Journal of Scientometric Research, 2026; 15(3):860-873.
https://www.jscires.org

Research Article

A Scientometric Approach to Theological Continuity: Measuring Stationarity in the Patrologia Graeca via Dynamic Time Warping

Evagelos Varthis[1,*], Kostas Domouchtsis[2], Ilias Chrysostomidis[1], Marios Poulos[1]

[1]School of Information Science and Informatics, Ionian University, Corfu, GREECE.
[2]Department of Computer, Informatics and Telecommunications Engineering, International University of Greece, Serres, GREECE.

ABSTRACT

This study presents a scientometric approach to examine continuity and structural change within the patristic hermeneutical tradition concerning the Books of Genesis, Psalms, Matthew, Luke and John, using a curated dataset linking Christian authors, their works, and biblical commentaries. By transforming biblical references drawn from chronologically ordered Church Fathers into a discrete historical time series, we assess whether theological engagement exhibits stable or shifting temporal structure. Conventional econometric stationarity tests are informative but limited for this corpus because of pronounced volume imbalances, most notably the disproportionate exegetical output of John Chrysostom, Cyril of Alexandria, and Athanasius of Alexandria. To address this limitation, we employ the framework of Reversible Stochastic Processes (RSP) together with Dynamic Time Warping (DTW) in a sliding-window design, allowing comparison of non-linear temporal patterns and localized variation. The findings suggest that the patristic citation signal is better described as structurally heterogeneous than as strictly stationary, and they indicate recurring intervals of elevated change of roughly 8-10 authors for the five most cited books. These patterns are consistent with a history shaped by uneven citation densities and doctrinal developments, and they illustrate the potential of signal-processing techniques for Digital Humanities research.

Keywords: Computational Patrology, Digital Humanities, Dynamic Time Warping, Reversible Stochastic Processes, Temporal Analysis of Exegesis.

Correspondence:
Evagelos Varthis
School of Information Science and Informatics, Ionian University, Corfu, GREECE.
Email: evarthis@ionio.gr
ORCID: 0000-0003-2970-2156

Received: 16-06-2026;
Revised: 06-08-2026;
Accepted: 21-09-2026.

INTRODUCTION

The intersection of Digital Humanities and Scientometrics has opened unprecedented avenues for the quantitative analysis of large-scale historical and theological corpora (Schreibman *et al*., 2004; Berry, 2012; Küng and Tracy, 1989). Early pioneers of computational linguistics envisioned such macroscopic analyses of religious texts (Busa, 1980), and modern techniques of macroanalysis (Jockers, 2013) and "distant reading" (Moretti, 2005) have proven invaluable for tracking thematic evolutions over long periods. Among these corpora, the *Patrologia Graeca (PG)* stands as a monumental repository of Christian thought comprising 170 volumes and 225,000 columns of text (Paparnakis, 2014; Varthis, 2022), mapping centuries of dogmatic evolution and exegetical tradition. A foundational element within this extensive textual corpus is the interpretation of the five biblical books most frequently cited by the authors of PG: Book of Genesis, Book of Psalms, Gospel of Matthew, Gospel of Luke, and Gospel of John. Exegetical engagement with these texts structured the development of interpretive techniques, hermeneutical frameworks, and modes of scriptural reasoning that underpinned early Christian anthropology and cosmology. Traditionally, the study of patristic exegesis has relied almost exclusively on qualitative, historical-critical methodologies. Yet biblical passages function within patristic literature as cited sources, while the Church Fathers themselves operate as citation agents. This structural analogy makes it possible to apply established bibliometric methods to patristic texts, including analyses of citation frequency distributions, co-citation networks, and other regularities commonly examined in quantitative scholarship. Furthermore, by treating the chronological succession of patristic authors as a discrete temporal axis, shifts in theological emphasis can be formally modeled and analyzed as a measurable time series from which the present time-series analysis derives patterns of continuity, periodicity, and structural change in patristic interpretation. This transformation raises a fundamental quantitative question: Is the historical transmission of theological interest a continuous, stationary process, or is it

OPEN ACCESS

Manuscript TECHNOMEDIA

DOI: 10.5530/jscires.20260068

Copyright Information :
Copyright Author (s) 2026 Distributed under Creative Commons CC-BY 4.0

Publishing Partner : Manuscript Technomedia. [www.mstechnomedia.com]

characterized by pronounced non-stationary structural breaks and paradigm shifts?

Standard econometric and time-series stationarity tests such as the Augmented Dickey Fuller (ADF) or Kwiatkowski-Phillip s-Schmidt-Shin (KPSS) tests often fail to capture the true dynamics of humanistic and historical data (Busetti and Taylor, 2003; Cavaliere and Xu, 2014; DeJong *et al*., 1992). These conventional metrics are highly sensitive to extreme outliers and massive volumetric variations-such as the unprecedented exegetical output of prominent figures like John Chrysostom or Athanasius of Alexandria-resulting in conflicting or misleading statistical interpretations. To address this important methodological gap, the present study introduces a novel application of the Reversible Stochastic Process (RSP) theory to PG bibliometric data. Following the Stationarity Distance framework proposed by Poulos (2016) and extended by Poulos (2020), we use Dynamic Time Warping (DTW) with a sliding-window procedure to evaluate the hermeneutical tradition. This algorithmic framework addresses some limitations of linear statistical tests, offering a flexible metric capable of detecting recurring patterns and localized structural variation.

The remainder of this paper is structured as follows: Section 2 details the dataset provenance, preprocessing, and author-level aggregation. Section 3 outlines the theoretical framework and the algorithmic steps of the RSP-DTW methodology. Section 4 presents the empirical results, including descriptive citation patterns, DTW peak profiles, peak-spacing estimates, and KPSS/ADF tests. Section 5 evaluates robustness through bootstrap, multi-window, and KPSS-bandwidth sensitivity analyses. Section 6 discusses the historical and theological implications, while Section 7 concludes the study.

Data and Preprocessing

Dataset Provenance

The empirical foundation of this study is an extensive dataset drawn from the broader corpus of patristic literature, with primary reference to the Patrologia Graeca, edited by Jacques-Paul Migne. The dataset is analogous to BiblIndex (2024), the well-known index of biblical citations across Christian literature, but focuses specifically on the patristic texts contained in the Patrologia Graeca and on authors associated with Eastern Christianity. The dataset was originally compiled more than fifty years ago by S. Sakkos and P. Koutlemanis, professors of theology at the Aristotle University of Thessaloniki, and others collaborators, see Table 1. It provides a structured mapping between biblical verses and the passages in which these verses are cited, discussed, or interpreted within the writings of Greek patristic authors in PG.

The dataset was further enhanced by Paparnakis and Domouchtsis (2017). An updated version is currently available as a web service (https://domis.sites.sch.gr), provided by K. Domouchtsis, offering a systematic online mapping between biblical verses and the corresponding passages in the Patrologia Graeca (PG) in which these.

Verses are discussed or commented upon. The enhanced dataset also includes direct hyperlinks to both the relevant patristic texts and the associated biblical verses, thereby enabling immediate access to the underlying textual sources. Thematic descriptions for the citations are also provided. The dataset has also been used for the photographic reprint of the PG by the Centre of Patristic Editions (1985-2010) in Athens, Greece. The dataset subset comprising the five biblical books analyzed in this study is publicly available through Zenodo for research use (https://doi.org/10.5281/zenodo.21212218).

Deduplication and Author-Level Aggregation

In the computational preprocessing stage, the source database contained 1,216,761 rows. Duplicate citation records were removed using the combined key of biblical verse code, PG volume number, Migne column number, and author name. After this deduplication step, 622,622 unique citation records were retained for the subsequent book- and author-level aggregation. The five books account for 39% of the 622,622 unique citation records having 242,992 rows. Records with missing author attribution, missing biblical-book identifiers, or ambiguous verse mappings were not included in the book-level time-series aggregation, because they could not be assigned unambiguously to an author-book citation count. Python scripts used for preprocessing the dataset, RSP-DTW computation, stationarity testing, and figure generation are available with the deposited research materials. More specifically, pandas library processes the ≈1.3M-row dataset (loading, deduplication, author-book aggregation) and applies the 5-point moving average to DTW distances. The *dtw* library computes the core RSP-DTW Stationarity Distance between each citation window and its reversed copy. *statsmodels* library provides the formal diagnostic framework - KPSS (stationarity) and ADF (unit root) - whose paired rejection forms the structural break signature. *scipy* library contributes peak detection (prominence ≥ 10) to locate DTW asymmetry spikes and percentile bootstrap for 95% confidence intervals on the mean peak spacing.

METHODOLOGY

Reversible Stochastic Process Framework

The present study -as aforementioned- introduces a quantitative approach grounded in the theory of Reversible. Stochastic Processes (RSP) for computing the *Stationarity Distance*, following Poulos (2016) and Poulos (2020).

Within this theoretical framework, a stochastic process-or time series-is defined as *strictly stationary* if its statistical properties remain invariant under time reversal. The complete theoretical and methodological workflow is summarized in Figure 1.

Formally, let denote the discrete-time citation raw-count process, where t indexes thecentury-based chronological sequence of authors. The process is said to be *time-reversible* if the finite-dimensional distributions satisfy $P(X_{t1},...,X_{tk}) = P(X_{tk},...,X_{t1})$ for all k and all ordered time indices $t_1 < \cdots < t_k$ (Kelly, 1979). Time reversibility has also been used as a statistical framework for identifying temporal asymmetry in empirical time series (Ramsey and Rothman, 1996). Equivalently, under the assumption of a stationary Markov chain, reversibility corresponds to the *detailed balance* condition $\pi_i P_{ij} = \pi_j P_{ji}$, where π is the stationary distribution and P is the transition matrix. In the present context, a strictly stationary citation process would exhibit statistical symmetry under time reversal: the distribution of citation counts when read forward in historical time should be indistinguishable from the distribution when read backward. Any deviation from this symmetry constitutes evidence of structural nonstationarity. Rather than testing for distributional equality globally (which is infeasible for a single finite realization), we operationalize the reversibility condition locally through a sliding-window DTW comparison between each sub-signal and its time-reversed counterpart.

Citation Signal Construction

To adapt this criterion to the present research and to detect *localized structural discontinuities* rather than relying on a global average measure, we implement a sliding-window methodology. The century-based ordering of the Church Fathers is treated as the discrete time axis n, generating a primary signal $x[n]$, where each value represents the total number of references for each book made by the n-th author. The algorithm therefore evaluates local, rather than global, departures from temporal symmetry.

Sliding-Window DTW Stationarity Distance

To measure the topological similarity and calculate the mathematical distance between the original and reversed signal segments, the Dynamic Time Warping (DTW) algorithm was selected.

For each step j, a sub-signal M is extracted of windows size W, its reversed counterpart N is constructed, and the DTW distance between them is calculated. A core segment of the Python implementation is presented in Table 2.

Unlike the simple Euclidean distance, which requires strict point-to-point correspondence and is consequently sensitive to small temporal shifts, DTW provides *elastic alignment*. This allows the algorithm to align data peaks that exhibit lag or distortion by compressing or stretching the time axis, thereby minimizing the total alignment cost (Sakoe and Chiba, 1978; Berndt and Clifford, 1994; Poulos, 2016).

The selection of DTW over alternative elastic distance measures was driven by two considerations specific to the patristic corpus. First, unlike Longest Common Subsequence (LCSS) or edit-distance-based measures which discard or penalize unmatched elements, DTW preserves the full chronological sequence of authors while permitting nonlinear alignment-a property useful for comparing doctrinal shifts that accelerate or decelerate across historical periods. Second, unlike simple cross-correlation which assumes a fixed linear lag structure, DTW accommodates the heterogeneous spacing of patristic authors, whose chronological distribution is irregular rather than uniformly sampled. Among the family of elastic measures, DTW was preferred because its accumulated cost matrix provides an interpretable local Stationarity Distance that quantifies differences between each window and its reversed version, rather than only testing for significance at a candidate break point.

Peak Detection and Periodicity Estimation

For reproducibility, all DTW distances were computed on the raw citation-count vectors within each sliding window using Python-based time-series tooling (Tavenard *et al.*, 2020). No amplitude normalization was applied, because the magnitude of citation concentration is treated as part of the historical signal rather than as noise. The DTW local cost was computed using the Euclidean distance. Because each observation consisted of a scalar citation count, the Euclidean distance reduces to the absolute difference between paired observations. The accumulated DTW distance, referred to as the Stationarity Distance, was recorded for each window. The DTW-distance sequence was then smoothed with a 5-point trailing moving average; peaks were identified using a minimum prominence threshold of 10 and a minimum peak distance of 1 index. The prominence threshold of 10 was chosen to retain peaks corresponding to meaningful departures from baseline symmetry while excluding high-frequency fluctuations attributable to single-author citation variation. The minimum peak distance of 1 index was selected to avoid merging genuinely distinct local maxima; because DTW distances are computed at each sliding step, adjacent indices can both represent genuine spikes at different positions along the author sequence.

Table 1: The original data structure compiled by P. Koutlemanis and S. Sakkos.

Book Code	Chapter	Verse	PG Column
59	1	04-5	644A
59	1	1	644A
59	5	1	644B
65	1	6	644B

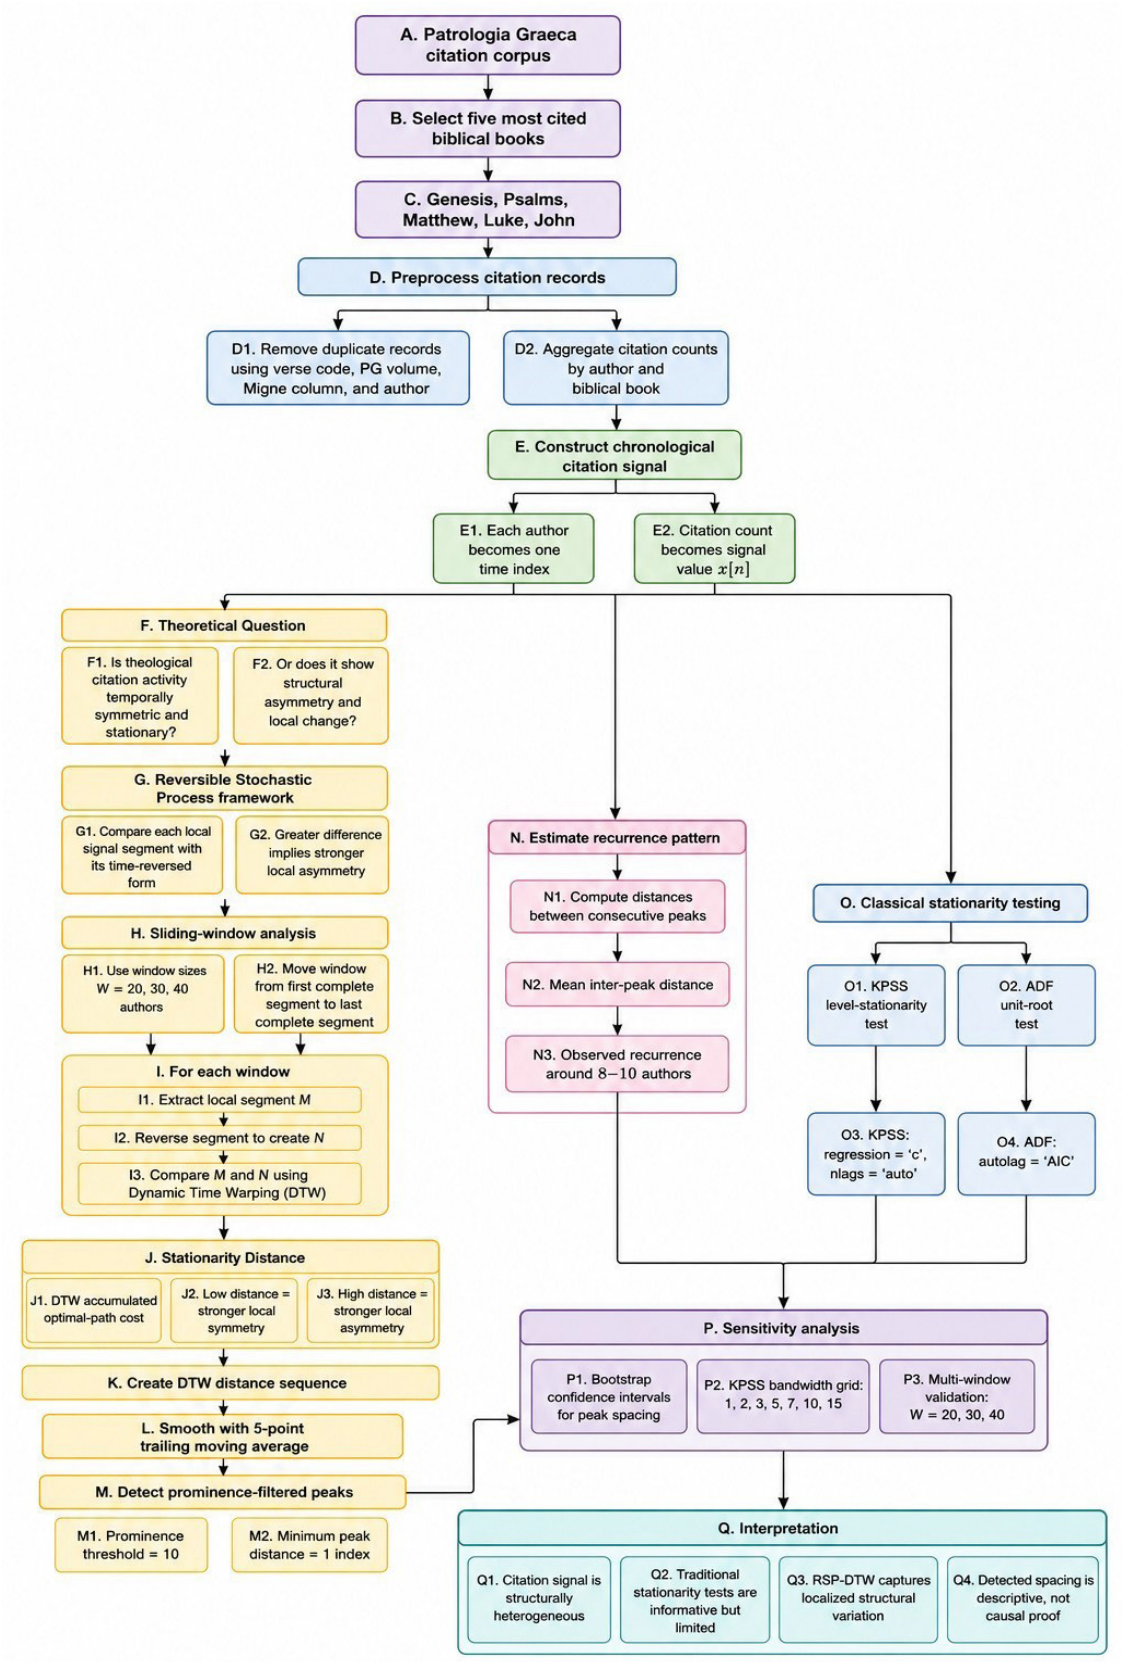


Figure 1: Flowchart of the theoretical and methodological workflow, from patristic citation data preprocessing to RSP-DTW analysis, stationarity testing, sensitivity validation, and interpretation.

Table 2: Python implementation of the Reversible Stochastic Processes (RSP), integrated with Dynamic Time Warping (DTW).

```
import numpy as np import pandas as pd
from scipy.signal import find_peaks from dtw import dtw

def get_sma_and_peaks(signal, window_size):
   A = []

# Sliding the window across the time series
  for j in range(0, len(signal) - window_size + 1, 1):
       M= signal[j: j + window_size]
       N = M[::-1]

# Time Reversal (RSP constraint)
        alignment = dtw(M, N, keep_internals=True, dist_method="euclidean')
        A.append(alignment.distance)

# 5-point trailing moving average
sma = pd.Series(A).rolling(window=5, min_periods=1).mean().values

# Peak detection with prominence filtering
peaks, _ = find_peaks(sma, prominence=10, distance=1)
return sma, peaks
```

Consecutive peak spacings were used to estimate local recurrence intervals. The same procedure was applied consistently across the five biblical-book series and across

window sizes of 20, 30, and 40 authors. Smaller windows (e.g., 10) tended to amplify author-level citation noise, while larger windows (e.g., 50) could smooth away local fluctuations.

To make the DTW calculation visually interpretable, Figure 2 presents a representative high-asymmetry window. The original citation segment is compared with its reversed counterpart, and the accumulated cost matrix displays the optimal warping path. Departures from the diagonal reference indicate nonlinear alignment between the forward and reversed sub-signals; such departures correspond to larger local Stationarity Distance values and therefore to stronger temporal asymmetry.

The computational procedure is summarized in the following discrete algorithmic steps:

Primary Signal Generation

The dataset is read by book, extracting the reference raw counting vector x for the total N authors. This process preserves the century-based chronological ordering used in the analysis.

Sliding Window Initialization

A fixed window size W (20, 30, or 40 authors) is defined, and the window is advanced one author at a time until the last complete window of length W has been evaluated.

Sub-signal Extraction and Time Reversal (RSP)

At each step j, a local sub-signal M of length W is extracted and reversed to form the sub-signal N, enabling the local RSP comparison.

Local DTW Execution

DTW calculates the alignment cost between M and N; the resulting accumulated optimal-path cost is recorded as the local Stationarity Distance.

Moving Average and Peak Detection

The sequence of Stationarity Distances is smoothed with a 5-point trailing moving average, and prominence-filtered peaks are used to estimate recurring intervals in the signal.

RESULTS

Descriptive Citation Patterns

The distribution of citations per biblical book is shown in Figure 3, while the top five books are presented in Table 3, alongside their citation counts and number of authors. The corresponding century-aware author-index signals used as DTW inputs are shown in Figure 4.

Bibliometric processing of the dataset identified both the distinct references to each biblical book and the unique contributing authors. These authors were arranged in chronological order,

and the resulting raw counting vector $x[n]$ captures the citation dynamics used as input for the DTW analysis.

RSP-DTW Peak Profiles and Peak-Spacing Results

The resulting smoothed DTW peak profiles are displayed in Figure 5. In the experiments, peak spacing for the top five books falls in the approximate range of 8 to 10 authors (Table 4), which is best interpreted as a descriptive recurrence rather than a precise periodic law.

KPSS and ADF Stationarity Tests

To systematically evaluate the intrinsic non-stationarity of the time series of Patristic hermeneutical citations and to overcome the constraints of purely observational measures, a formal hypothesis testing framework was applied to the discrete signal x[n] for each of the top five books. To evaluate these hypotheses, two standard econometric tests for stationarity and unit roots were applied to the series: the (ADF) test and the (KPSS) test (Table 5).

Table 3: Citation count and number of authors.

Bible Book	Citations	Authors
Genesis	48,969	259
Psalms	45,905	305
Matthew	66,199	321
Luke	37,187	281
John	44,732	289

Formal hypothesis models for the historical time series

KPSS Test

- **Null Hypothesis (H_0):** The theological interest sequence is stationary; its mean and variance remain constant over time, and no structural breaks significantly disrupt the series.
- **Alternative Hypothesis (H_1):** The sequence is non-stationary, exhibiting changes in mean, variance, or structural anomalies over time.

ADF Test

- **Null Hypothesis (H_0):** The theological interest sequence contains a unit root, i.e., it behaves as a **random walk**.
- **Alternative Hypothesis (H_1):** The sequence is stationary (no unit root), with deviations from a random walk.

In the implementation, KPSS was run as a level-stationarity test with a constant term (regression="c') and automatic Newey-West lag selection (nlags="auto'), while ADF was run with AIC-based lag selection (autolag="AIC'). The results of these analyses are summarized in Table 6.

As reported in Table 6, the KPSS and ADF tests deliver different but not incompatible inferences: the KPSS test rejects the null hypothesis of stationarity, while the ADF test rejects the null hypothesis of a unit root. Such divergence is common when a series contains structural breaks or regime shifts, because the assumptions underlying both tests are violated in different ways (Perron, 2006; Zivot and Andrews, 1992; Paparoditis and Politis, 2018). In this context, the results are consistent with a citation

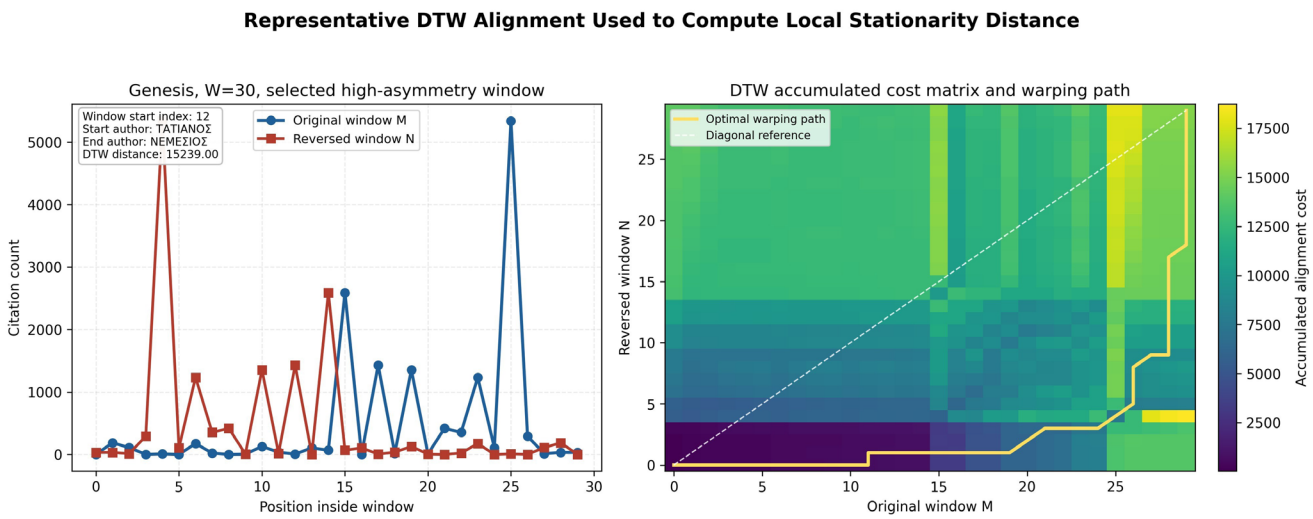


Figure 2: Representative DTW alignment for a high-asymmetry Genesis window (W =30). The left panel compares the original citation sub-signal with its time-reversed counterpart, while the right panel shows the accumulated DTW cost matrix and optimal warping path used to compute the local Stationarity Distance.

Table 4: Mean distance between non-stationarity peaks under different window sizes (*i*).

Bible Book	i =20	i =30	i =40	Mean
Genesis	7.76	6.91	7.96	7.54
Psalms	7.58	8.22	10.50	8.77
Matthew	8.25	9.53	10.40	9.39
Luke	7.44	8.33	9.43	8.40
John	8.03	7.26	9.80	8.37

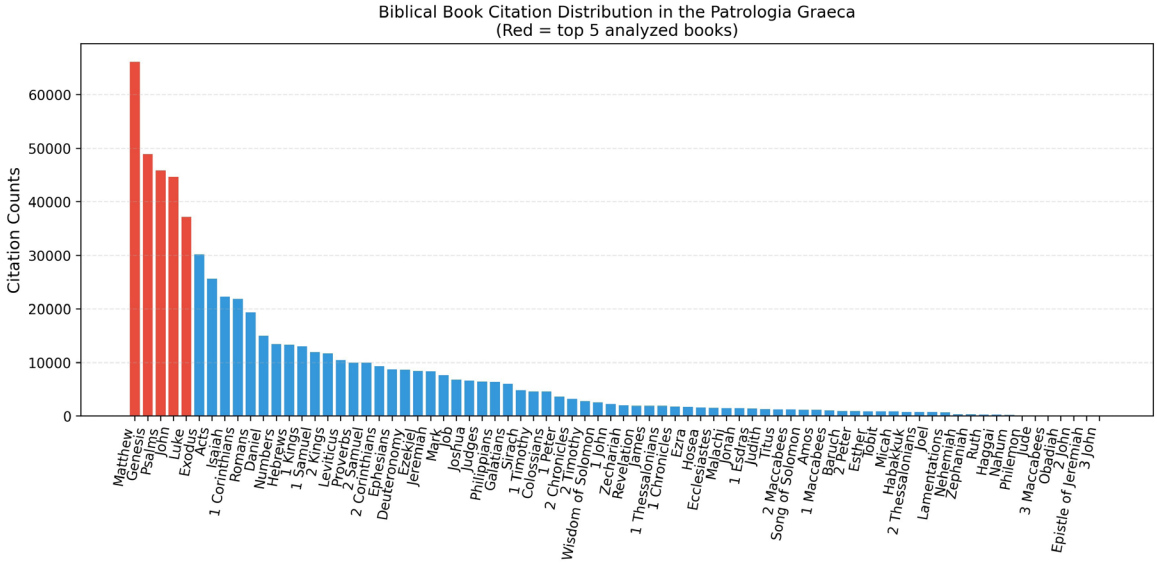


Figure 3: Citation count distribution per book.

process that is heterogeneous across time and strongly influenced by a small number of highly prolific authors, rather than with a single smooth stationary mechanism. The series is therefore best interpreted as break-sensitive and locally heterogeneous, not as evidence for a uniform stochastic process across the entire chronological corpus. Traditional non-parametric tests like KPSS are particularly sensitive to such magnitude variations, which limits their ability to capture the underlying dynamics of the series. By applying the Reversible Stochastic Process (RSP) metric in combination with Dynamic Time Warping (DTW), the analysis provides a complementary description of local asymmetry beyond the classical tests. The resulting DTW distances are consistent with departures from symmetry in the citation signals for Genesis, Psalms, Matthew, Luke, and John.

Robustness and Sensitivity Analysis

To evaluate the robustness of the reported findings against methodological parameterization, three complementary sensitivity tests were conducted: bootstrap confidence intervals on the DTW-derived periodicity, a grid search over KPSS bandwidth parameters, and multi-window DTW validation.

Bootstrap Confidence Intervals on DTW Periodicity

The peak-spacing distances $d_1, d_2, \ldots, d_k$ between consecutive prominence-filtered DTW peaks provide the main empirical indication of the recurrent author-spacing pattern. The sample mean $\bar{d}$ is subject to sampling variability: repeating the analysis on a slightly different author subset or with small data perturbations could shift the estimate. To quantify this uncertainty without assuming a parametric distribution, the percentile bootstrap (Efron, 1979) was applied with $B = 5000$ resamples. For each resample, k differences were drawn with replacement from the observed spacing vector and the mean was computed. The 95% confidence interval was taken as the 2.5th and 97.5th percentiles of the resulting bootstrap distribution.

KPSS Parameter Sensitivity

The KPSS test statistic depends on the Newey-West long-run variance estimator $\hat{\sigma}^2 = \hat{\gamma}_0 + 2\sum_{j=1}^{l} w(j,l)\,\hat{\gamma}_j$, where l is the bandwidth parameter and $w(j,l) = 1 - j/(l+1)$ is the Bartlett kernel. Here, $\hat{y}_j$ denotes the estimated autocovariance of the citation signal at lag j, with $\hat{y}_0$ corresponding to the ordinary variance. The bandwidth controls how many autocovariance terms enter the variance estimate. To verify that the rejection of the stationarity null hypothesis is not an artifact of a specific bandwidth choice, the KPSS test was re-run across the full range $l = 1, 2, 3, 5, 7, 10, 15$ for each of the five biblical books.

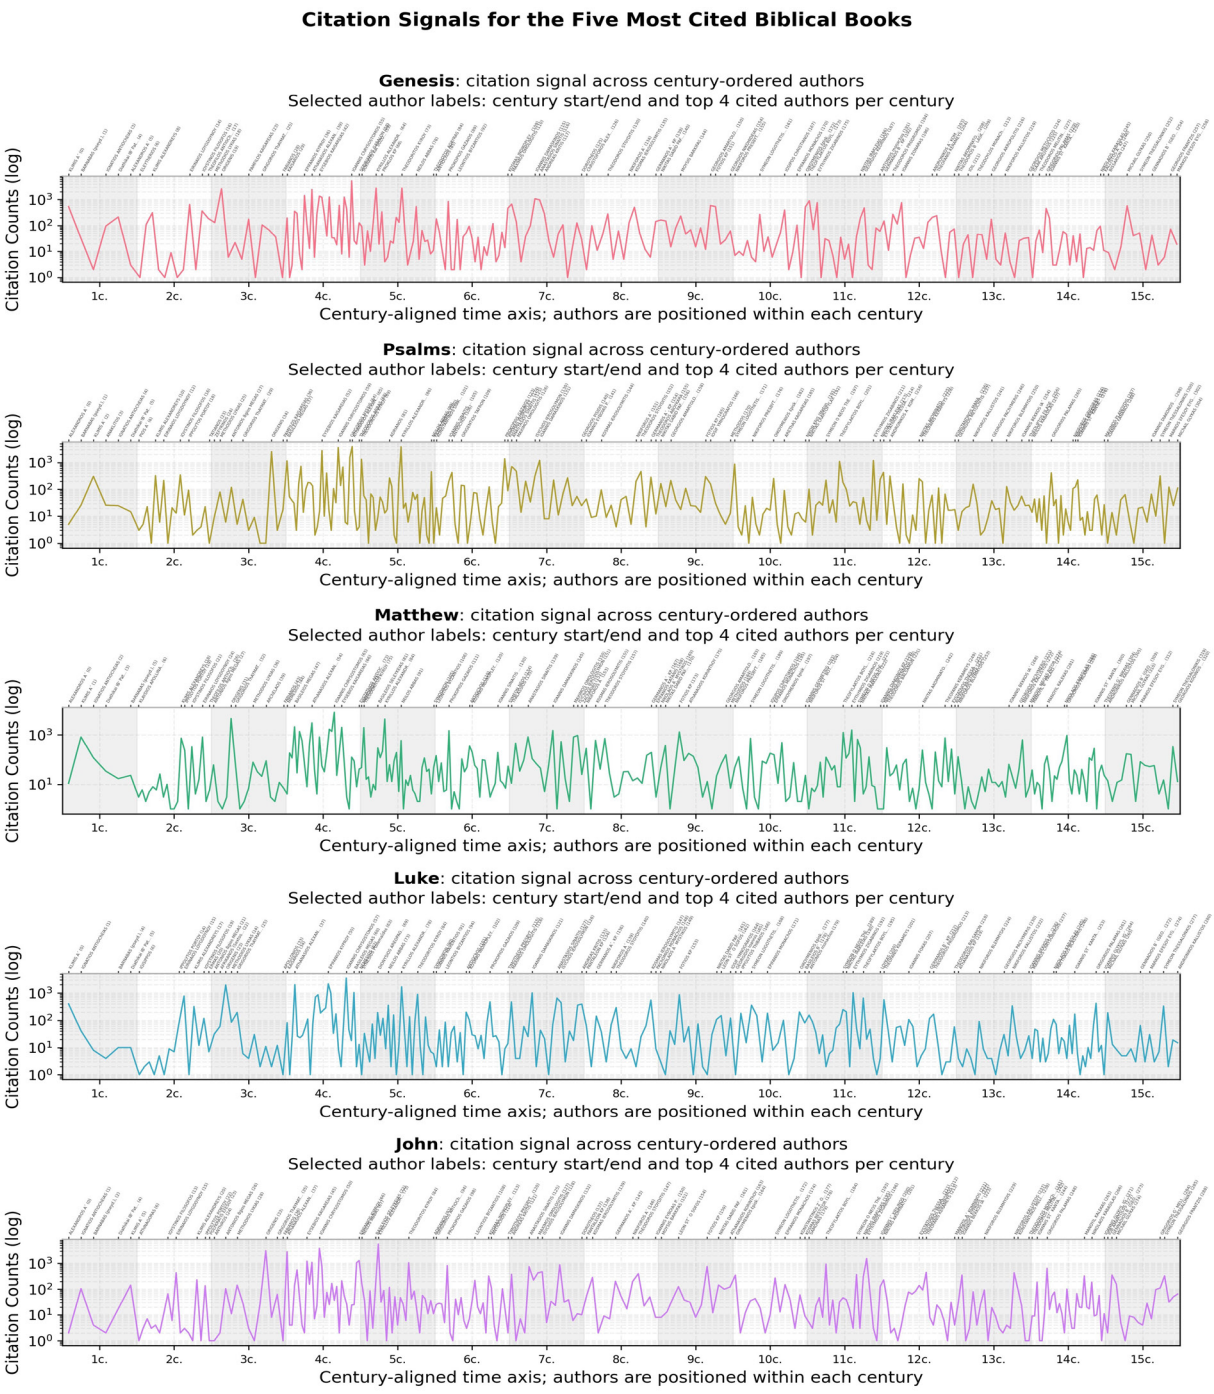


Figure 4: Authors signal for the top five books.

Table 5: Most active authors per biblical book.

Author	Genesis	Psalms	Matthew	Luke	John
John Chrysostom	5,339	3,575	8,121	3,532	3,874
Cyril of Alexandria	2,788	3,831	4,338	1,694	5,652
Origen	2,589	2,538	4,555	1,981	3,133
Athanasius of Alexandria	2,435	3,900	3,046	2,025	2,868

Table 6: Results of KPSS and ADF stationarity and unit-root tests.

Book	KPSS	*p*	KPSS (α =0.05)	ADF	*p*	ADF (α =0.05)
Genesis Psalms	1.05 0.98	<0.01 <0.01	Reject stationarity Reject stationarity	−3.26	0.02	Reject unit root Reject unit root
Matthew	0.74	<0.01	Reject stationarity	−4.39 −4.23	<0.0.1 <0.0.1	Reject unit root
Luke	0.98	<0.01	Reject stationarity	−4.42	<0.0.1	Reject unit root
John	0.73	0.01	Reject stationarity	−16.67	<0.0.1	Reject unit root

Table 7: Bootstrap sensitivity of DTW periodicity across window sizes.

Book	*W* =20	*W* =30	*W* =40	Overall mean
Genesis	7.76 (6.31, 9.38)	6.91 (5.47, 8.47)	7.96 (6.19, 10.04)	7.54
Psalms	7.58 (6.33, 9.03)	8.22 (6.56, 10.00)	10.50 (7.92, 13.08)	8.77
Matthew	8.25 (6.69, 9.92)	9.53 (7.40, 11.90)	10.40 (8.24, 12.64)	9.39
Luke	7.44 (6.47, 8.47)	8.33 (6.85, 9.81)	9.43 (7.22, 12.00)	8.40
John	8.03 (6.58, 9.45)	7.26 (5.88, 8.76)	9.80 (7.20, 12.80)	8.37

Table 8: KPSS *p*-values across bandwidth parameters.

l	Genesis	Psalms	Matthew	Luke	John
1	0.010*	0.010*	0.010*	0.010*	0.010*
2	0.010*	0.010*	0.010*	0.010*	0.010*
3	0.010*	0.010*	0.010*	0.010*	0.010*
5	0.010*	0.010*	0.010*	0.010*	0.014*
7	0.010*	0.010*	0.012*	0.010*	0.017*
10	0.010*	0.017*	0.018*	0.010*	0.022*
15	0.012*	0.025*	0.027*	0.021*	0.034*

* $p<0.05$ (rejects stationarity).

Multi-Window DTW Validation

The sliding-window DTW methodology requires selecting a window width W. To confirm that the detected periodicity is a structural property of the data rather than an artifact of window parameterization, the full RSP-DTW pipeline was executed at three window sizes: $W = 20$, $W = 30$, and $W = 40$ authors.

Summary of Sensitivity Results

Table 7 presents the mean inter-peak distances with 95% bootstrap confidence intervals across all three window sizes. Under the randomized within-century ordering used for the final analysis, the periodicity converges to approximately 8-10 authors for the top five books regardless of parameterization. The bootstrap CIs overlap substantially (approximately 5-13 authors), indicating that the recurrent author-spacing pattern is relatively stable across the tested settings.

For clarity, Table 7 therefore serves two complementary validation purposes: it reports the multi-window DTW validation across $W = 20$, $W = 30$, and $W = 40$, and it also reports the bootstrap uncertainty intervals associated with each window-specific periodicity estimate.

The KPSS level-stationarity test rejects the null hypothesis ($p<0.05$) for all five books across all tested bandwidths from $l = 1$ through $l = 15$ (see Table 8). The small numerical differences between the p-values in Table 8.

those reported in Table 6 arise because the latter uses automatic Newey-West bandwidth selection (nlags="auto'), which chooses a data-dependent integer lag that may not appear among the preselected grid points in Table 8. Thus, the sensitivity analysis supports the broader evidence of local non-stationarity and indicates that the KPSS rejection is not driven by a single bandwidth choice.

In summary, the sensitivity analyses suggest that the main findings are not an artifact of a single window size or bandwidth choice.

Classical spectral methods and change-point detection provide useful complementary perspectives, but they address different questions from the RSP-DTW framework used here. Spectral

methods are most appropriate for approximately uniformly sampled signals and for detecting sinusoidal or frequency-like regularities, whereas the present citation signal is indexed by chronologically ordered authors with uneven historical spacing and highly unequal citation magnitudes. Similarly, change-point methods are designed to identify discrete regime boundaries (Bai and Perron, 2003; Lavielle, 2005), while the present analysis aims to quantify local departures from time-reversal symmetry across the full author sequence. For this reason, spectral and change-point approaches are treated as complementary future extensions rather than replacements for the RSP-DTW Stationarity Distance analysis.

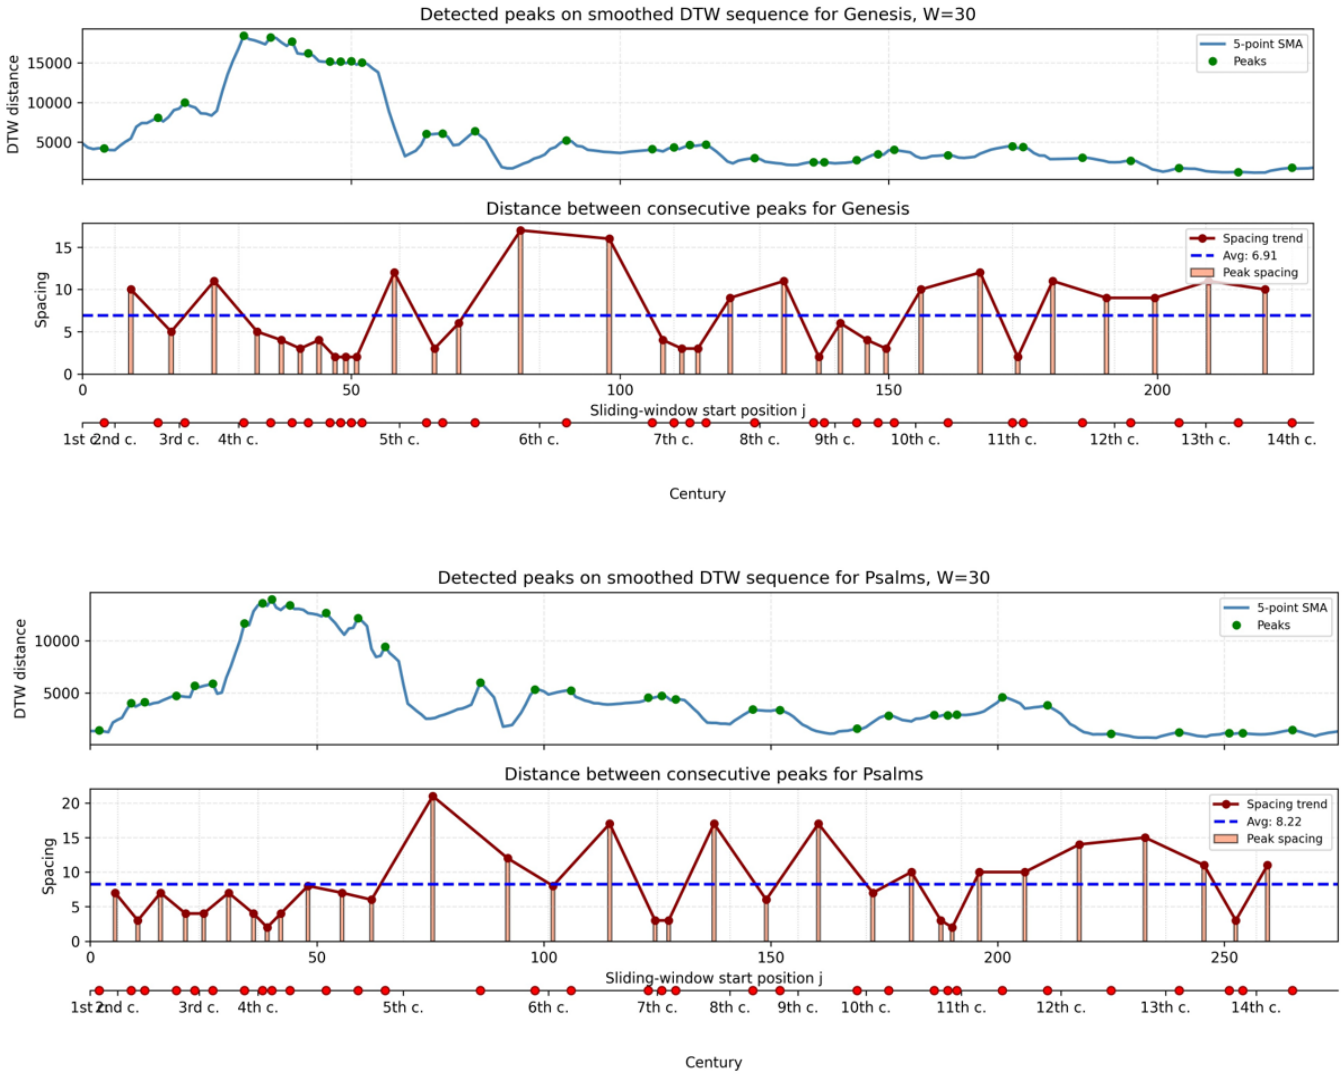
DISCUSSION

Interpretation of Local Non-Stationarity

This study introduced a scientometric approach to the quantitative analysis of Patristic literature, transforming the citations of textual volume of the Patrologia Graeca into a discrete historical time series. The primary objective was to examine structural continuity and stationarity in the hermeneutical tradition concerning Genesis, Psalms, Matthew, Luke, and John.

The experimental results suggest that the evolution of theological thought is not well described by a simple linear or stationary process. Instead, the citation signal shows several pronounced structural asymmetries, especially in the periods associated with John Chrysostom, Cyril of Alexandria, and Athanasius of Alexandria (Table 5). Their exegetical corpus is better understood as a dominant citation concentration than as an isolated statistical outlier. In this setting, traditional tests such as ADF and KPSS should be interpreted cautiously, because extreme variance shifts can reduce their usefulness. Within this methodological gap, the application of the Reversible Stochastic Process (RSP) integrated with Dynamic Time Warping (DTW) via a sliding window proved useful for describing local changes in the series.

A notable outcome of the present research is that the Moving Average peaks are not randomly distributed. Instead, they show a recurring spacing of roughly 8-10 authors. This pattern is compatible with the idea that doctrinal and interpretive emphasis



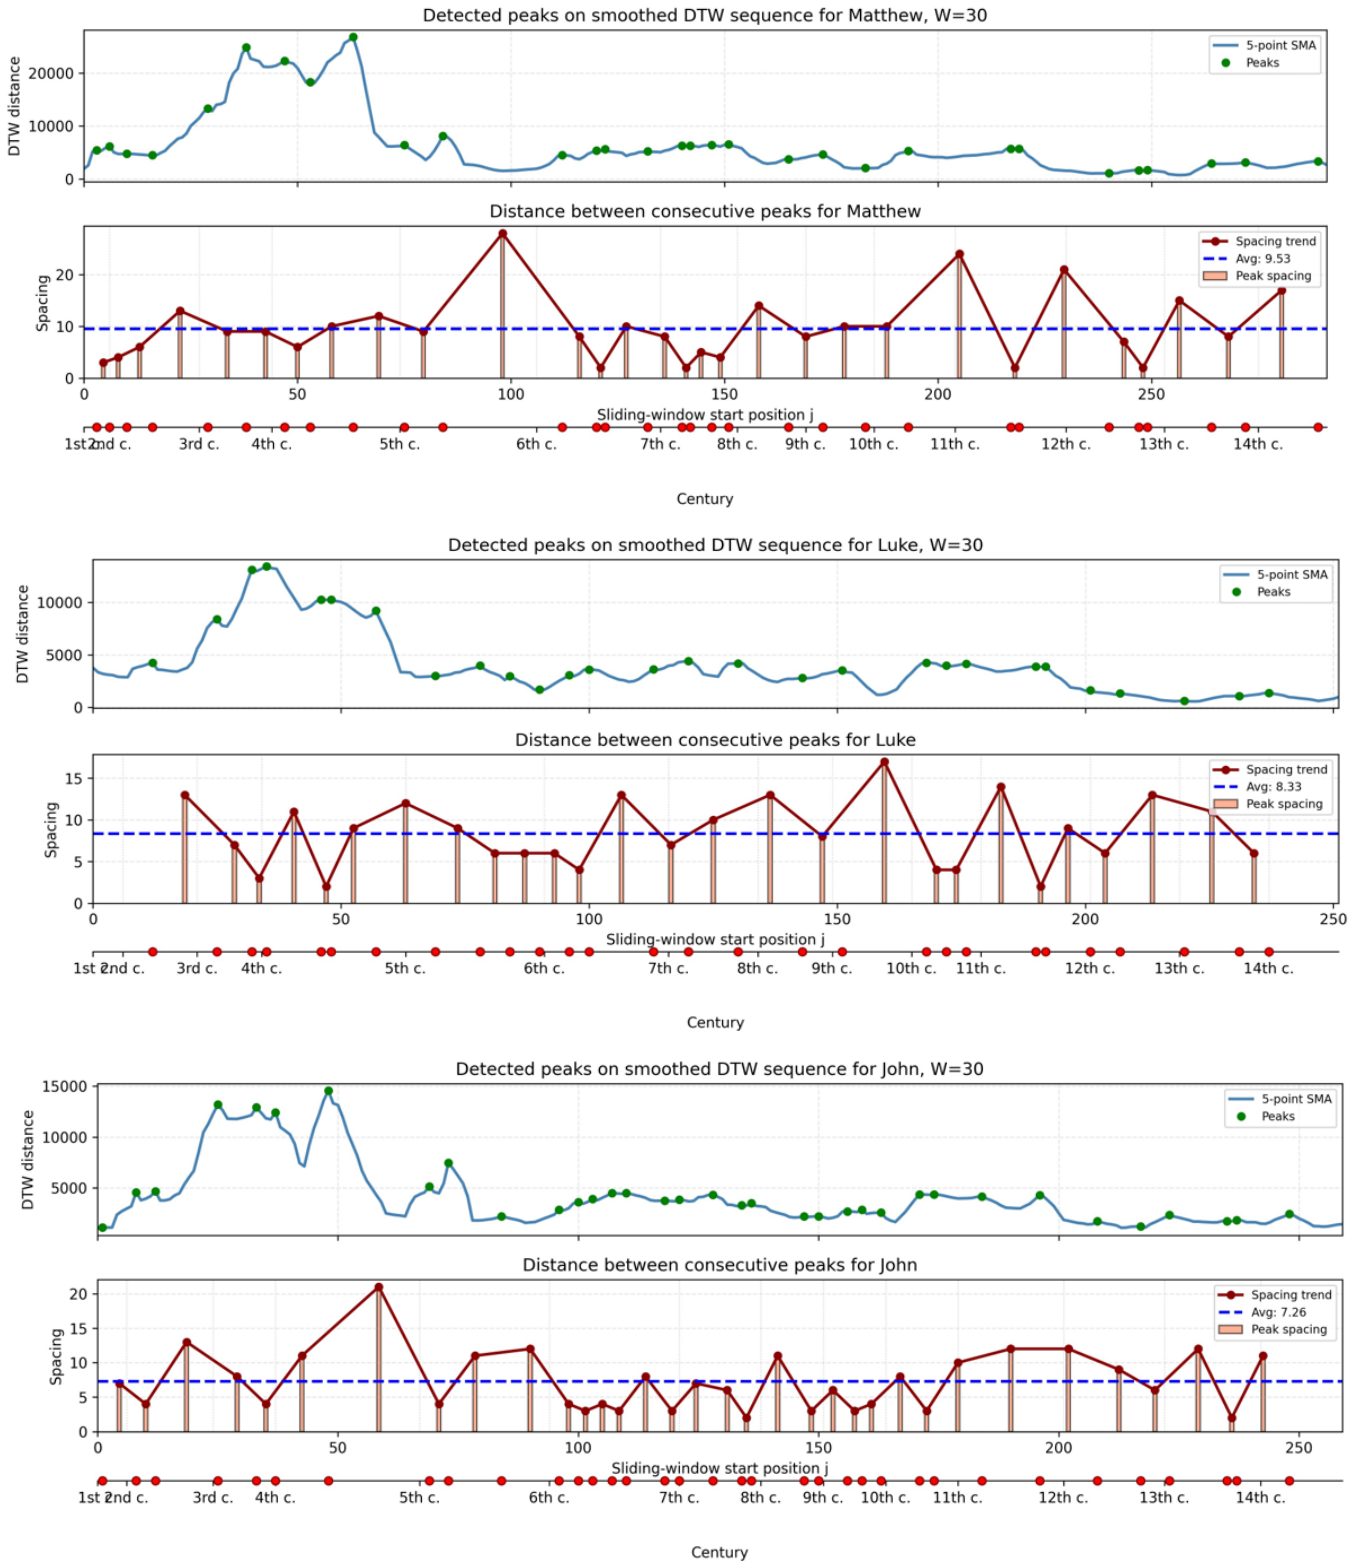


Figure 5: Peaks of the smoothed computed DTW signal for the top five books (W =40).

changes unevenly over time, although the present analysis does not establish a causal historical mechanism. The observed spacing may be compared with broader theoretical discussions of paradigm change in theology, including the volume edited by Küng and Tracy (1989).

Also, the RSP-DTW framework is not inherently domain-specific: it can be applied to any chronologically ordered citation corpus, including legal codices, philosophical traditions, or scientific literature, provided that a curated author-to-reference mapping is available.

Historical and Theological Context

Simultaneously, the statistical prominence of John Chrysostom, Cyril of Alexandria, and Athanasius of Alexandria in the time series is consistent with qualitative analyses by patristic scholars such as Robert C. Hill (1999) and Frances Young (1997). These scholars emphasize that Chrysostom's exegetical work, including his homiletic engagement with Genesis, represented a major development within the Antiochene tradition, both quantitatively and qualitatively. At the same time, the Church faced significant doctrinal and institutional crises, beginning with the Arian controversy addressed at the First Council of Nicaea (325) and vigorously opposed by Athanasius of Alexandria, and continuing through the Meletian Schism (360-415) in Antioch.

Later Christological debates involved Cyril of Alexandria, culminating in the Council of Ephesus (431) against Nestorianism and the Council of Chalcedon (451), which defined the two natures of Christ. Subsequent controversies included the Monothelite crisis (c. 630-681), opposed by Maximus the Confessor and resolved at the Third Council of Constantinople (680-681), and the Iconoclastic controversy (726-843), defended by John of Damascus and ultimately settled at the Second Council of Nicaea (787) and in 843 with the Triumph of Orthodoxy. Even after this period, the Church faced political and jurisdictional tensions, including the Photian Schism (863-867), involving Photius I of Constantinople and the papacy, as well as disputes over the Christianization of the Slavs, which raised questions of liturgical language and ecclesiastical authority. Collectively, these crises suggest that from the fourth to the ninth centuries, the development of doctrine in the Byzantine Church was shaped by intellectual and theological debates as well as by political and institutional struggles.

From the eleventh to the fifteenth century, the same pattern of doctrinal, liturgical, and institutional tension continued in later Byzantine and Orthodox history, as summarized visually in Figure 6. The Great Schism of 1054 crystallized disputes over papal primacy, the *Filioque*, and ecclesiology. In the twelfth century, the Byzantine Eucharistic controversy of 1156-1157 focused on the nature of Christ's sacrifice in the Eucharist. The thirteenth century then intensified East-West tensions through the Fourth Crusade and the Sack of Constantinople in 1204, followed by the Council of Lyons II in 1274 and the anti-unionist crisis of 1274-1282, in which resistance to ecclesiastical union with Rome became a major marker of Orthodox identity. In the fourteenth century, the Hesychast controversy and the Palamite councils of 1341, 1347, and 1351 shifted attention toward the distinction between divine essence and energies, the theology

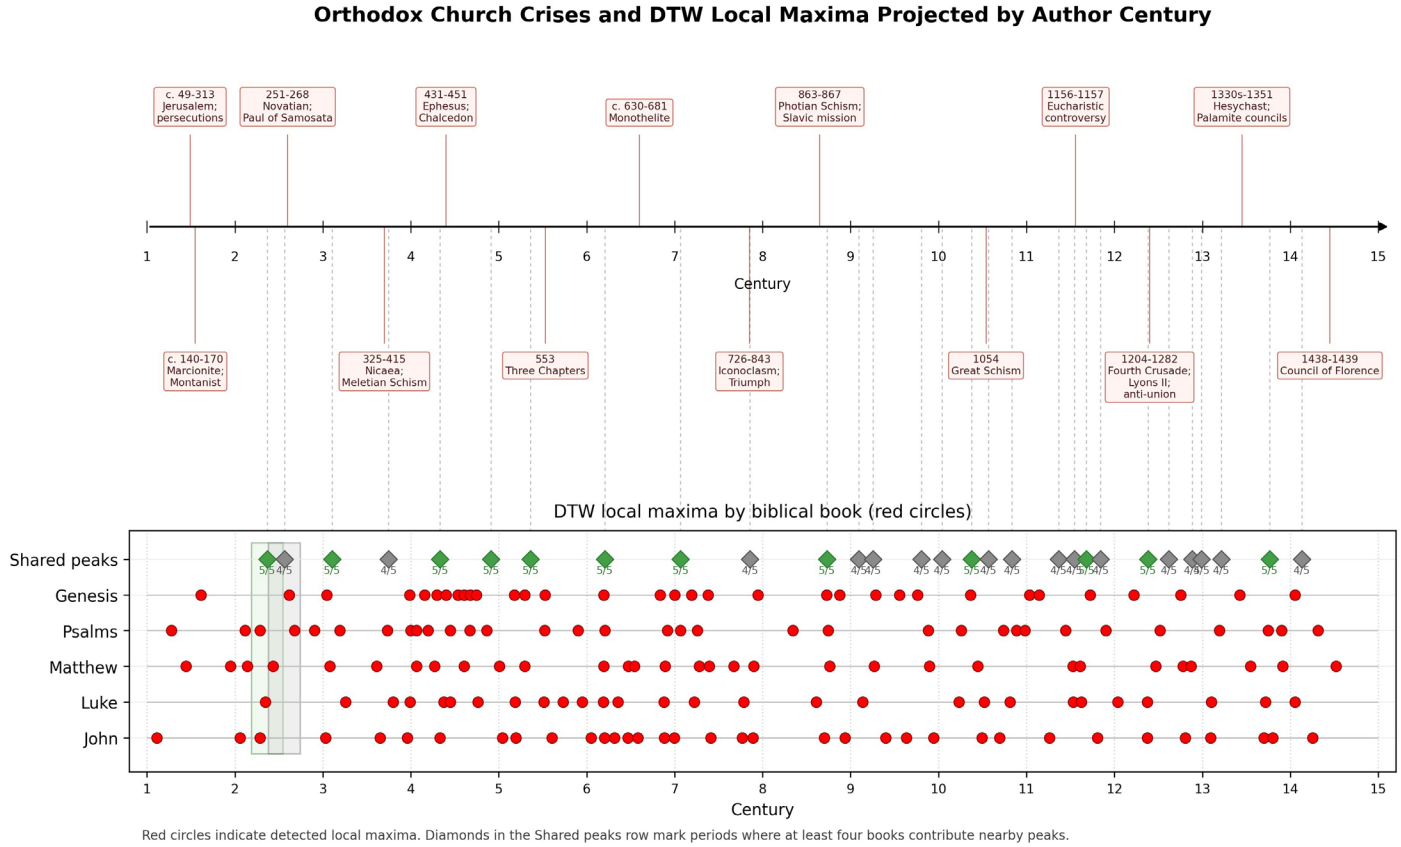


Figure 6: Combined temporal axis of major doctrinal and institutional crises affecting the Orthodox Church from the first to the fifteenth century, with DTW local maxima projected by author century for Genesis, Psalms, Matthew, Luke, and John. Green diamonds indicate periods where all five books contribute nearby local maxima, while gray diamonds indicate periods where four of the five books contribute.A pattern emerges linking increases in citations to periods of crisis in ecclesiastical history.

of prayer, and the experience of divine light. Finally, the Council of Florence in 1438- 1439 represented another attempted union with Rome shortly before the final fall of Constantinople. These later crises show that theological development remained bound to institutional pressure, political rupture, and debates over ecclesial authority well beyond the patristic peak centuries detected in the citation signal.

LIMITATIONS AND FUTURE WORK

Several limitations of the present study suggest directions for future research. A first limitation concerns the chronological ordering of authors within individual centuries. Because precise birth dates or floruit dates are not available for all authors in the corpus, the chronological ordering of authors within the same century cannot be established with full certainty. To avoid imposing an artificial intra-century sequence based on citation frequency or alphabetical order, the primary chronological constraint was therefore defined at the century level. Within each century, authors were randomly ordered using a fixed seed, so that the procedure remained reproducible while acknowledging uncertainty in finer temporal placement. This choice prevents highly cited authors from determining their own position in the signal and avoids treating alphabetic order as historical evidence. The resulting author-ordered citation series should therefore be interpreted as a century-constrained chronological signal, with intra-century ordering treated as an uncertainty rather than as a known historical fact.

Second, the DTW-based sliding-window methodology has quadratic time complexity ($O(W^2)$ per window), which is negligible for the current corpus of approximately 300 authors but would require acceleration through lowerbounding techniques (Keogh and Ratanamahatana, 2005) or scalable subsequence-search methods (Rakthanmanon *et al*., 2012) when applied to all books of Scripture cited throughout the full *Patrologia Graeca* (170 volumes) or to larger heterogeneous bibliographic databases.

Third, the current analysis treats each biblical book independently as a univariate citation signal. Future work could extend the framework to multivariate time-series representations or co-citation network models, capturing interactions between books and authors through time-evolving graph representations.

Future work should also compare the RSP-DTW profile with complementary signal-processing diagnostics, including spectral or pseudoperiodic analysis and formal change-point detection, to evaluate whether the same intervals of local asymmetry are recovered by methods based on frequency structure or discrete regime boundaries.

Finally, the dataset inherits the survivorship bias of the *Patrologia Graeca* compilation; texts not preserved or selected by Migne are absent, and future work could assess sensitivity by supplementing the corpus with non-PG patristic fragments.

CONCLUSION

This study demonstrated that the History of Dogma and Patrology can be examined through computational and scientometric methods without reducing theological interpretation to purely numerical description. By transforming biblical references in the Patrologia Graeca into author-ordered citation signals, the paper proposed a reproducible framework for identifying temporal structure, local asymmetry, and discontinuities in patristic exegetical activity.

The results show that the citation history of Genesis, Psalms, Matthew, Luke, and John is not adequately described as a simple stationary process. The KPSS and ADF tests indicate that the signals contain both stationary components and structural disturbances, while the RSP-DTW analysis reveals localized departures from time-reversal symmetry. Across multiple window sizes, the detected DTW peak spacings repeatedly fall near an interval of approximately 8- 10 authors. This pattern should not be interpreted as a strict historical law, but as a descriptive recurrence suggesting that major shifts in exegetical intensity appear at uneven but measurable intervals in the author sequence.

Substantively, the findings suggest that the development of patristic interpretation followed a pattern of sustained continuity punctuated by episodes of intensified theological and literary production associated with major crises in ecclesiastical history. The prominence of John Chrysostom, Cyril of Alexandria, Athanasius of Alexandria, and Origen supports qualitative scholarship that treats them as central to Christian exegetical tradition. The computational results complement historical-theological interpretation by locating citation concentration and structural asymmetry.

Methodologically, RSP combined with DTW provides a useful tool for chronologically ordered textual corpora. It detects local changes and compares forward and reversed citation structures, making it suitable for uneven author distributions, strong citation concentrations, and irregular temporal spacing.

Future work may extend the analysis to more biblical books or the full Patrologia Graeca, construct multivariate citation signals, and compare RSP-DTW with spectral, change-point, network, or topic-modeling methods. The study thus offers both a focused analysis and a broader path for integrating patristics, scientometrics, and signal processing.

ABBREVIATIONS

PG: Patrologia Graeca; **ADF:** Augmented Dickey-Fuller; **KPSS:** Kwiatkowski-Phillips-Schmidt-Shin; **RSP:** Reversible Stochastic Process; **DTW:** Dynamic Time Warping; **DGPC:** Digital Greek Patristic Catena; **LCSS:** Longest Common Subsequence; **AIC:** Akaike Information Criterion.

CONFLICT OF INTEREST

The authors declare that there is no conflict of interest.

REFERENCES

Bai, J., & Perron, P. (2003). Computation and analysis of multiple structural change models. Journal of Applied Econometrics, 18(1), 1-22. https://doi.org/10.1002/jae.659

Berndt, D. J., & Clifford, J. (1994). Using dynamic time warping to find patterns in time series. In Proceedings of the AAAI-94 Workshop on Knowledge Discovery in Databases (pp. 359–370). AAAI Press. https://cdn.aaai.org/Workshops/1994/WS-94-03/WS94-03-031.pdf

Berry, D. M. (Ed.). (2012). Understanding digital humanities. Palgrave Macmillan. https://doi.org/10.1057/9780230371934

BiblIndex. (n.d.). BiblIndex: Index of biblical references in Christian literature. Sources Chrétiennes–HiSoMA. Retrieved October 9, 2025, from https://www.biblindex.org/

Busa, R. (1980). The annals of humanities computing: The index Thomisticus. Computers and the Humanities, 14(2), 83-90. https://doi.org/10.1007/BF02403798

Busetti, F., & Taylor, A. R. (2003). Variance shifts, structural breaks, and stationarity tests. Journal of Business &Economic Statistics, 21(4), 510-531. https://doi.org/10.1198/073500103288619269

Cavaliere, G., & Xu, F. (2014). Testing for unit roots in bounded time series. Journal of Econometrics, 178(2), 259-272. https://doi.org/10.1016/j.jeconom.2013.08.026

Centre of Patristic Editions. (1985-2010). Photographic reprint of the Patrologia Graeca [Catalog]. Athens, Greece:Centre of Patristic Editions.

DeJong, D. N., Nankervis, J. C., Savin, N. E., & Whiteman, C. H. (1992). The power problems of unit root tests in time series with autoregressive errors. Journal of Econometrics, 53(1-3), 323-343. https://doi.org/10.1016/ 0304-4076(92)90090-E

Chrysostom, J. (1986). Homilies on Genesis, 1–17 (R. C. Hill, Trans.; The Fathers of the Church, Vol. 74). Catholic University of America Press. https://doi.org/10.2307/j.ctt32b2jg

Jockers, M. L. (2013). Macroanalysis: Digital methods and literary history. University of Illinois Press. https://doi.org/10.5406/illinois/9780252037528.001.0001

Kelly, F. P. (1979). Reversibility and stochastic networks. Wiley.

Keogh, E., & Ratanamahatana, C. A. (2005). Exact indexing of dynamic time warping. Knowledge and Information Systems, 7(3), 358-386. https://doi.org/10.1007/s10115-004-0154-9

Küng, H., & Tracy, D. (Eds.).(1989). Paradigm change in theology: A symposium for the future. Crossroad Publishing Company.

Lavielle, M. (2005). Using penalized contrasts for the change-point problem. Signal Processing, 85(8), 1501-1510.https://doi.org/10.1016/j.sigpro.2005.01.012

Moretti, F. (2005). Graphs, maps, trees: Abstract models for a literary history. Verso.

Paparnakis, A. (2014). Digital Greek Patristic Catena (DGPC): A brief presentation [Paper presentation]. International Workshop on Computer Aided Processing of Intertextuality in Ancient Languages, Villeurbanne, Lyon, France.

Paparnakis, A., & Domouchtsis, C. (2017). Digital Greek patristic catena (DGPC), a brief presentation. Journal of Data Mining & Digital Humanities, 2017, Project presentations. https://doi.org/10.46298/jdmdh.4001

Paparoditis, E., & Politis, D. N. (2018). The asymptotic size and power of the augmented Dickey-Fuller test for a unit root. Econometric Reviews, 37(9), 955-973. https://doi.org/10.1080/00927872.2016.1178887

Perron, P. (2006). Dealing with structural breaks. In T. C. Mills & K. Patterson (Eds.),Palgrave handbook of econometrics, 1(2), 278-352. Palgrave Macmillan. https://doi.org/10.1057/9780230280830_32

Poulos, M. (2016). Determining the stationarity distance via a reversible stochastic process. PLoS ONE, 11(10), e0164110. https://doi.org/10.1371/journal.pone.0164110

Poulos, M. (2020). Assessing the stationarity of web bounce rates using dynamic time warping and reversible stochastic processes. Expert Systems, 37(4), e12530. https://doi.org/10.1111/exsy.12502

Ramsey, J. B., & Rothman, P. (1996). Time irreversibility and business cycle asymmetry. Journal of Money, Credit and Banking, 28(1), 1-21. https://doi.org/10.2307/2077963

Rakthanmanon, T., Campana, B., Mueen, A., Batista, G., Westover, B., Zhu, Q., *et al.* (2012). Searching and mining trillions of time series subsequences under dynamic time warping. In Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 262-270). ACM. https:// doi.org/10.1145/2339530.2339576

Sakoe, H., & Chiba, S. (1978). Dynamic programming algorithm optimization for spoken word recognition. IEEE transactions on acoustics, speech, and signal processing, 26(1), 43-49.

Transactions on Acoustics, Speech, and Signal Processing, 26(1), 43-49. https://doi.org/10.1109/TASSP.1978.1163055

Schreibman, S., Siemens, R., & Unsworth, J. (Eds.).(2004). A companion to digital humanities. Blackwell. https: //doi.org/10.1002/9780470999875

Tavenard, R., Faouzi, J., Vandewiele, G., Divo, F., Androz, G., Holtz, C., *et al.* (2020). Tslearn, a machine learning toolkit for time series data. Journal of Machine Learning Research, 21(118), 1-6.

Varthis, E., Tzanavaris, S., Giarenis, I., Papavlasopoulos, S., Drakakis, M., & Poulos, M. (2022). Semantic enrichment on large scanned collections through their "satellite texts": the paradigm of Migne's Patrologia Graeca.

Information Discovery and Delivery, 50(2),217-234.

Young, F. M. (1997). Biblical exegesis and the formation of Christian culture. Cambridge University Press.

Zivot, E., & Andrews, D. W. K. (1992). Further evidence on the great crash, the oil-price shock, and the unit-root hypothesis. Journal of Business & Economic Statistics, 10(3), 251-270. https://doi.org/10.1080/07350015.1992.10509904

Cite this article: Varthis E, Domouchtsis K, Chrysostomidis I, Poulos M. A Scientometric Approach to Theological Continuity: Measuring Stationarity in the Patrologia Graeca via Dynamic Time Warping. J Scientometric Res. 2026;15(3):860-73.