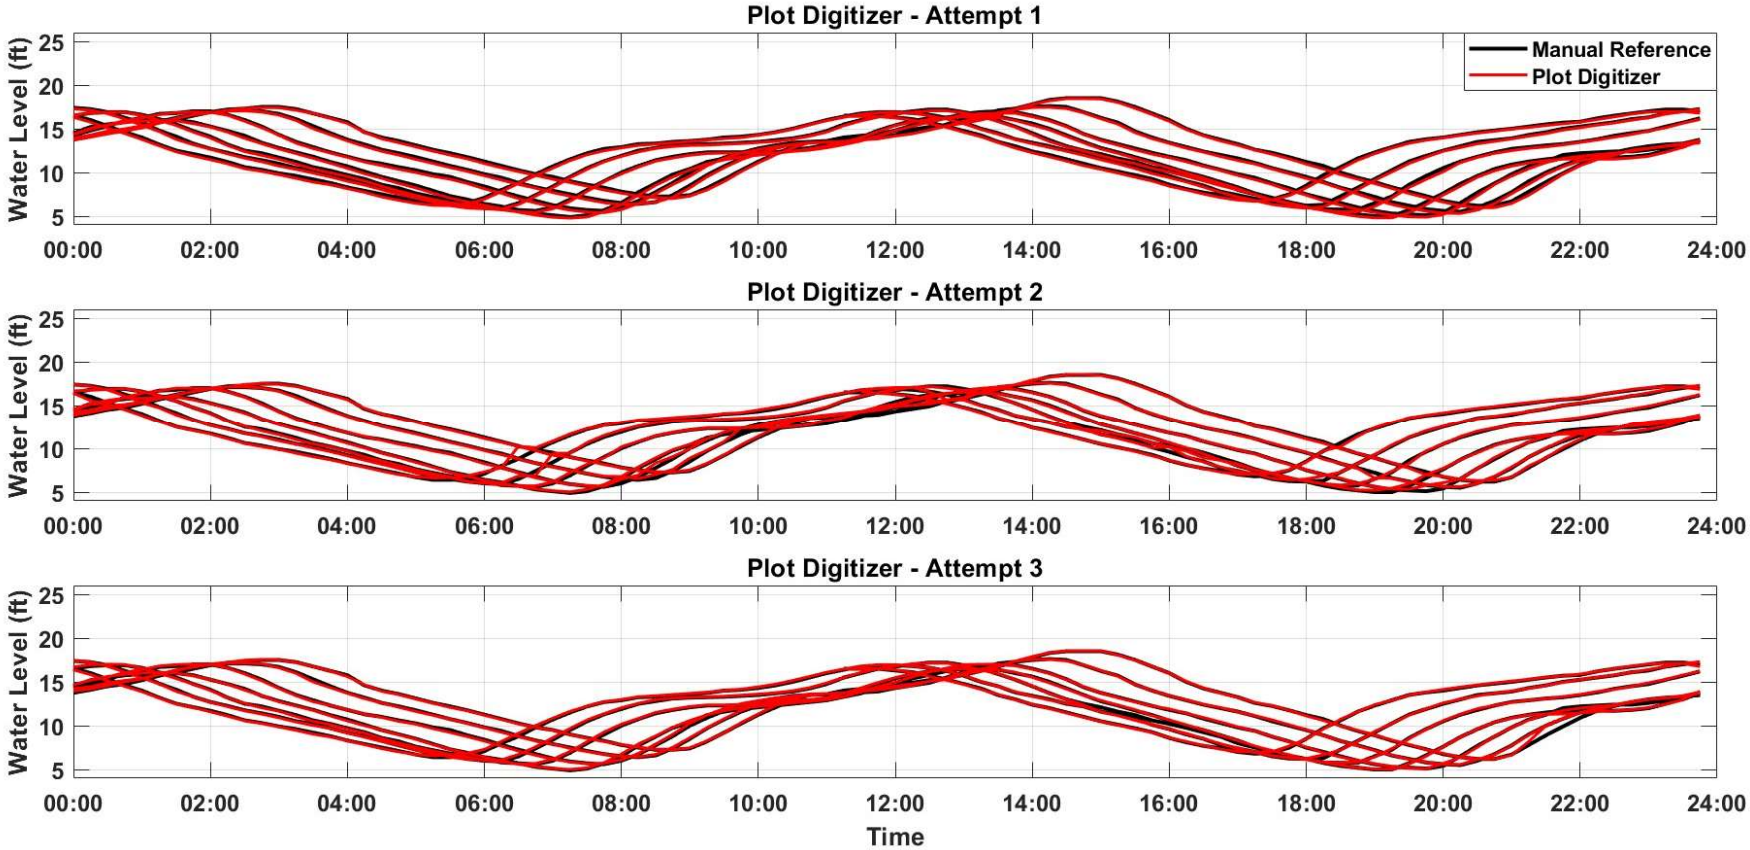

Figure 31. *Attempts 1-3 of using Plot Digitizer to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

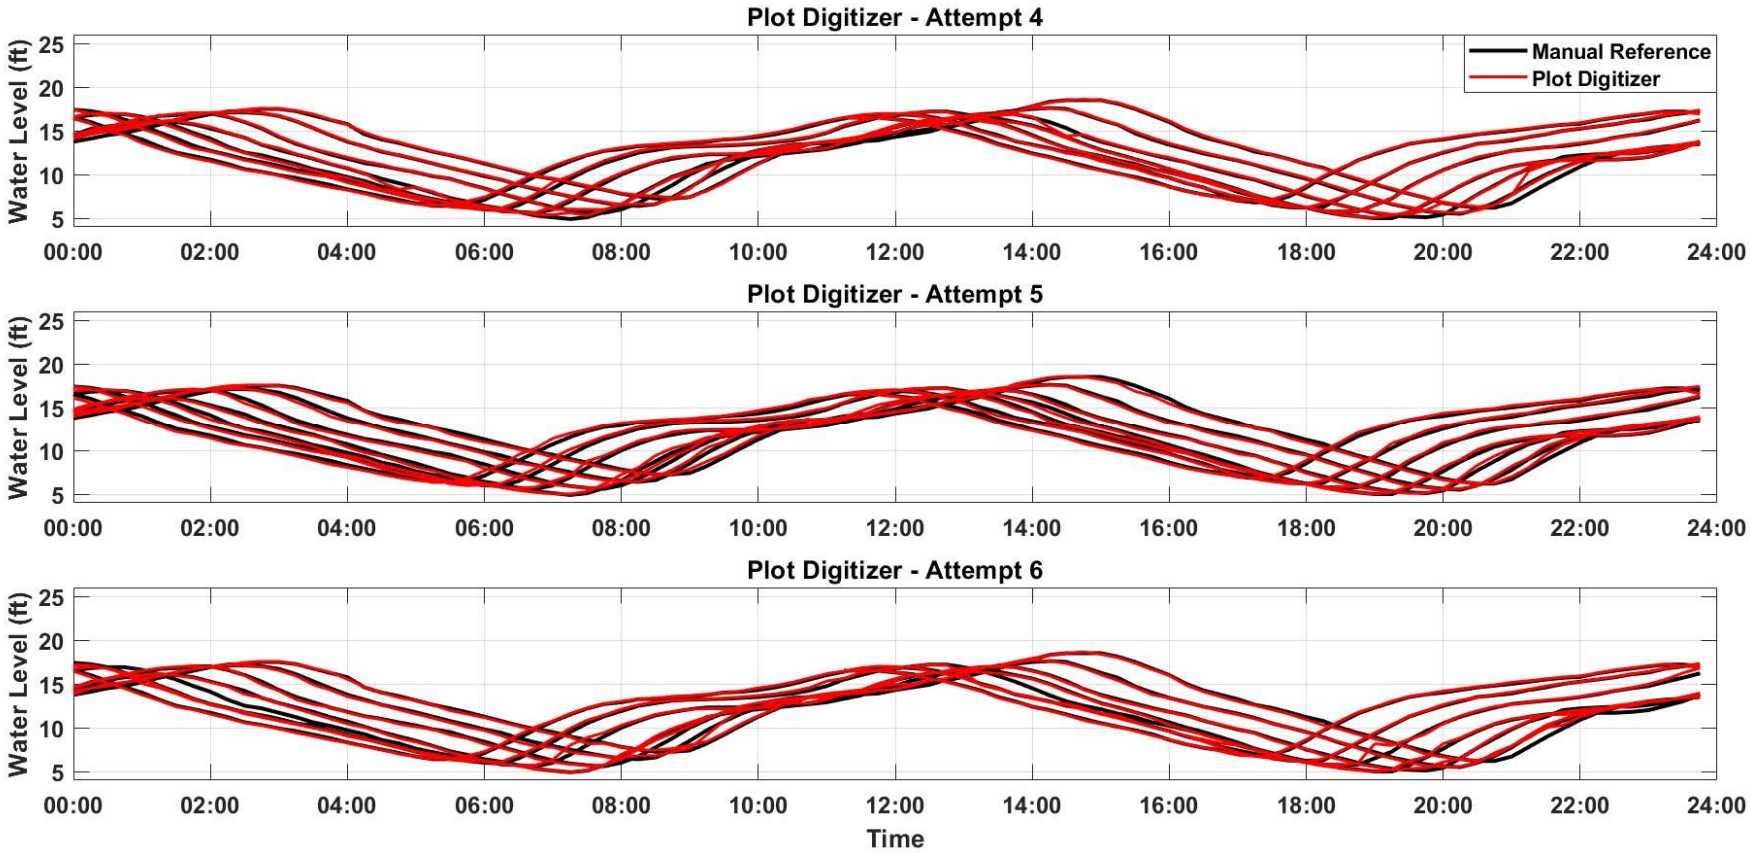

Figure 32. *Attempts 4-6 of using Plot Digitizer to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

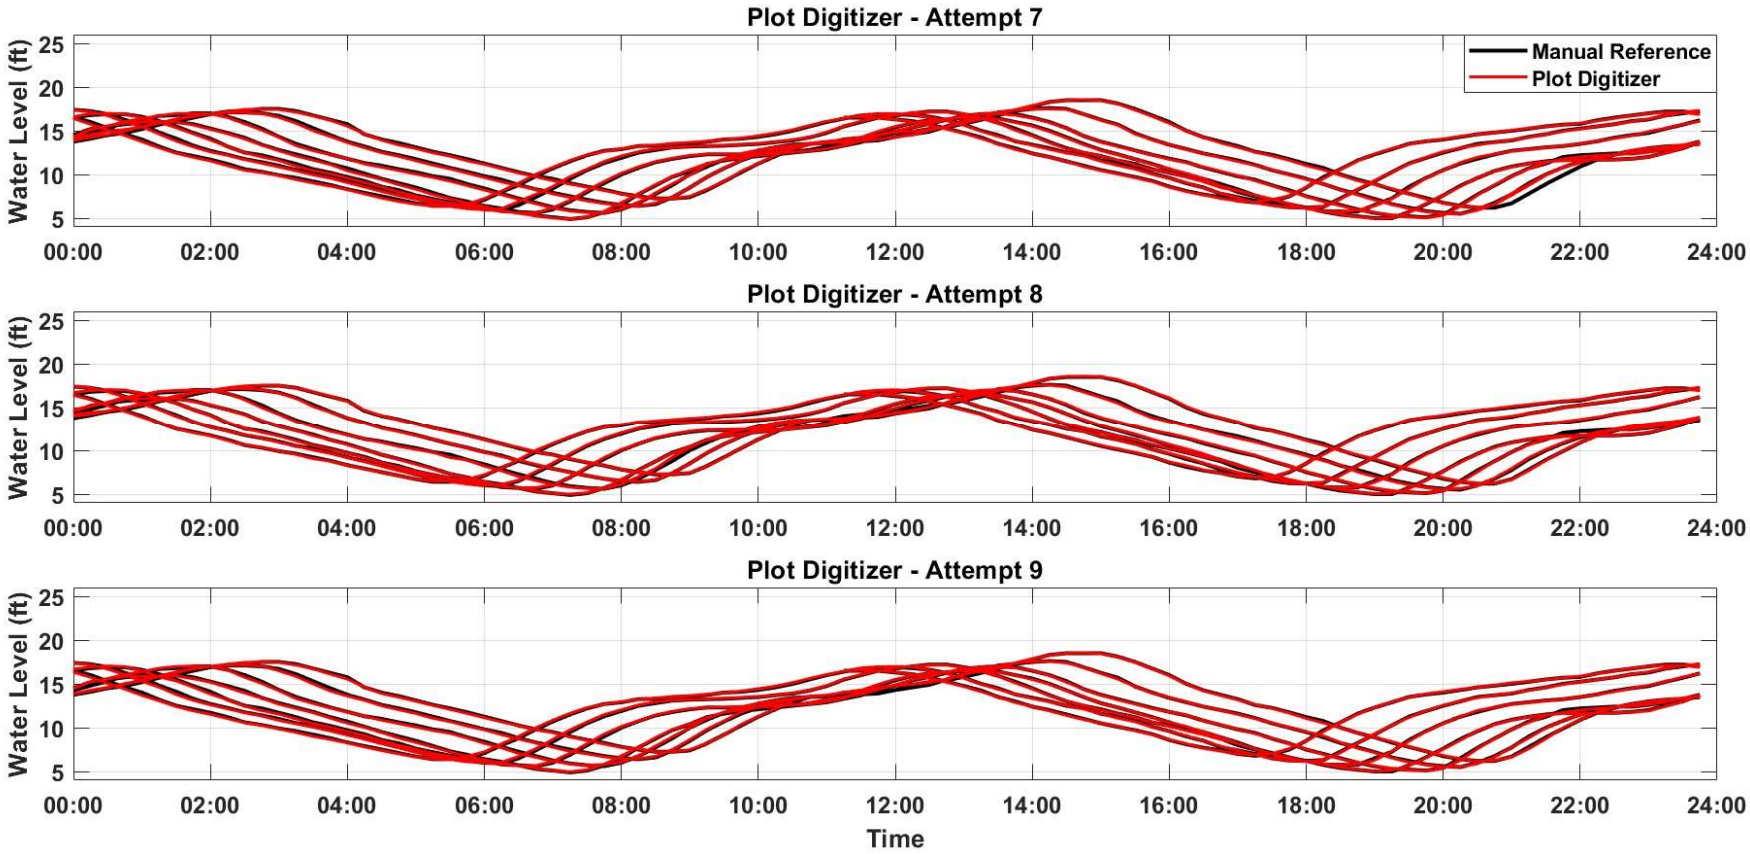

Figure 33. *Attempts 7-9 of using Plot Digitizer to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

3.4. WebPlotDigitizer

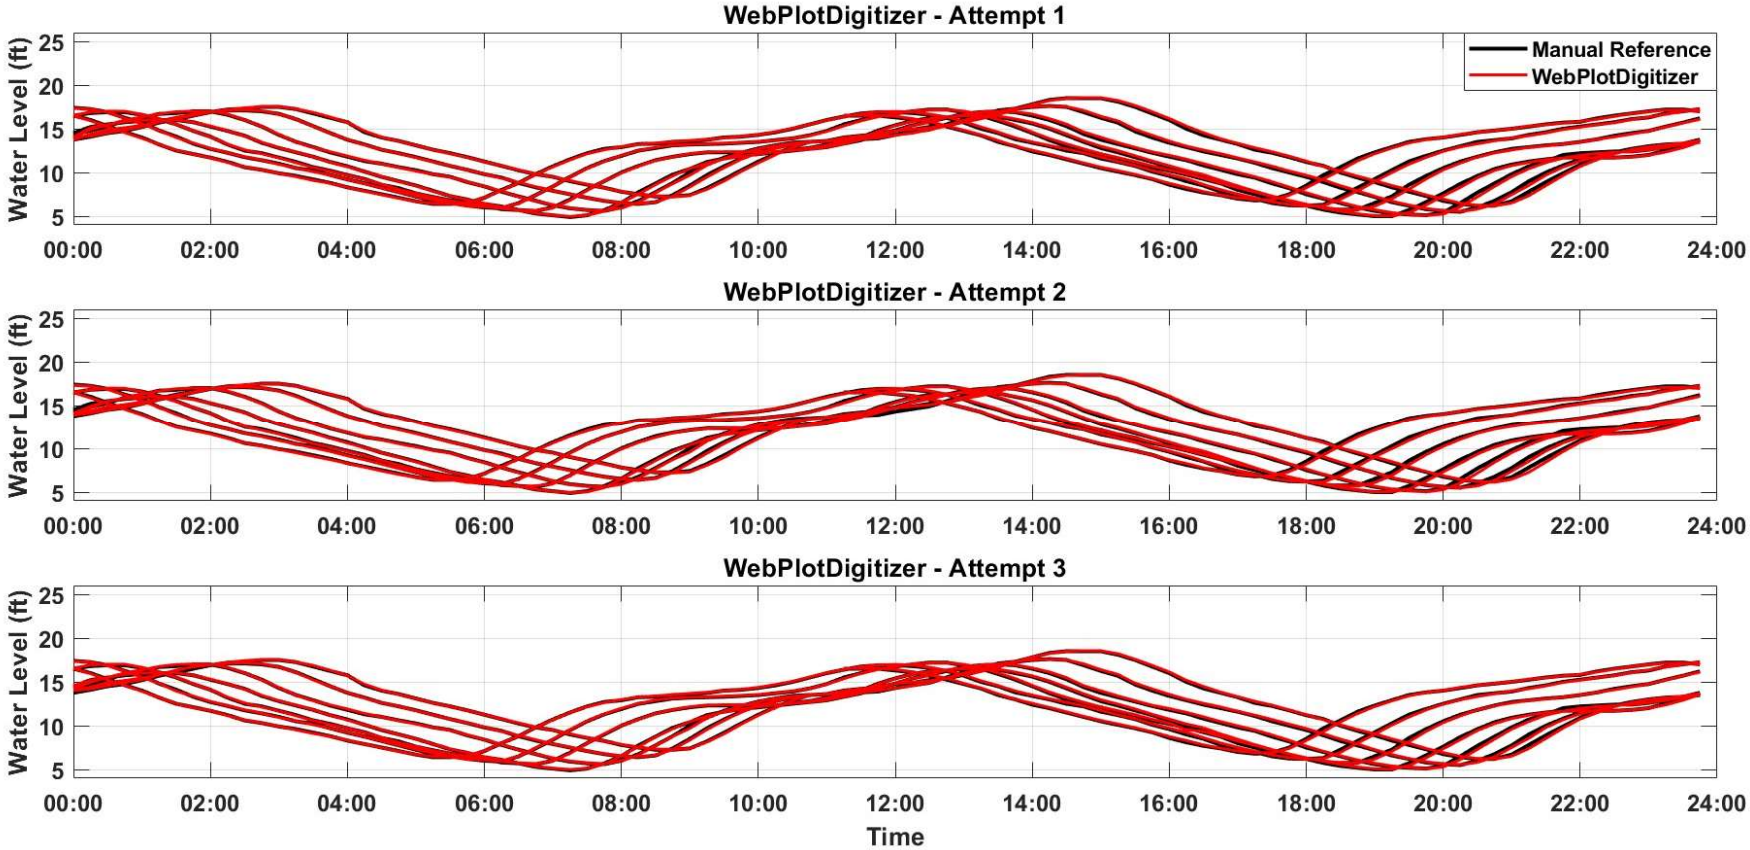

Figure 34. *Attempts 1-3 of using WebPlotDigitizer to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

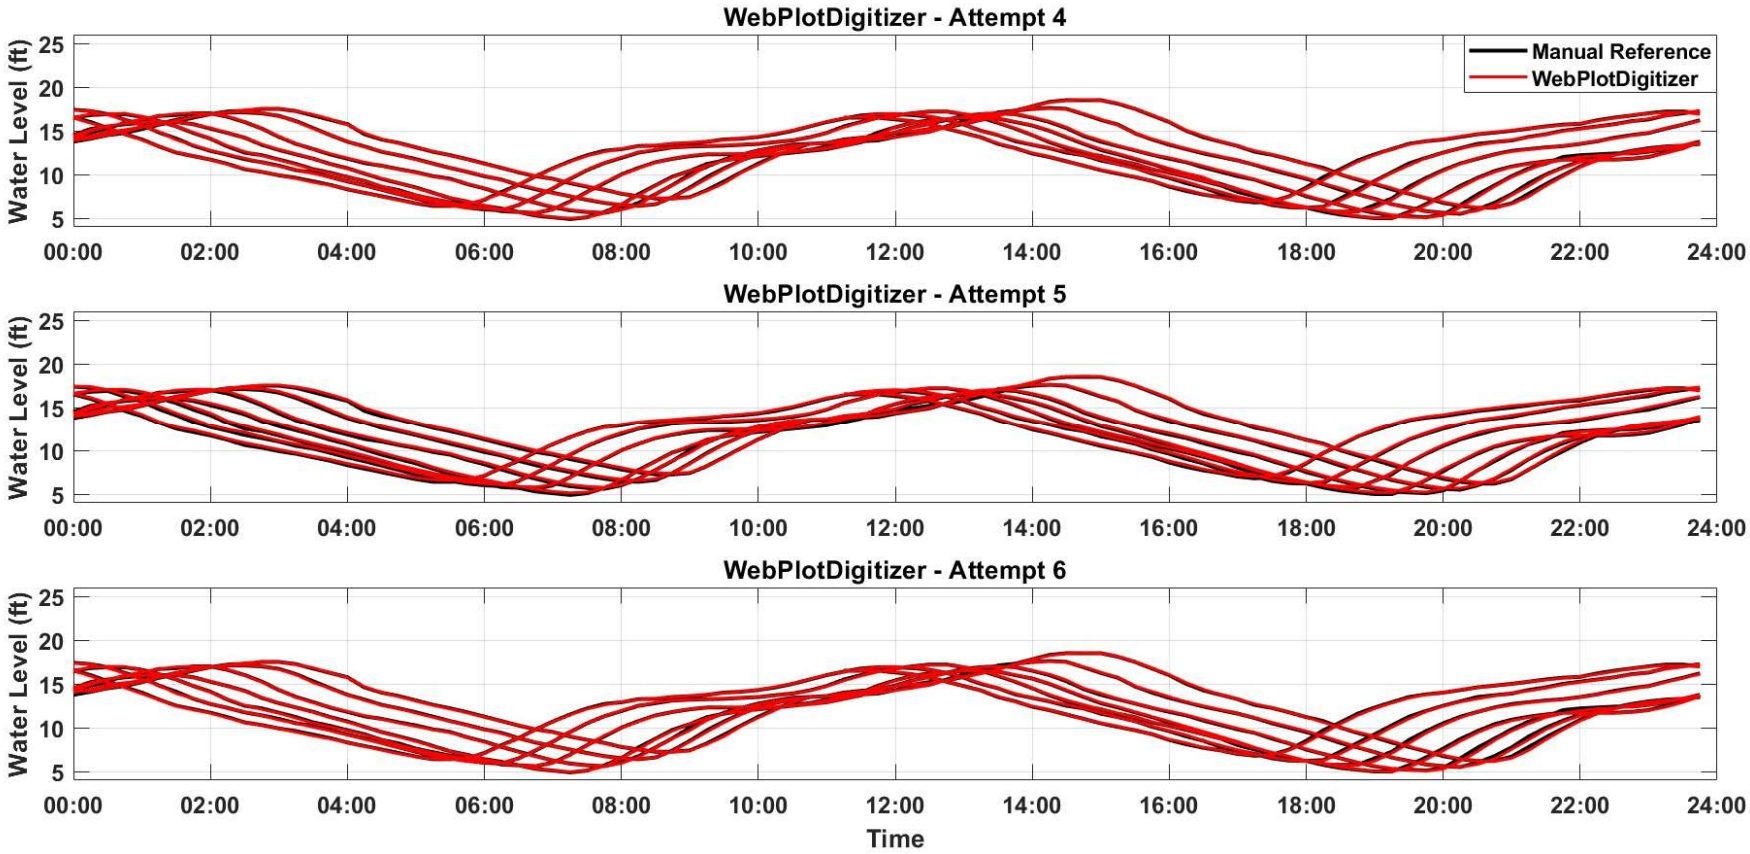

Figure 35. *Attempts 4-6 of using WebPlotDigitizer to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

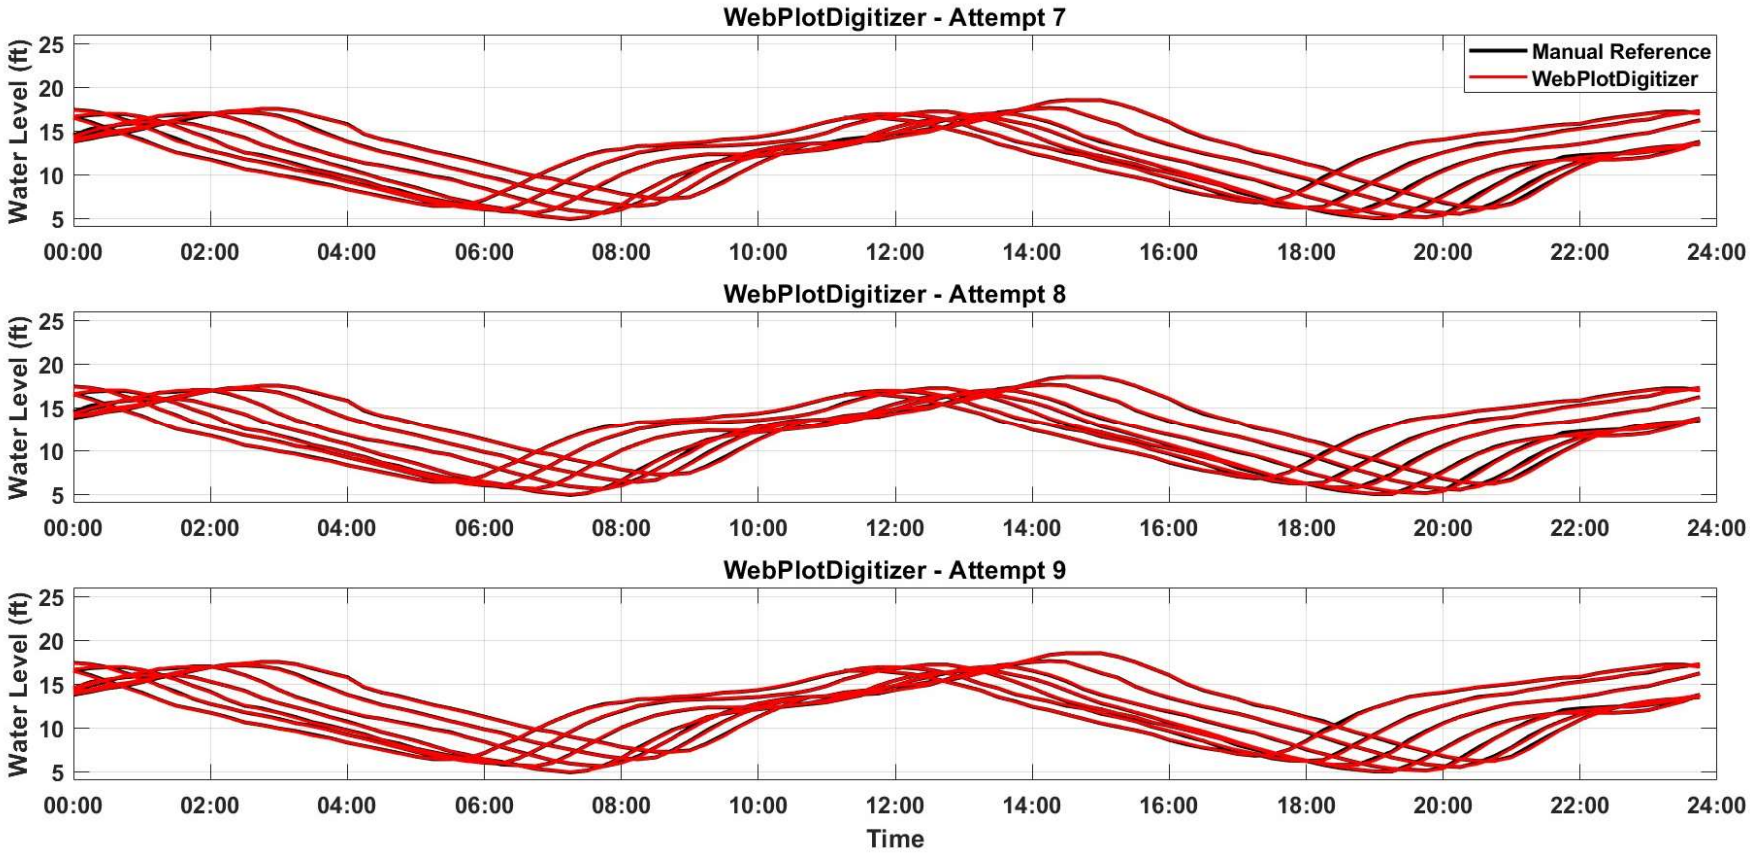

Figure 36. *Attempts 7-9 of using WebPlotDigitizer to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

4. LLM Comparison Plots

4.1. ChatGPT

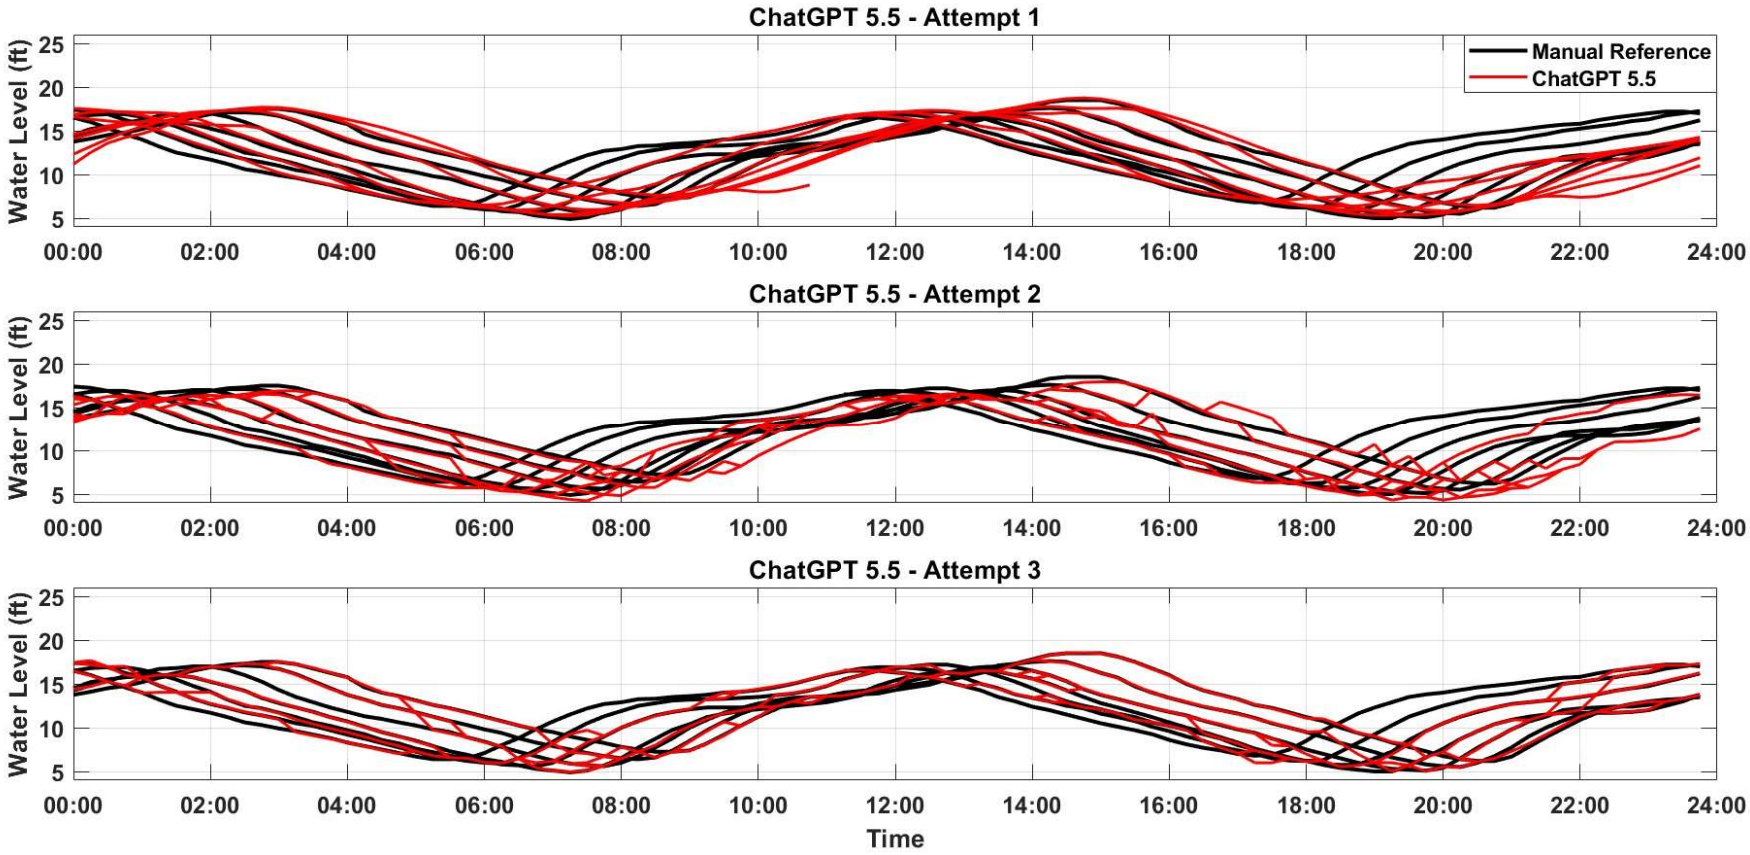

Figure 37. *Attempts 1-3 of using OpenAI ChatGPT Plus (GPT-4o) to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

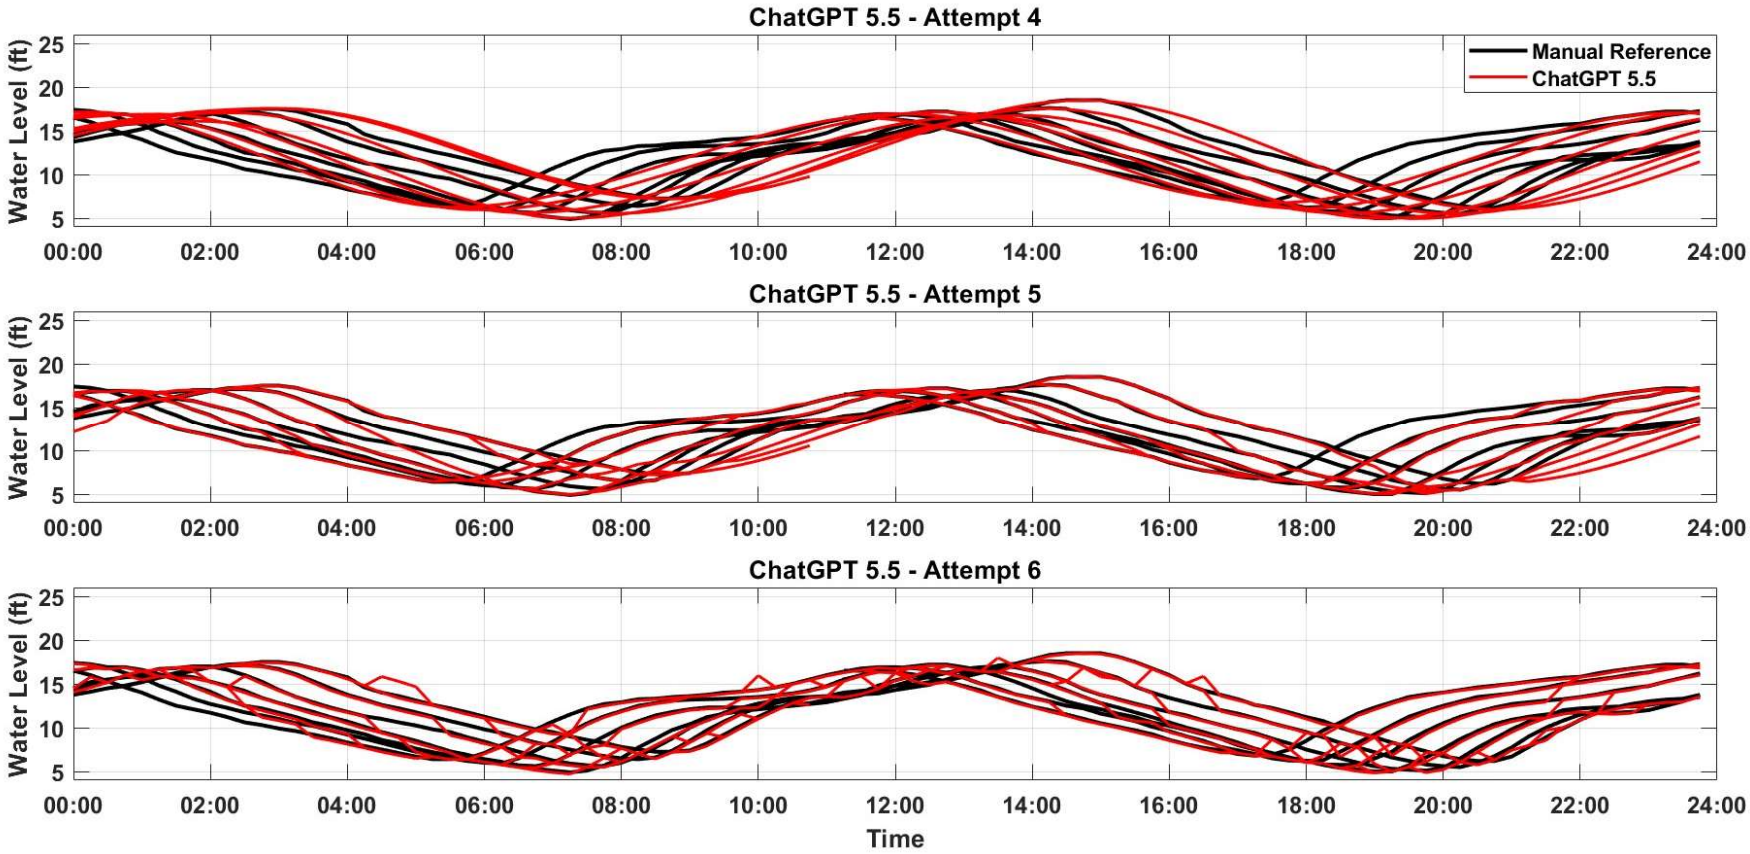

Figure 38. *Attempts 4-6 of using OpenAI ChatGPT Plus (GPT-4o) to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

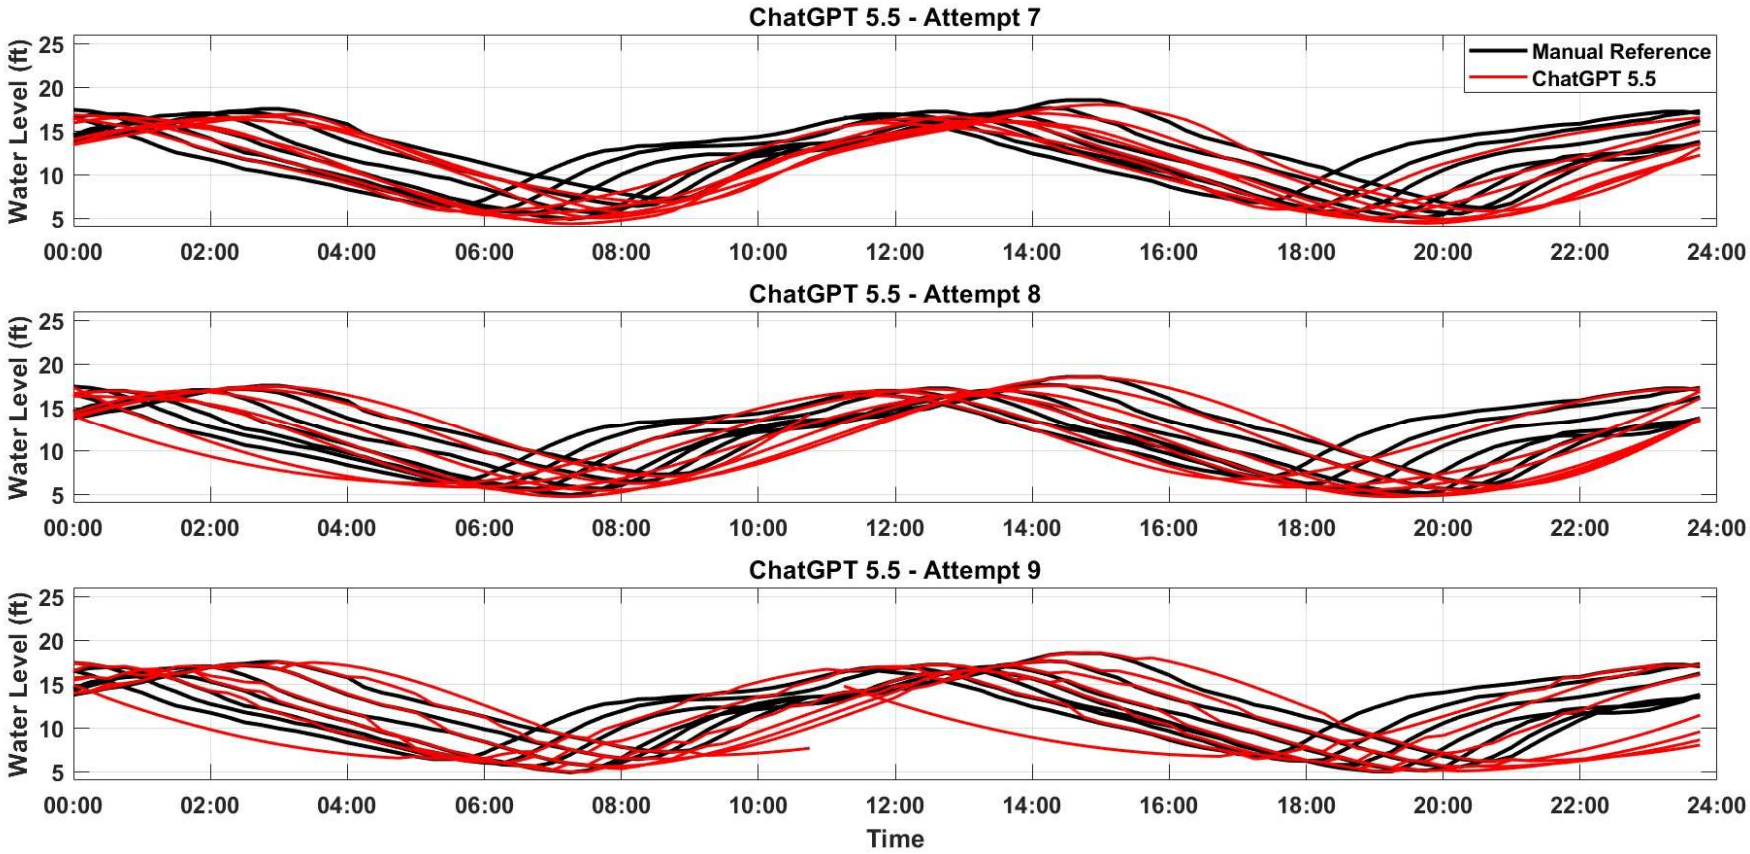

Figure 39. *Attempts 7-9 of using OpenAI ChatGPT Plus (GPT-4o) to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

4.2. Gemini

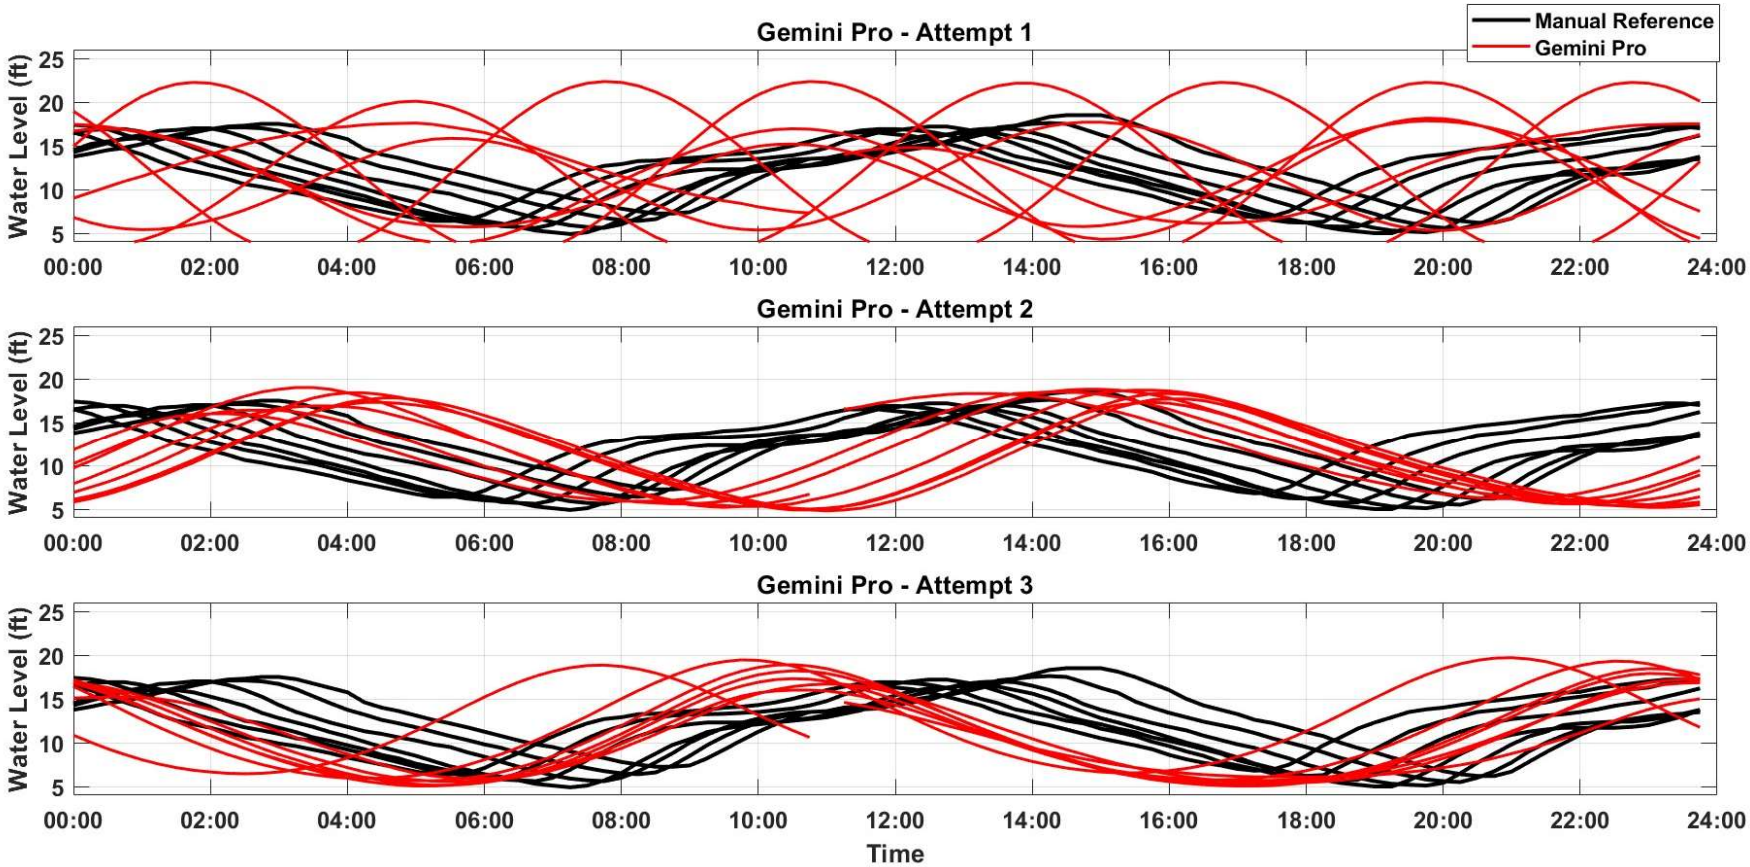

Figure 40. *Attempts 1-3 of using Google Gemini Advanced (Gemini 3.1Pro) to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

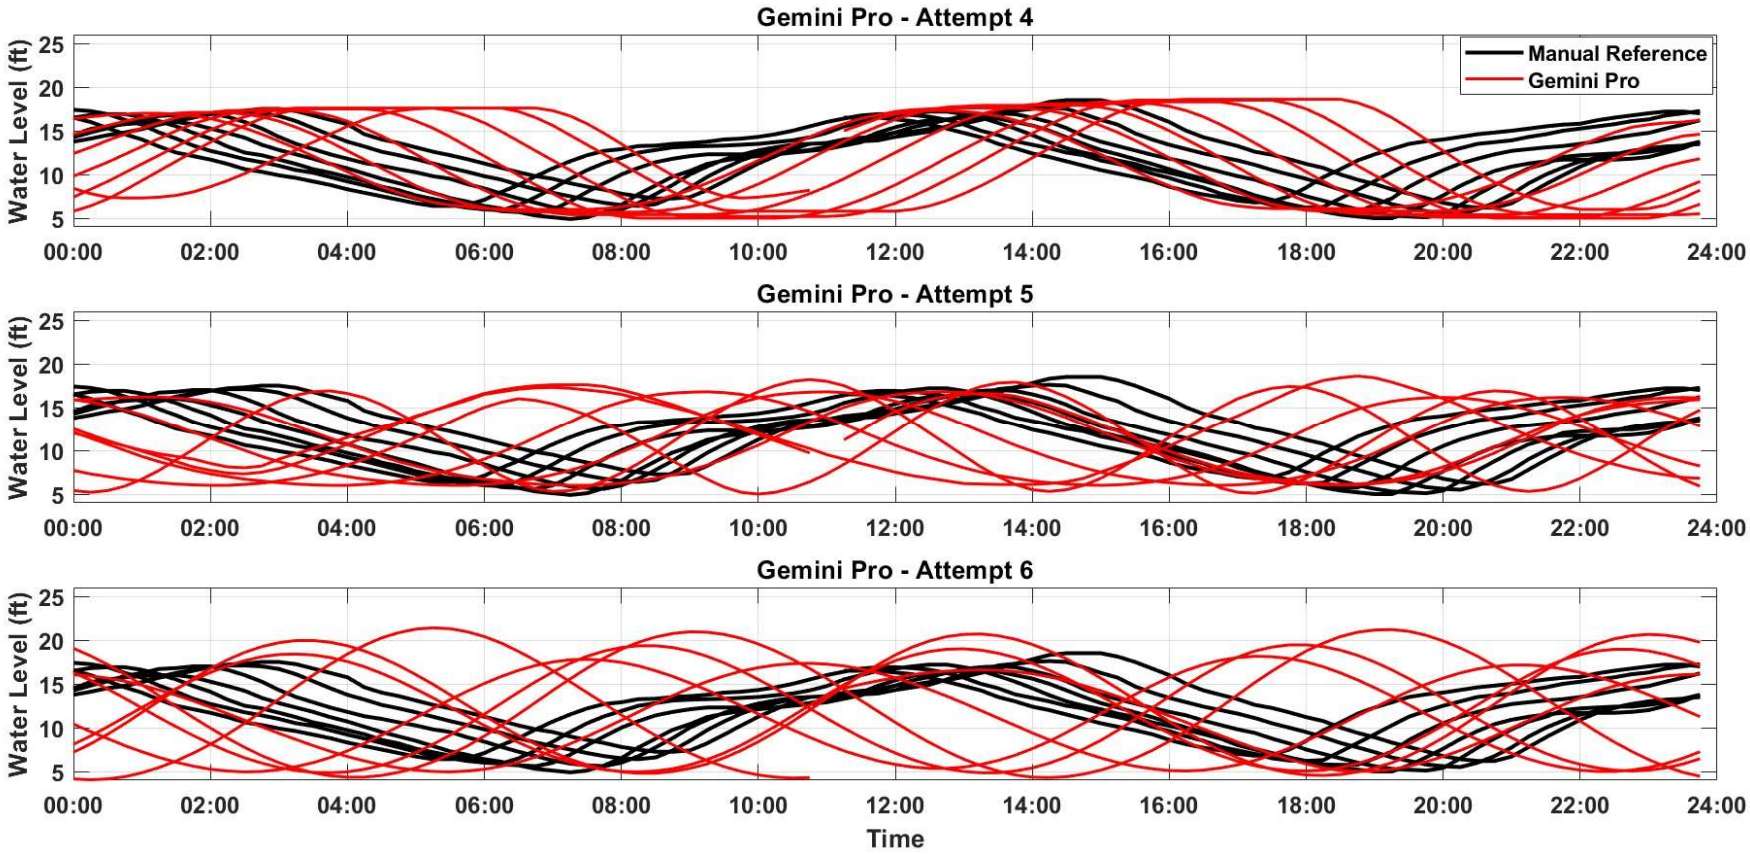

Figure 41. *Attempts 4-6 of using Google Gemini Advanced (Gemini 3.1Pro) to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

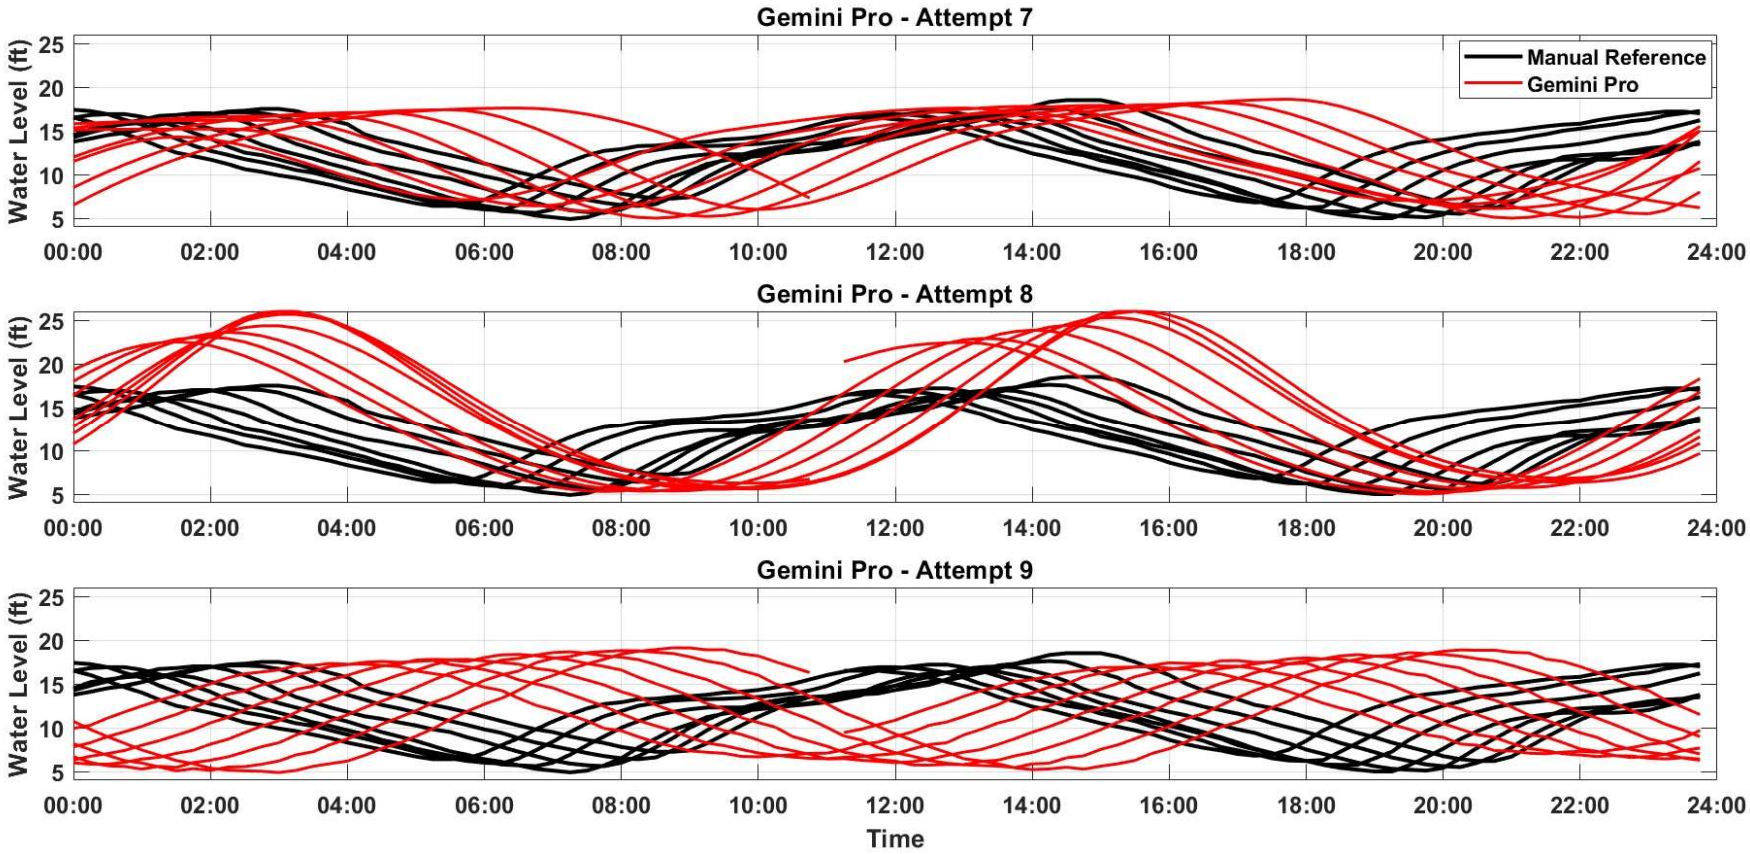

Figure 42. *Attempts 7-9 of using Google Gemini Advanced (Gemini 3.1Pro) to digitise the test marigram, where the red lines show the recovered data and the black shows the reference data.*

Supplementary Material 3

1. All Categories

1.1. Overall Statistical Differences Between Categories

Table 1. *Friedman test results comparing all method categories for accuracy and efficiency. A p-value below 0.05 indicates a statistically significant difference between at least two categories. Kendall's W indicates the effect size, with larger values representing stronger differences between categories.*

	Friedman test P-value	**Kendall's W**
Accuracy	0.001	0.825
Efficiency	0.001	0.911

1.2. Pairwise Comparisons of Accuracy

Table 2. *Pairwise comparisons of accuracy between categories using Bonferroni-corrected Wilcoxon signed-rank tests, p-values below 0.05 indicate a significant difference between those categories. Mean differences are calculated as Category A − Category B.*

Category A	**Category B**	**Mean RMSE Difference (feet)**	**P-value (Bonferroni Corrected)**
Manual	Point-and-Click	-0.044	0.047
Manual	Trace	-0.114	0.234
Manual	LLM	-5.288	0.047
Trace	Point-and-Click	0.069	0.469
Trace	LLM	-5.174	0.047
Point-and-Click	LLM	-5.244	0.047

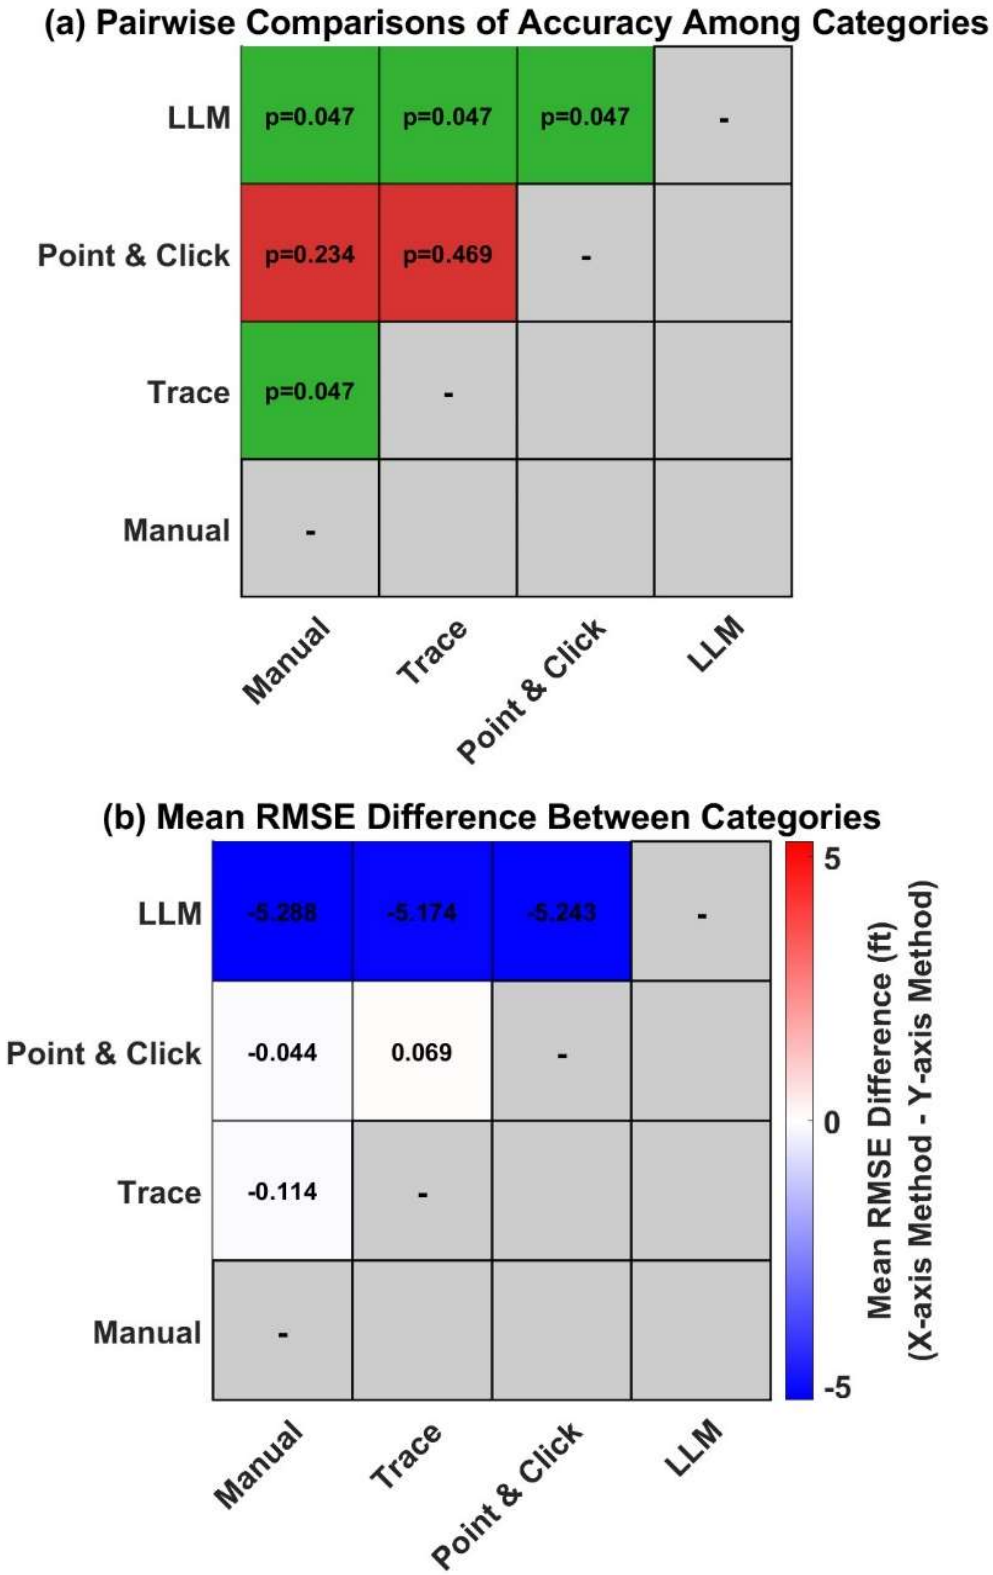

Figure 3. (a) *Pairwise comparisons of accuracy between categories using Bonferroni-corrected Wilcoxon signed-rank tests, where green indicates there is a significant difference* **(b)** *Mean RMSE differences between categories, calculated as the x-axis tool minus the y-axis tool. Grey cells represent redundant self-comparisons or duplicated comparisons.*

1.3. Pairwise Comparisons of Efficiency

Table 3. *Pairwise comparisons of efficiency between categories using Bonferroni-corrected Wilcoxon signed-rank tests, p-values below 0.05 indicate a significant difference between those categories. Mean differences are calculated as Category A − Category B.*

Category A	Category B	Mean Difference of Completion Time (minutes)	P value (Bonferroni Corrected)
Manual	Point-and-Click	104.29	0.023
Manual	Trace	98.431	0.023
Manual	LLM	138.82	0.023
Trace	Point-and-Click	5.862	1.000
Trace	LLM	40.387	0.023
Point-and-Click	LLM	34.525	0.023

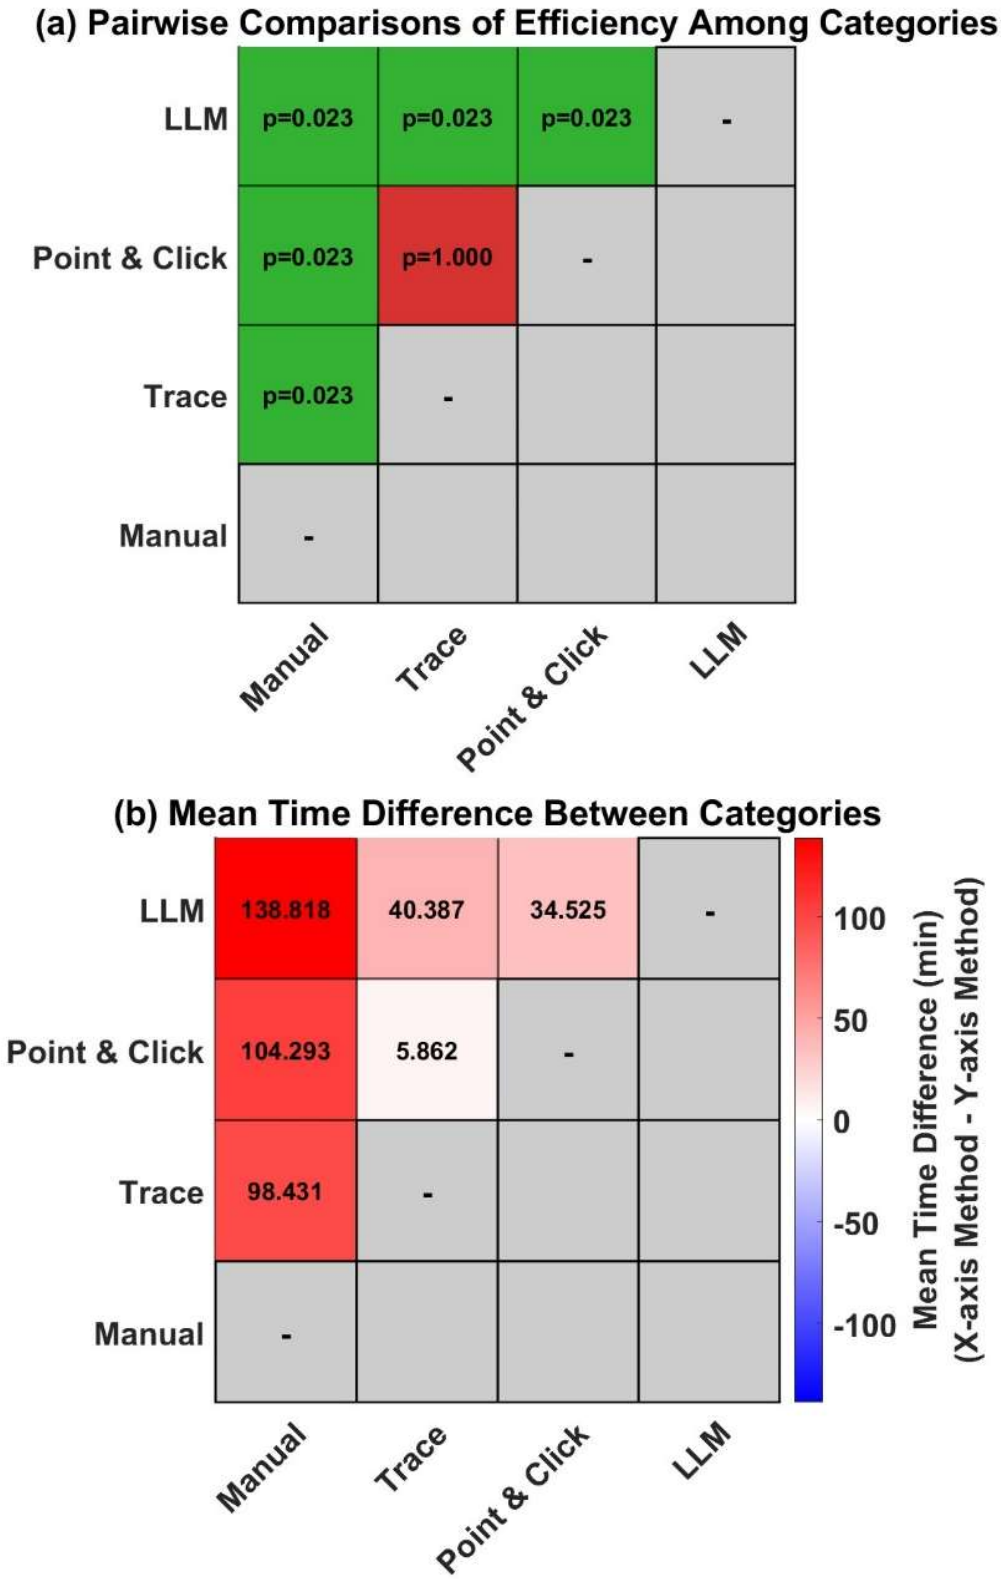

Figure 2. (a) *Pairwise comparisons of efficiency between categories using Bonferroni-corrected Wilcoxon signed-rank tests, where green indicates there is a significant difference* **(b)** *Mean differences in completion time between categories, calculated as the x-axis tool minus the y-axis tool. Grey cells represent redundant self-comparisons or duplicated comparisons.*

2. Point-and-Click Category

2.1. Overall Statistical Differences Between Point-and-Click Tools

Table 4. *Friedman test results comparing all point-and-click tools for accuracy and efficiency. A p-value below 0.05 indicates a statistically significant difference between at least two tools. Kendall's W indicates the effect size, with larger values indicating stronger differences between categories.*

	Friedman test P-value	**Kendall's W**
Accuracy	0.462	0.105
Efficiency	0.252	0.145

2.2. Pairwise Comparisons of Accuracy

Table 5. *Pairwise comparisons of accuracy between point-and-click tools using Bonferroni-corrected Wilcoxon signed-rank tests, p-values below 0.05 indicate a significant difference between those tools. Mean differences are calculated as Category A − Category B.*

Tool A	**Tool B**	**Mean Difference (feet)**	**P value (Bonferroni Corrected)**
GetData Graph Digitizer	Engauge Digitizer	-0.032	1.000
GetData Graph Digitizer	DigitizeIt	-0.022	1.000
GetData Graph Digitizer	GRABIT	0.008	1.000
GetData Graph Digitizer	GraphGrabber	-0.072	1.000
GetData Graph Digitizer	Dagra	-0.045	1.000
GetData Graph Digitizer	OriginLab	-0.105	1.000
Engauge Digitizer	DigitizeIt	0.010	1.000
Engauge Digitizer	GRABIT	0.040	1.000
Engauge Digitizer	GraphGrabber	-0.039	1.000
Engauge Digitizer	Dagra	-0.013	1.000
Engauge Digitizer	OriginLab	-0.073	1.000
DigitizeIt	GRABIT	0.030	1.000
DigitizeIt	GraphGrabber	-0.050	1.000
DigitizeIt	Dagra	-0.023	1.000
DigitizeIt	OriginLab	-0.083	1.000
GRABIT	GraphGrabber	-0.080	1.000
GRABIT	Dagra	-0.053	1.000
GRABIT	OriginLab	-0.113	1.000
GraphGrabber	Dagra	0.027	1.000
GraphGrabber	OriginLab	-0.033	1.000
Dagra	OriginLab	-0.060	1.000

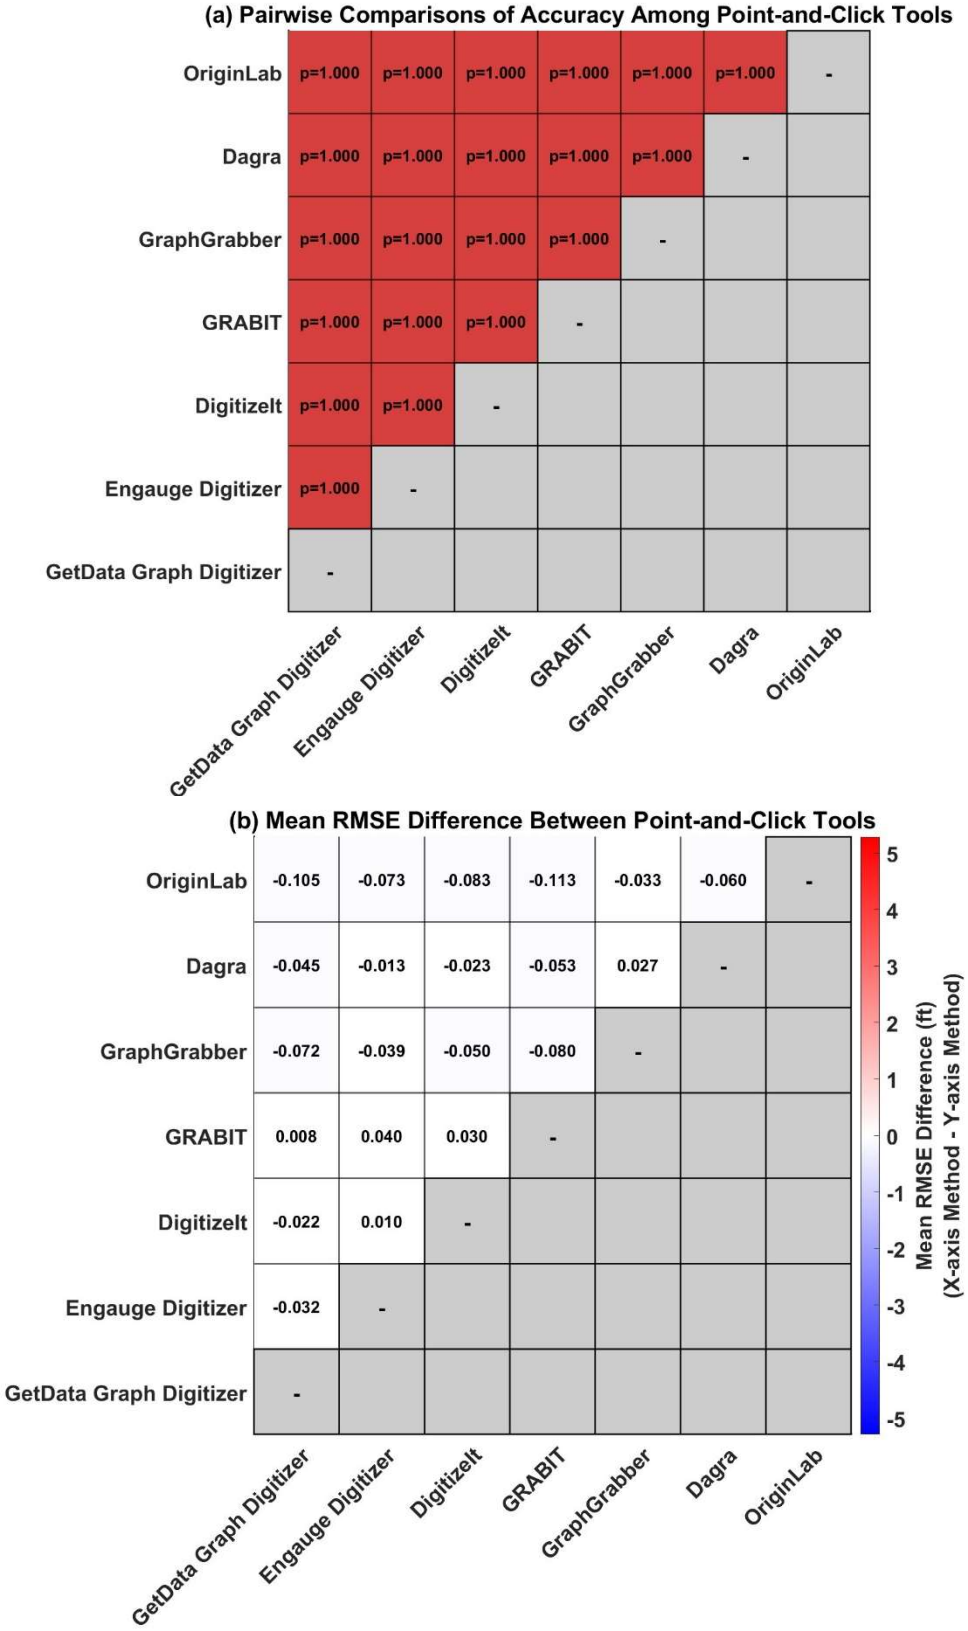

Figure 3. (a) *Pairwise comparisons of accuracy between point-and-click tools using Bonferroni-corrected Wilcoxon signed-rank tests, where green indicates there is a significant difference* **(b)** *Mean RMSE differences between point-and-click tools, calculated as the x-axis tool minus the y-axis tool. Grey cells represent redundant self-comparisons or duplicated comparisons.*

2.3. Pairwise Comparisons of Efficiency

Table 6. *Pairwise comparisons of efficiency between point-and-click tools using Bonferroni-corrected Wilcoxon signed-rank tests, p-values below 0.05 indicate a significant difference between those tools. Mean differences are calculated as Category A − Category B.*

Tool A	Tool B	Mean Difference (minutes)	P value (Bonferroni Corrected)
GetData Graph Digitizer	Engauge Digitizer	9.301	0.41
GetData Graph Digitizer	DigitizeIt	3.422	1.000
GetData Graph Digitizer	GRABIT	-1.531	1.000
GetData Graph Digitizer	GraphGrabber	3.993	0.82
GetData Graph Digitizer	Dagra	5.572	1.000
GetData Graph Digitizer	OriginLab	2.853	1.000
Engauge Digitizer	DigitizeIt	-5.879	1.000
Engauge Digitizer	GRABIT	-10.832	1.000
Engauge Digitizer	GraphGrabber	-5.308	1.000
Engauge Digitizer	Dagra	-3.729	1.000
Engauge Digitizer	OriginLab	-6.448	1.000
DigitizeIt	GRABIT	-4.953	1.000
DigitizeIt	GraphGrabber	0.571	1.000
DigitizeIt	Dagra	2.150	1.000
DigitizeIt	OriginLab	-0.569	1.000
GRABIT	GraphGrabber	5.524	1.000
GRABIT	Dagra	7.103	1.000
GRABIT	OriginLab	4.384	1.000
GraphGrabber	Dagra	1.579	1.000
GraphGrabber	OriginLab	-1.140	1.000
Dagra	OriginLab	-2.719	1.000

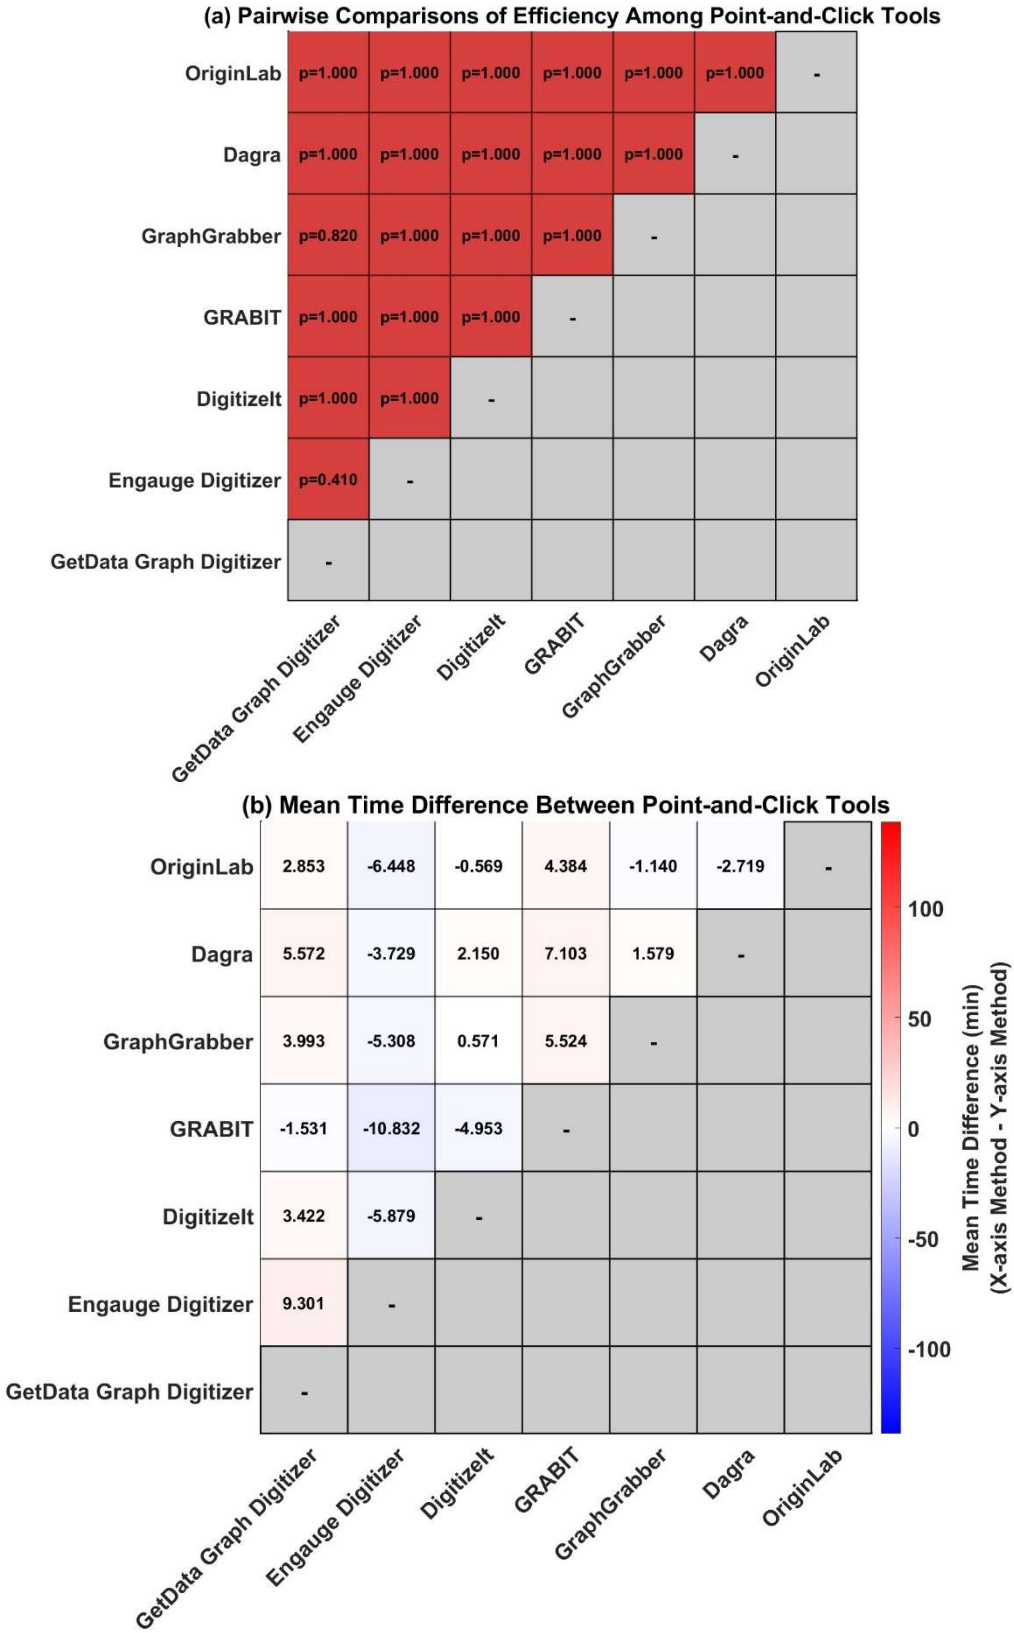

Figure 4. (a) *Pairwise comparisons of efficiency between point-and-click tools using Bonferroni-corrected Wilcoxon signed-rank tests, where green indicates there is a significant difference* **(b)** *Mean differences in completion time between point-and-click tools, calculated as the x-axis tool minus the y-axis tool. Grey cells represent redundant self-comparisons or duplicated comparisons.*

3. Trace Tools Category

3.1. Overall Statistical Differences Between Trace Tools

Table 7. *Friedman test results comparing all trace methods for accuracy and efficiency. A p-value below 0.05 indicates a statistically significant difference between at least two tools. Kendall's W indicates the effect size, with larger values representing stronger differences between categories.*

	Friedman test P-value	**Kendall's W**
Accuracy	0.001	0.644
Efficiency	0.001	0.635

3.2. Pairwise Comparisons of Accuracy

Table 8. *Pairwise comparisons of accuracy between trace tools using Bonferroni-corrected Wilcoxon signed-rank tests, p-values below 0.05 indicate a significant difference between those tools. Mean differences are calculated as Category A − Category B.*

Tool A	**Tool B**	**Mean RMSE Difference (feet)**	**P value (Bonferroni Corrected)**
NUNIEAU	WebPlotDigitizer	0.054	0.070
NUNIEAU	Plot Digitizer	-0.080	0.338
NUNIEAU	DigitGraph	-0.074	1.000
WebPlotDigitizer	Plot Digitizer	-0.133	0.023
WebPlotDigitizer	DigitGraph	-0.128	0.023
Plot Digitizer	DigitGraph	0.005	1.000

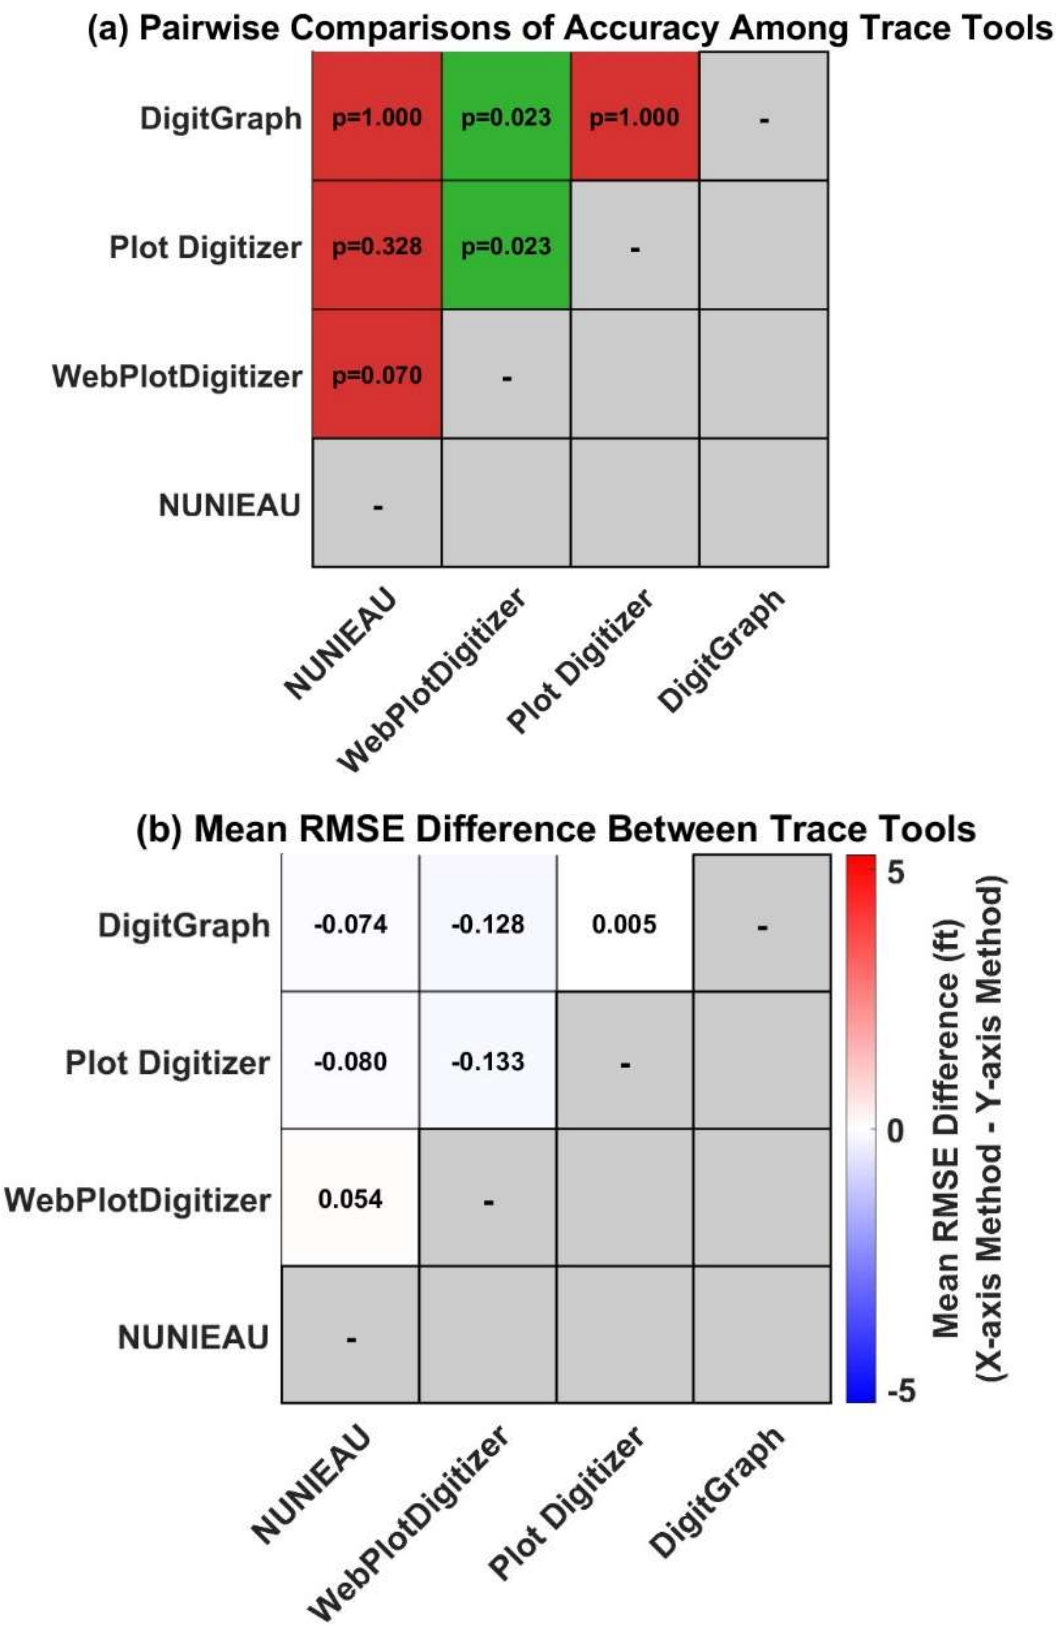

Figure 5. (a) *Pairwise comparisons of accuracy between trace tools using Bonferroni-corrected Wilcoxon signed-rank tests, where green indicates there is a significant difference* **(b)** *Mean RMSE differences between trace tools, calculated as the x-axis tool minus the y-axis tool. Grey cells represent redundant self-comparisons or duplicated comparisons.*

3.3. Pairwise Comparisons of Efficiency

Table 9. *Pairwise comparisons of efficiency between trace tools using Bonferroni-corrected Wilcoxon signed-rank tests, p-values below 0.05 indicate a significant difference between those tools. Mean differences are calculated as Category A − Category B.*

Tool A	Tool B	Mean Difference (minutes)	P value (Bonferroni Corrected)
NUNIEAU	WebPlotDigitizer	15.156	0.773
NUNIEAU	Plot Digitizer	40.711	0.023
NUNIEAU	DigitGraph	31.347	0.023
WebPlotDigitizer	Plot Digitizer	25.556	0.164
WebPlotDigitizer	DigitGraph	16.191	0.984
Plot Digitizer	DigitGraph	-9.364	0.984

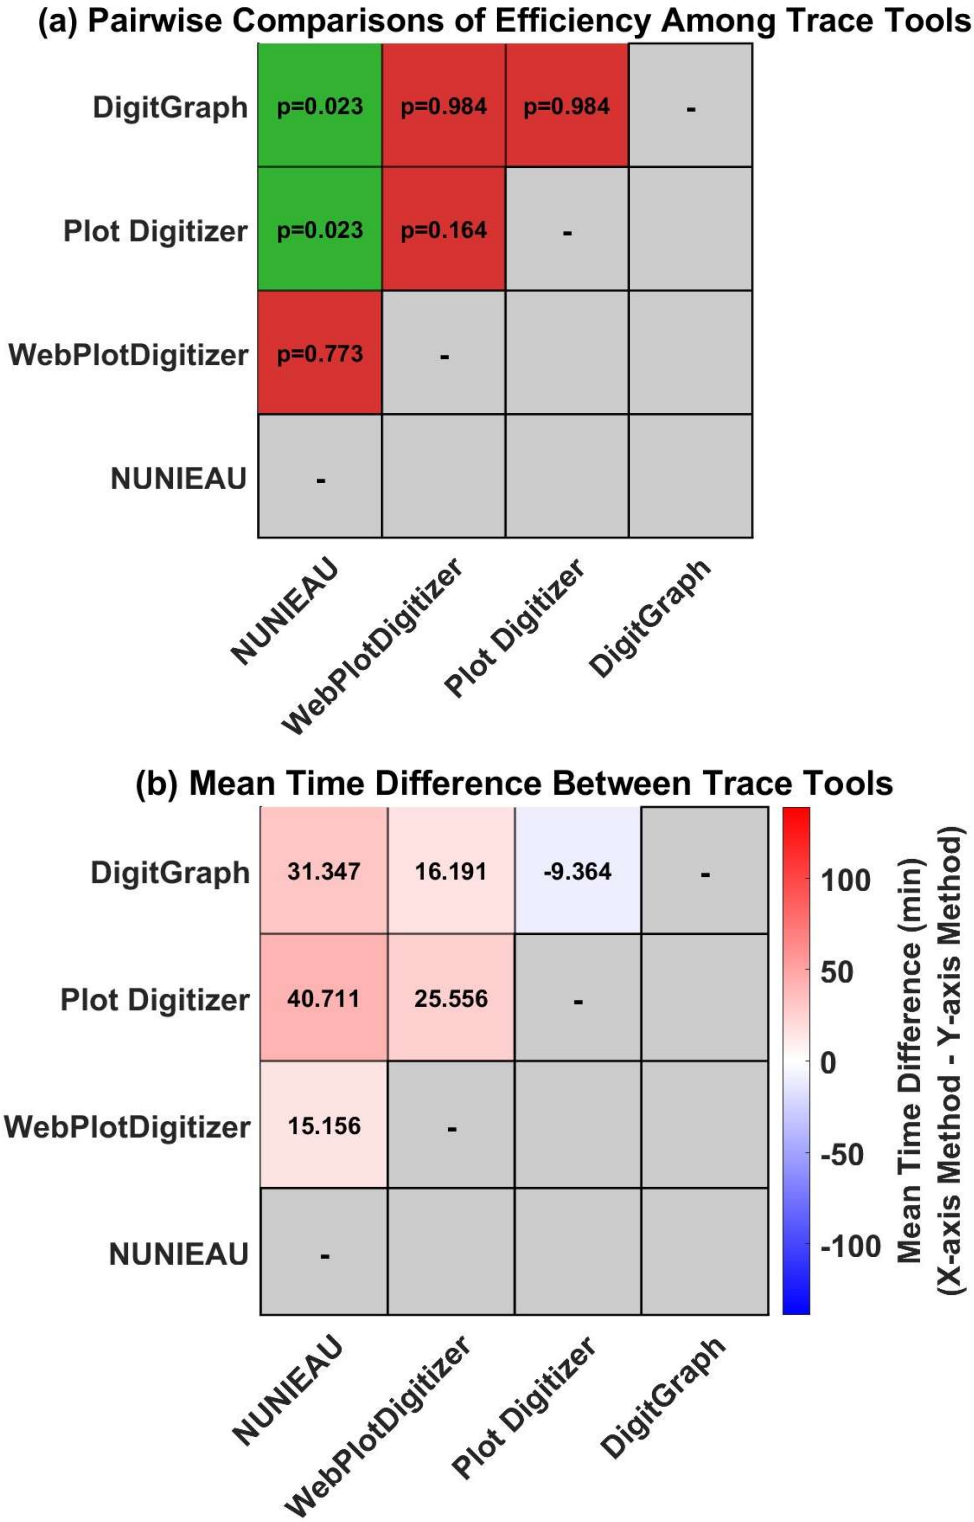

Figure 6. (a) *Pairwise comparisons of efficiency between trace tools using Bonferroni-corrected Wilcoxon signed-rank tests, where green indicates there is a significant difference* **(b)** *Mean differences in completion time between trace tools, calculated as the x-axis tool minus the y-axis tool. Grey cells represent redundant self-comparisons or duplicated comparisons.*

4. LLM Category

4.1. Pairwise Comparisons of Accuracy

Table 10. *Pairwise comparisons of accuracy between LLMs using Bonferroni-corrected Wilcoxon signed-rank tests, p-values below 0.05 indicate a significant difference between those LLMs. Mean differences are calculated as Category A − Category B.*

Tool A	Tool B	Mean RMSE Difference (feet)	P value (Bonferroni Corrected)
ChatGPT	Gemini	0.330	0.004

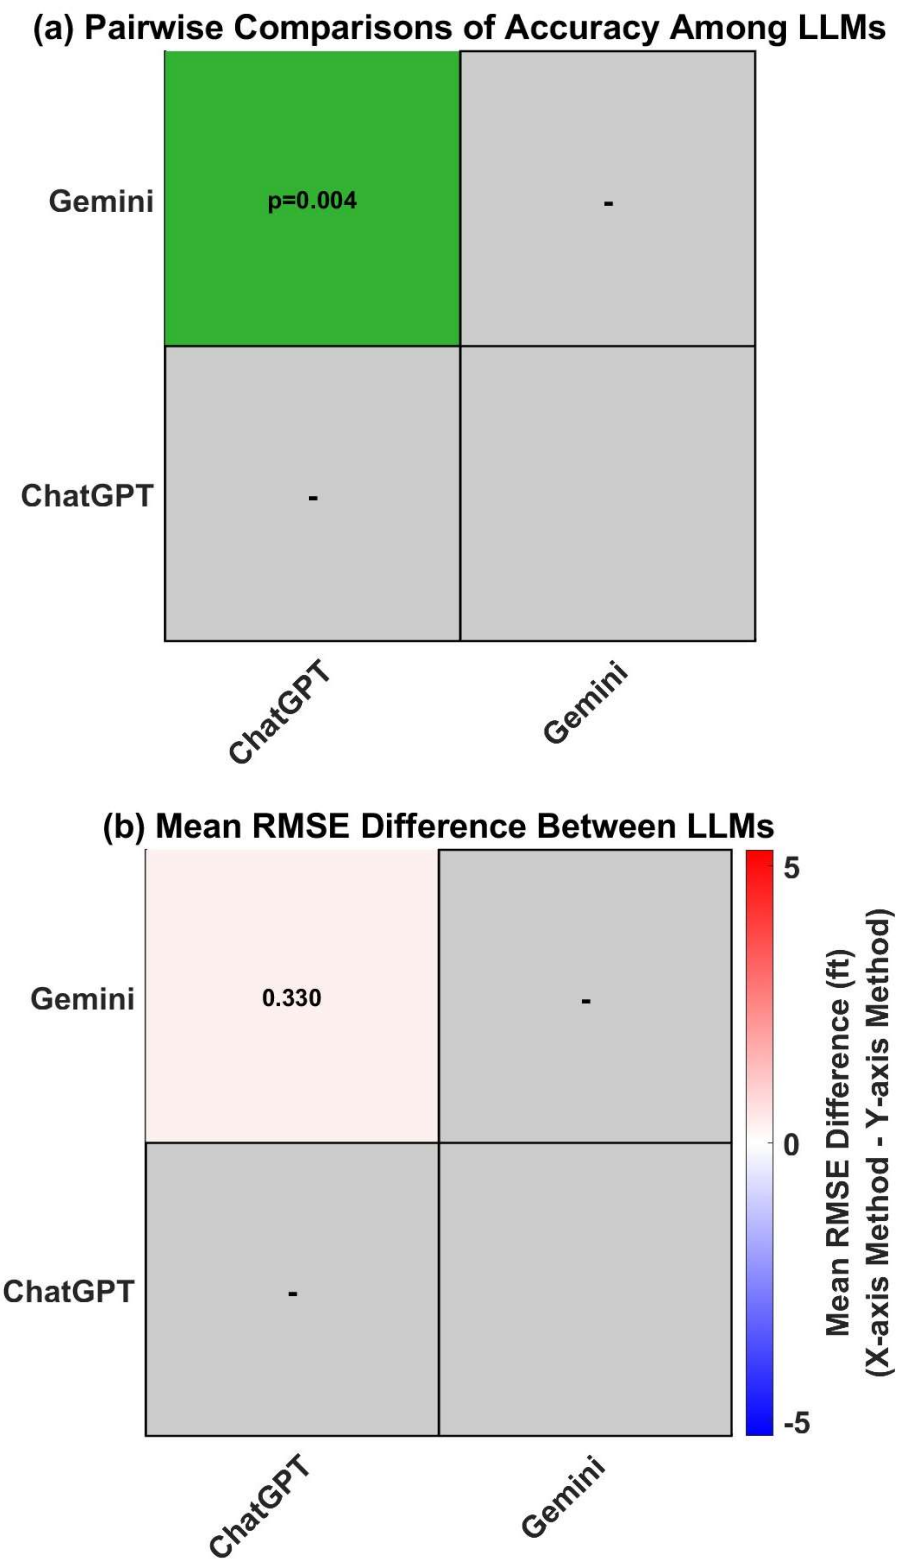

Figure 7. (a) *Pairwise comparisons of accuracy between LLMs using Bonferroni-corrected Wilcoxon signed-rank tests, where green indicates there is a significant difference* **(b)** *Mean RMSE differences between LLMs, calculated as the x-axis tool minus the y-axis tool. Grey cells represent redundant self-comparisons or duplicated comparisons.*

4.3. Pairwise Comparisons of Efficiency

Table 11. *Pairwise comparisons of efficiency between LLMs using Bonferroni-corrected Wilcoxon signed-rank tests, p-values below 0.05 indicate a significant difference between those LLMs. Mean differences are calculated as Category A − Category B.*

Tool A	Tool B	Mean Difference (minutes)	P value (Bonferroni Corrected)
ChatGPT	Gemini	-7.276	0.004

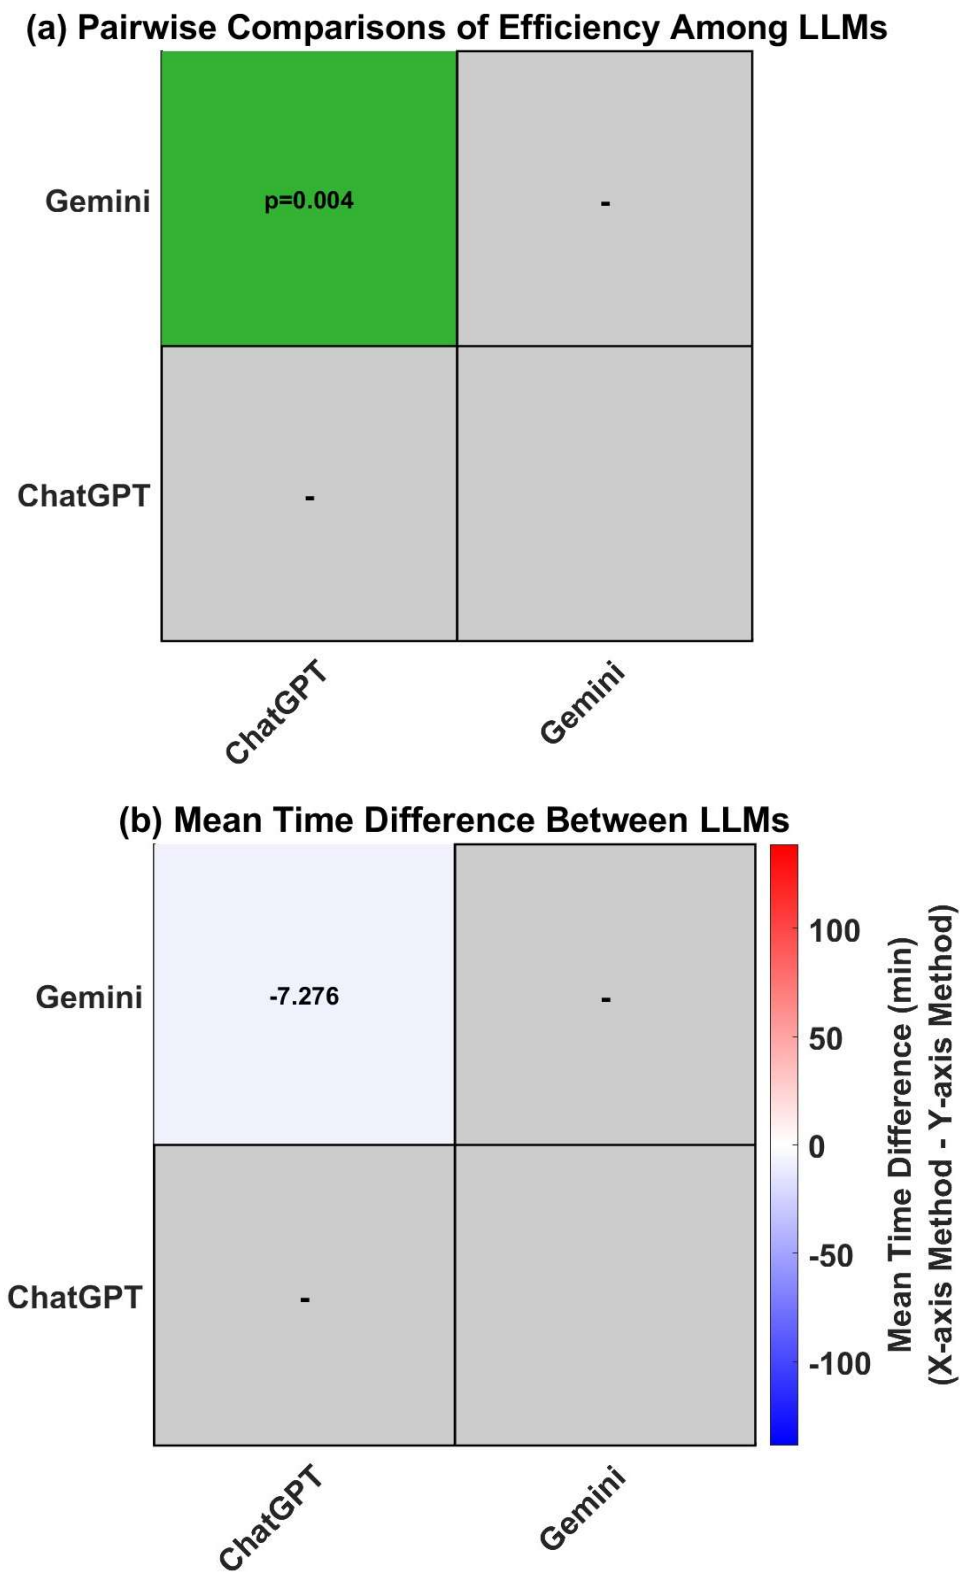

Figure 8. (a) *Pairwise comparisons of efficiency between LLMs using Bonferroni-corrected Wilcoxon signed-rank tests, where green indicates there is a significant difference* **(b)** *Mean differences in completion time between LLMs, calculated as the x-axis tool minus the y-axis tool. Grey cells represent redundant self-comparisons or duplicated comparisons.*