



HAL
open science

Temporal dynamics of emotions in Italian online soccer fandoms

Salvatore Citraro, Giovanni Mauro, Emanuele Ferragina

► **To cite this version:**

Salvatore Citraro, Giovanni Mauro, Emanuele Ferragina. Temporal dynamics of emotions in Italian online soccer fandoms. EPJ Data Science, 2026, <10.1140/epjds/s13688-026-00640-7>. <hal-05597612>

HAL Id: hal-05597612

<https://sciencespo.hal.science/hal-05597612v1>

Submitted on 20 Apr 2026

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Temporal dynamics of emotions in Italian online soccer fandoms

Received: 13 June 2025

Accepted: 5 March 2026

Published online: 15 April 2026

Cite this article as: Citraro, S., Mauro, G., Ferragina, E. Temporal dynamics of emotions in Italian online soccer fandoms. *EPJ Data Science* (2026). <https://doi.org/10.1140/epjds/s13688-026-00640-7>

Salvatore Citraro, Giovanni Mauro, Emanuele Ferragina

We are providing an unedited version of this manuscript to give early access to its findings. Before final publication, the manuscript will undergo further editing. Please note there may be errors present which affect the content, and all legal disclaimers apply.

If this paper is publishing under a Transparent Peer Review model then Peer Review reports will publish with the final article.

Temporal Dynamics of Emotions in Italian Online Soccer Fandoms

Salvatore Citraro^{1*}, Giovanni Mauro² and Emanuele Ferragina³

^{1*}CNR-ISTI, Via Giuseppe Moruzzi, 1, Pisa, Italy.

²Scuola Normale Superiore, Pisa, Italy.

³SciencesPo, Rue Saint Guillaume, 27, Paris, France.

*Corresponding author(s). E-mail(s): salvatore.citraro@isti.cnr.it;

Contributing authors: giovanni.mauro@sns.it;

emanuele.ferragina@sciencespo.fr;

Abstract

This study investigates the emotional dynamics of Italian soccer fandoms through computational analysis of user-generated content from official Instagram accounts of 83 teams across Serie A, Serie B, and Lega Pro during the 2023-24 season. By applying sentiment analysis to fan comments, we extract temporal emotional patterns and identify distinct clusters of fan bases with similar pre-season expectations. Drawing from complex systems theory, we characterize joy as displaying anti-bursty temporal distributions, while anger is marked by pronounced bursty patterns. Our analysis reveals significant correlations between these emotional signals, preseason expectations, socioeconomic factors, and final league rankings. In particular, the burstiness metric emerges as a meaningful correlate of team performance, **as supported by stepwise regression**. These findings offer novel insights into the relationship between fan emotional expression and team outcomes, suggesting potential avenues for research in sports analytics, social media dynamics, and fan engagement studies.

Keywords: time series, burstiness, emotions, soccer

1 Introduction

Computational methods in social sciences have revolutionized our understanding of human behavior, offering new insights into phenomena such as the structure of social

networks and large-scale collective changes [1–4]. An important aspect of this transformation is the capability to study sentiment and emotion at scale, allowing researchers to trace collective emotional dynamics by identifying group-based and recurring narratives [5–8]. In this context, soccer fandoms provide a unique and rich environment to explore group-based emotions [9, 10]. Whilst computational methods have been successfully applied in understanding soccer from the perspective of player performance, game outcomes, injuries, scoring probability [11–14], much less attention has been paid to the emotional experiences of fans. Recent studies [15–23] have started to investigate this domain, but a significant gap remains in the application of rigorous computational methods to fully capture the nature and dynamics of emotions in online soccer fandoms.

Social media have transformed the way fans engage, reshaping traditional soccer cultures [24] by shifting communities from physical to digital spaces and, in the process, deterritorializing fandoms [25]. Yet, it remains to be established the extent to which digital spaces are different from traditional and physical spaces. Online platforms might help expand the reach and prestige of top teams (as shown by the growing globalisation of merchandising and the attention European top clubs receive in countries where the interest for football was not considerable in the recent past) [25]; even so, the nature of fan identity, especially within lower leagues, remains strongly tied to local and community-based features [26]. Moreover, it is unclear whether online fan interactions reflect traditional emotional dynamics related to sport dynamics, which are characterized by intense post-game discussions involving enthusiasm, frustrations, opinions, *what-ifs* [27]. This digital shift raises questions about how online interactions compare with the long-standing traditions of soccer fandoms.

In this work, we explore the emotional dynamics of digital soccer fandoms within the Italian league system, including Serie A, Serie B, and Lega Pro, during the 2023–24 season. We aim to capture the complex patterns underlying the social behavior of Italian fans online, by answering the following research questions: *(i)* Do different pools of fans exhibit similar patterns in how they emotionally support their teams?; *(ii)* Can these eventual patterns, together with preseason expectations and other deeply rooted socioeconomic factors, correlate with the seasonal performance of the relative teams?

By representing emotions as measurable time-dependent signals, we address question *(i)* employing approaches stemming both from unsupervised machine learning and complex system analysis, namely clustering methods for time series [28] and burstiness for describing in/homogeneous temporal patterns in signal sequences [29]. Segmenting users’ emotional trends, our clustering analysis uncovers groups of similar fandoms sharing homogeneous features that proxy preseasonal fandoms’ expectations. Burstiness reveals regularities in expressions of joy and heterogeneous patterns of bursty/anti-bursty anger, related to the specific characteristics of each fandom.

Finally, to address question *(ii)*, whether these measures influence the relationship between fandom behavior and the seasonal performance of their respective teams, we employ multiple regression analysis. This approach enables us to assess how emotional patterns, such as burstiness, correlate with final team rankings. By analyzing these dynamics, we demonstrate how emotional signals can help explain variations in team performance.

The remainder of the paper is organized as follows. Section 2 describes data collection and the methods used for clustering, burst detection, and multi-regression. Section 3 focuses on the results achieved using the previously described data and methods. Section 4 concludes the paper by discussing the results.

2 Materials and Methods

2.1 Data Collection and Pre-processing

We collect user comments from the official Instagram accounts of 83 Italian teams that participated in the 2023-24 seasons of Serie A, Serie B, and Lega Pro; the top three divisions in the Italian soccer system. **Among the 83 teams, 18 compete in Serie A, 20 in Serie B, and 45 in Lega Pro, which is further divided into three regional groups: 16 teams in Group A, 14 in Group B, and 15 in Group C.** Teams' accounts share content such as match updates, including squad list, half- and full-time scores, highlights, etc. We refer to this type of content as "posts". The frequency of posting activity varies between teams; cf. A.1.

Users write comments in response to teams' posts and express their opinions, feelings, and reactions to the post content. We refer to users' activities as "comments". Considering that some posts can contain one or a very few number of comments, we detect individual teams' thresholds to discard non-relevant posts. In other words, we retain posts that have at least a number of comments equal to the median of the distribution of the team account, cf. A.1.

We perform an emotion detection task using *feel-it*¹, a python library tailored for the Italian language [30]. Comments are classified into one of four base emotions: joy, anger, sadness, and fear. Consequently, we count the percentage of emotions conveyed by the comments on each post.

Each post is further labeled according to the absence/non maximal (0) or maximal (1) presence of an emotion, for each emotion. For example, a post containing 75% of comments expressing anger, 15% expressing joy, 5% sadness, and 5% fear, is labelled 0 in the dimensions of joy, sadness and fear, and 1 in the dimension of anger, thus being considered as an "emotional event" of this latter dimension. See Appendix A.2 for other information about this representation.

Formally, let \mathcal{E} be the set of emotions $\mathcal{E} = \{joy, anger, sadness, fear\}$.

Definition 1 (Emotional Time Series Dataset) An emotional time series dataset $\mathcal{P} = \{P_1, P_2, \dots, P_n\} \in \mathbb{R}^{n \times |\mathcal{E}| \times m}$ is a collection of n emotional time series.

Definition 2 (Emotional Time Series) An emotional time series P is a collection of $|\mathcal{E}|$ emotional signals $P = \{p_{joy}, p_{anger}, p_{sadness}, p_{fear}\} \in \mathbb{R}^{|\mathcal{E}| \times m}$.

Notably, m refers to the length of the time series, and the quantity is expressed in days, cf. Appendix A.1.

¹<https://github.com/MilaNLProc/feel-it>

Definition 3 (Emotional Signal) An emotional signal, like p_{joy} , is a sequence of m real-valued observations as the percentage of comments classified as joy within a post, i.e., $p_{joy} = [p_1, p_2, \dots, p_m] \in \mathbb{R}^m$; the same holds for the other three emotions.

Moreover, for an event-based emotional representation, we consider a binary labeling of emotional events for each emotion in \mathcal{E} , where an emotional event is defined by the maximal presence of a specific emotion.

Definition 4 (Binary Emotional Signal) Given a post P_i , its emotional signal $p_e \in P_i$ for an emotion $e \in \mathcal{E}$ is transformed into a binary emotional signal $s_e = [s_1, s_2, \dots, s_m] \in \{0, 1\}^m$, assuming value 1 if emotion e is maximal within the post, 0 otherwise.

This transformation results in a dataset of event-based emotional time series.

Definition 5 (Event-based Emotional Time Series) An event-based emotional time series S is a collection of $|\mathcal{E}|$ binary emotional signals $S_i = \{s_{joy}, s_{anger}, s_{sadness}, s_{fear}\}$.

Definition 6 (Event-based Emotional Time Series Dataset) An event-based emotional time series dataset $\mathcal{S} = \{S_1, S_2, \dots, S_n\} \in \mathbb{R}^{n \times |\mathcal{E}| \times m}$ is a collection of n event-based emotional time series.

These time series representations guarantee the privacy of individual users [31]. Textual content from users can not be traced or monitored, as all comments are aggregated into an emotion distribution of the posts, and our analysis is limited to the post level.

2.2 External Features

Additionally, our analysis incorporates key contextual variables for each time series: per capita income (PCI), unemployment rate (U), welfare spending (W)², geographical location (GEO), team heritage R_H ³, and team market value (MV)⁴. These data enrich our analytical framework by providing quantifiable measures of the socioeconomic landscape and competitive environment surrounding each fandom. Specifically, PCI, U, W, and GEO capture fundamental socioeconomic conditions of the team's supporter base at the city level, while MV is a quantitative representation of pre-season performance expectations. This multidimensional approach allows us to explore how emotional dynamics interact with broader socioeconomic and competitive contexts.

Teams' market values (MV) are reasonable proxies for fandoms' expectations because they consider team's investments in key players, financial successes, brand

²PCI, U, and W have been collected using Istat data: <https://www.istat.it/>

³This variable is gathered following the Italian criterion for team reputation: https://it.wikipedia.org/wiki/Tradizione_sportiva

⁴Since market values fluctuate throughout the season, we consider only their values at the beginning of the 2023/24 seasons — specifically, the day before each league's kickoff, retrieving them from <https://www.transfermarkt.it/>

strength, and overall reputation. This choice should guarantee a more reliable representation of fandoms' expectations. Moreover, to evaluate the alignment between pre-season expectations and end-of-season outcomes, we consider the difference Δ_{mv} between a team's expected rank, based on its market value, R_{mv} , and its final rank, R , as $\Delta_{mv} = R_{mv} - R$. If Δ_{mv} is negative, the team underperformed relative to its expected rank R_{mv} ; otherwise, the team outperformed relative to it. For example, since *Spezia*, ranked first in terms of market value in Serie B, finished the season in 15th place, Δ_{mv} is -15, indicating that expectations were not met. Conversely, since *Catanzaro*, ranked 12th in market value, finished 5th, Δ_R is +7, reflecting a performance exceeding expectations.

Similar logic can be applied to further contextualize socioeconomic features. For instance, wealth disparities can shape infrastructures, youth development, and team investments. Thus, to evaluate an alignment between a rank based on such wealth factors and end-of-season outcomes, we consider the difference Δ_{pci} between a team's expected rank, in terms of per capita income, R_{PCI} , and its final rank, R , as $\Delta_{pci} = R_{PCI} - R$.

Formally, each time series P can be further described as $P = (\mathbf{p}, \mathbf{a})$, where \mathbf{p} represents the emotional signals as sequence of observations, and \mathbf{a} is a vector of k attributes associated with the entire time series, i.e., contextual information that does not change over time.

2.3 Time Series Clustering

We aim to partition the time series dataset \mathcal{P} defined in Section 2.1 into clusters of emotionally aligned fandoms. The dataset previously defined is multivariate, since each post contains the four emotional dimensions. Without losing information, for this task we can work on the univariate time series, with $|\mathcal{E}| = 1$, choosing *joy*, being the most relevant dimension.

Definition 7 (Emotional Time Series Dataset Clustering) Clustering an emotional time series dataset \mathcal{P} involves partitioning \mathcal{P} into k clusters, $\mathcal{C} = \{\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_k\}$, such that $\bigcup_{j=1}^k \mathcal{C}_j = \mathcal{P}$, i.e., all time series are assigned to a cluster, and $\mathcal{C}_i \cap \mathcal{C}_j = \emptyset$ for $i \neq j$, i.e., all clusters are disjoint.

There exist several clustering methods for time series [28], including adaptations of classic clustering algorithms, such as k-means and hierarchical clustering. An important difference when dealing with time series is the distance calculation. Euclidean distance is often insufficient because it can not handle temporal distortions. A more relevant similarity measure is Dynamic Time Warping (DTW), which aligns time series by allowing flexible warping of the time axis.

In the experimental section, we use hierarchical clustering with DTW as the distance measure to partition the dataset. Hierarchical clustering involves iteratively merging clusters based on pairwise DTW distances, producing a dendrogram for flexible analysis, cf. A.2. We also consider the k-means algorithm for time series, which partitions data by minimizing the within-cluster DTW distance to a representative

centroid, updating cluster assignments iteratively. The full details and results of k-means clustering are reported in A.2 to avoid overloading the main experimental section.

Note that for this task, we further aggregate values per day, ensuring that all elements are conform to a daily temporal resolution, as teams have varying posting frequencies — e.g., some teams may post 10 times a day while others only once.

2.4 Burstiness and Memory

Let τ be the inter-event time between two consecutive events. We aim to quantify whether the distribution of the inter-event times $P(\tau)$ deviate from the random activity pattern characterized by the exponential distribution. In our cases, $P(\tau)$ is the distribution of the waiting time τ that elapses before a target emotion appears again in the emotional sequence. We use the mean value of the distribution, $\mu_\tau = \frac{1}{n_\tau} \sum_{i=1}^{n_\tau} \tau_i$, where n_τ is the length of the sequence, and the standard deviation, $\sigma_\tau = \sqrt{\frac{\sum_{i=1}^{n_\tau} (\tau_i - \mu_\tau)^2}{n_\tau}}$, to measure the coefficient of variation r as follows:

$$r = \frac{\sigma_\tau}{\mu_\tau} \quad (1)$$

The variation r is a measure of the dispersion of a distribution compared to its mean, useful to introduce the burstiness parameter B as follows:

$$B = \frac{\sigma_\tau - \mu_\tau}{\sigma_\tau + \mu_\tau} = \frac{r - 1}{r + 1} \quad (2)$$

Burstiness B ranges from -1 and 1, where 1 indicates the bursty time series (σ_τ much higher than μ_τ), -1 the periodic one, and 0 indicates the random activity pattern.

Being B affected by the finite number of events in the time series [32], it can be used a variation of B for finite event sequences of size n as follows:

$$B_n = \frac{\sqrt{n+1} \cdot r - \sqrt{n-1}}{(\sqrt{n+1} - 2)r + \sqrt{n-1}} \quad (3)$$

This modification allows to analyze time series with a relatively small number of events, which can be the case of low posting activity of a team account. In the experiments, we will use the quantity B_n described in Eq. 3 to measure burstiness.

Moreover, we use the memory parameter M [29] as follows:

$$M = \frac{1}{n-1} \sum_{i=1}^{n-1} \frac{(\tau_i - \mu_1)(\tau_{i+1} - \mu_2)}{\sigma_1 \sigma_2}, \quad (4)$$

where μ_1 and μ_2 , and σ_1 and σ_2 are sample mean and sample standard deviation of τ_i 's values and τ_{i+1} 's values, with $(i = 1, \dots, n_\tau - 1)$. The memory coefficient M is similar to the autocorrelation of a time series at lag=1. **This** measure can be extended to a generic lag k [33], by replacing 1 in $n_\tau - 1$ with k , and letting calculate the correlation between two inter-event times τ and τ' separated by k events (instead of 1 event) to better capture long-range correlations. Memory M ranges from -1 and 1, where 1 indicates short/long inter-event times followed by short/long ones, -1 indicates short-/long inter-event times followed by long/short ones, and 0 indicates no correlation.

In the experiments, we will use the quantity M described in Eq. 4 to measure the memory parameter, which is consistent with its original definition [29].

2.5 Regression and Model Comparison

To investigate the eventual relationship between emotional dynamics and external features, we employ multiple regression analysis. Multiple regression is a statistical technique that is used to model the relationship between a dependent variable and multiple independent variables. It extends simple linear regression by allowing the inclusion of several **features**, enabling the analysis of how each independent variable contributes to explaining the variation in the dependent variable while controlling for the effects of other variables.

The general form of a multiple regression model is:

$$Y = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \cdots + \beta_p X_p + \epsilon, \quad (5)$$

where Y is the dependent variable, e.g., the final ranking R , and X_p are the independent variables, such as the socioeconomic variables or the burstiness parameter as described in Eq. 3; β_0 is the intercept, β_p are the coefficients representing the effect of each independent variable on Y , and ϵ is the error term, capturing unexplained variability.

To evaluate the performance of the regression model, we use the coefficient of determination, R^2 , and the root mean squared error, $RMSE$.

R^2 measures the proportion of variance in the dependent variable that is explained by the independent variables. It ranges from 0 to 1, where higher values indicate better fit. R^2 is calculated as:

$$R^2 = 1 - \frac{SS_{\text{res}}}{SS_{\text{tot}}}, \quad (6)$$

where SS_{res} is the residual sum of squares and SS_{tot} is the total sum of squares.

$RMSE$ measures the average deviation of the predicted values from the actual values. It provides a sense of the model's prediction error in the units of the dependent variable. $RMSE$ is calculated as:

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (Y_i - \hat{Y}_i)^2}, \quad (7)$$

where Y_i is the actual value, \hat{Y}_i is the predicted value, and n is the number of observations. Lower $RMSE$ values indicate better predictive accuracy.

Moreover, to assess whether the inclusion of additional features leads to a statistically significant improvement in model performance, we can compare different regression models with an F-test. Two regression models are nested when one model (the reduced model) is a special case of the other (the full model) obtained by setting one or more coefficients to zero. For example, if the full model includes variables X_1, X_2, X_3 and the reduced model includes only X_1, X_2 , the models are nested.

The null hypothesis of the F-test states that the additional features in the full model do not provide a significant improvement in explanatory power. The F-statistic is defined as follows:

$$F = \frac{(\text{SS}_{\text{res}}^{(r)} - \text{SS}_{\text{res}}^{(f)}) / (p_f - p_r)}{\text{SS}_{\text{res}}^{(f)} / (n - p_f - 1)}, \quad (8)$$

where $\text{SS}_{\text{res}}^{(r)}$ and $\text{SS}_{\text{res}}^{(f)}$ are the residual sum of squares of the reduced and full models, respectively, p_r and p_f are the number of features in the reduced and full models, and n is the number of observations. A statistically significant F-statistic indicates that the additional variables included in the full model significantly improve the fit.

3 Results

3.1 Common Features within Clusters

While partitioning the dataset, we expect that fandoms within the same cluster exhibit similar trends in their emotional behavior compared to fandoms in other clusters. As mentioned in Section 2.3, we apply the hierarchical clustering algorithm, selecting $k = 3$ number of clusters by looking at the resulting dendrogram, cf. A.3, while the results for the k-means clustering are left for A.3.

Figure 1 and Figure 2 highlight the feature distributions in the three identified clusters. Indeed, to gain insights into such groups, we analyze how features, as described in Section 2.2, distribute within groups, recalling that clustering does not ensure or aims to maximize feature homogeneity, as it operates exclusively on the emotional time series data. First of all, from Figure 1, there emerges a highly homogeneous cluster in terms of geographical distribution — cluster id, $C2$ — which consists of 15 fandoms from North Italy, and only 3 and 1 fandoms from Center and South Italy, respectively; these are *Empoli*, *Perugia*, *Pontedera*, from Center, and *Catanzaro*, from South. Interestingly, only 2 fandoms from Serie A are present in this cluster: *Empoli* and *Monza*. Among the 6 teams from Serie B, five of them — *Parma*, *Como*, *Venezia*, *Cremonese*, and *Catanzaro* — are ranked in the top six of the league's final ranking, plus *Cittadella*. The socioeconomic attributes of this cluster falls, on average, within the lowest tertile, Q1, with the same applying to MV, also in Q1. Moreover, teams within this group, on average, have finished in the top part of the ranking, and generally show higher Δ_{pci} , and, to a lesser extent, higher Δ_{mv} (Figure 2). Remember that positive differences in both variables indicate final results that exceeded pre-season expectations.

Moving on to the other groups, the cluster labeled with id $C1$ consists of 7, 4 and 10 fandoms from North, Center and South Italy, respectively, whereas cluster id $C3$ consists of 20, 13 and 10 fandoms from North, Center, and South Italy, respectively (Figure 1). In other words, the geographical distribution reveals a higher concentration of teams from the South in $C1$, and from the North in $C3$. There is no notable homogeneity in terms of leagues, suggesting that emotional differences are not influenced by league hierarchies. Regarding PCI and MV, the two clusters show mixed distributions without clear patterns, though the first one contains a higher proportion of Q2 bins. As seen in Figure 2, it is evident that the first cluster, $C1$, predominantly features negative moods, with both the values of Δ_{pci} and Δ_{mv} falling below zero. These negative

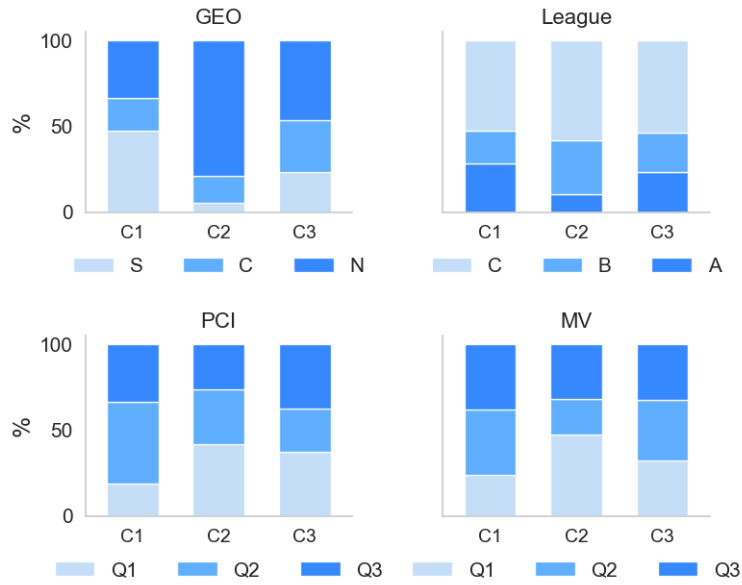


Fig. 1: Distribution of categorical features across the three clusters identified by the hierarchical clustering. The labels *S*, *C*, *N* in the GEO legend represent *South*, *Center*, and *North Italy*, while *C*, *B*, *A* in the League legend stands for Lega Pro, Serie B, and Serie A, and *Q1*, *Q2*, *Q3* correspond to the tertiles of the binned average per capita incomes (PCI) and market values (MV), with *Q1* representing the lowest category and *Q3* the highest.




Fig. 2: Distribution of continuous features across the three clusters identified by the hierarchical clustering.

differences indicate that final results did not meet pre-season expectations. The third cluster, *C3*, displays less homogeneity, except for a slight negative mood in Δ_{pci} .

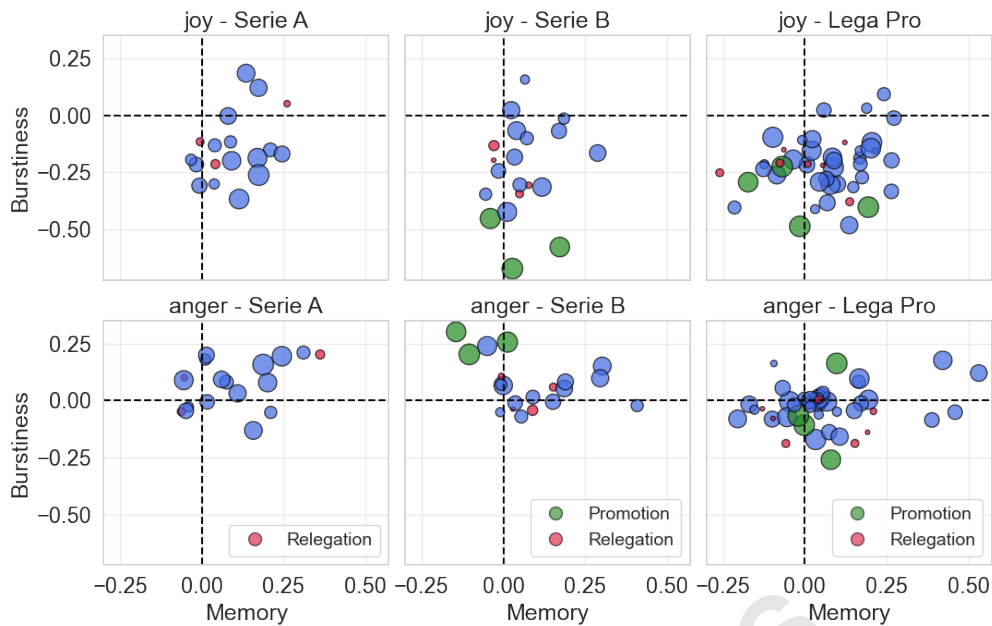


Fig. 3: Focus on *joy* and *anger*: Points represent pairs of burstiness parameter and memory coefficient values. Sizes map the final rank of the teams; moreover, green/red points highlight a promotion/relegation to a higher/lower league.

3.2 Loyalty as Anti-Bursty Patterns and Memory Effects

In this section, we investigate whether fandoms' emotional signals manifest bursty properties as other human dynamics [34], or whether they exhibit random, or non-random but regular, dynamics. Figure 3 shows burstiness vs. memory diagrams, grouped by leagues, for two target emotions: joy and anger. We do not consider sadness and fear, in this main analysis, since they rarely come as the most frequent emotions within posts, see A.2.

Joy and anger unveil different event-based dynamics. Joy rarely displays a bursty trend, and only a few fandoms have a burstiness greater than 0. In Serie A, these teams are *Fiorentina*, *Lazio*, and *Salernitana*, the latter one also relegated to Serie B at the end of the season. In Serie B, the teams are *Bari* and *Sampdoria*. In Lega Pro, *Ancora*, *Catania*, *Foggia*, and *SPAL*. Notably, these fandoms also manifest a positive memory effect. The other fandoms, in all leagues, overall depict anti-bursty emotional dynamics with respect to joy. This can indicate consistent loyalty and supportive behavior for the respective teams. Serie B has more supportive fandoms, also including the four teams relegated to Lega Pro at the end of the season. This suggests a complex bond with the teams that goes beyond the final rankings and match results, possibly linked to pre-season expectations, cf. 3.1.

In contrast to joy, anger displays more heterogeneous behaviors, with memory effects varying more widely across teams. Only a few fandoms highlight a burstiness

value less than 0 among Serie A and Serie B, namely *Lazio*, and *Udinese* in Serie A, and *Bari*, *Cittadella*, *Lecco*, *Pisa*, and *Ternana* in Serie B, with *Lecco* and *Ternana* being relegated to Lega Pro, and *Bari* having disputed the play-out matches against *Ternana*. These behaviors are consistent with the bursty patterns observed for joy for these teams, e.g., *Lazio* has both the highest values of joy-burstiness and the lowest values of anger-burstiness. This indicates a constant feeling of anger, while joy can be relative to single matches. Notably, Lega Pro's burstiness values differ from the other two leagues. An anti-bursty anger behavior is observed among the relegated teams. Moreover, memory with respect to anger uncovers a wide heterogeneous behavior, making it difficult to identify a common pattern across all fandoms. In Serie B, for instance, *Como*, *Cremonese* and *Venezia* show the most negative memory effect, whereas *Cittadella*, *Palermo* and *Sampdoria*, the most positive memory effect.

3.3 Joy burstiness and team performance

In this section, we examine the association between final team ranking R and the set of correlates identified in Section 2.2: team heritage R_H , per capita income PCI , market value MV , joy burstiness B_{joy} , unemployment rate U , welfare spending W , geographical location GEO . Our main objective is to combine socioeconomic variables with complex sentiment-based measures (i.e., emotion burstiness) to find correlations with teams' performance.

To identify the most relevant descriptors of R , we applied backward elimination stepwise regression. Starting with a full model containing all seven features ($R^2 = 0.51$), we iteratively tested the removal of each variable using F-tests to compare nested models. The elimination process proceeded as follows: First, U was removed ($F = 0.118, p = 0.73$); second, W was removed ($F = 0.708, p = 0.40$); third, GEO ($F = 1.327, p = 0.25$) was removed, as their inclusion did not significantly improve the model. Thus, the final model retained four statistically significant features: R_h , PCI , MV , B_{joy} ($R^2 = 0.49$). This final model has the form

$$R \sim \beta_0 + \beta_1 R_H + \beta_2 PCI + \beta_3 MV + \beta_4 B_{joy} + \epsilon, \quad (9)$$

with the following coefficient values: $\beta_1 = 1.21$, $\beta_2 = -1.33$, $\beta_3 = -2.59$, $\beta_4 = 2.47$.

Next, we are particularly interested in the contribution of joy burstiness. This can be assessed using a stepwise forward selection approach, where the full model is represented in Equation (9) and the reduced model, excluding a specific variable such as B_{joy} can take the reduced form

$$R \sim \beta_0 + \beta_1 R_H + \beta_2 PCI + \beta_3 MV + \epsilon \quad (10)$$

The inclusion of B_{joy} leads to a substantial and statistically significant improvement in model fit ($F = 25.60, p = 0.00$), indicating that joy burstiness captures significant variation in final rankings.

Clearly, it is important to note that while regression supports this correlation between joy burstiness and final rankings, the causal direction of the relationship remains ambiguous. The observed association may reflect team performance influencing fan emotional patterns rather than fan emotions driving team outcomes. In other words, successful teams may generate more consistent expressions of joy among

their supporters, resulting in lower burstiness values, while underperforming teams may produce more sporadic joy reactions thus higher burstiness values. Therefore, joy burstiness should be interpreted primarily as a correlate of team performance. Disentangling causal pathways would require longitudinal analysis that lie beyond the scope of these observations.

4 Discussion and Conclusion

We analyzed the complex emotional dynamics of online soccer fandoms in Italy’s Serie A, B, and Lega Pro, processing comments from Instagram to quantify users’ emotions, such as joy and anger, throughout the 2023-24 season. Using clustering for partitioning the resulting emotional time series, we uncovered groups of fandoms featuring similar pre-season expectations and geographical distribution, exclusively by looking at such emotional signals. Then, we used measures from complex systems analysis, namely burstiness and memory [29], unveiling, overall, bursts of anger and steady feelings of joy across the 83 fandoms considered. **Finally, we investigated the role of burstiness in joy — proxying complex fan behavior — when correlating multiple variables to teams’ final ranking, such as pre-season expectations and socioeconomic features.**

Although correlates such as goals scored, expected goals (xG)⁵, and complex statistics that capture offensive and defensive effectiveness [11] are already established predictors of team rankings, our research focused on the other side of the equation: patterns of fan behavior. Rather than attempting to replace previous predictors, we investigated whether social characteristics of fan communities could also correlate with team outcomes, especially given that both the clustering and burstiness analyses revealed non-random and meaningful emotional patterns. The current analysis is confined to a single season, and one year alone might provide a restricted view of the potential long-term burstiness in users’ behavior. However, including team heritage as a dependent variable does not improve the correlations, reaching $R^2 = 0.47$ when removing it, compared to $R^2 = 0.50$ as outlined in Section 3.3. Hence, in the future, we plan to incorporate additional features that may more accurately reflect team traditions and other complex fan characteristics, e.g., more refined metrics that can account for fan pre-season expectations [22]. **Future work will also include a more in-depth analysis examining the individual role of socio-economic features and investigating their correlation with other behaviors, such as stadium attendance, as explored in previous studies [35].** Moreover, the bursty nature of emotional dynamics has received limited attention outside cognitive and psychological studies, with only preliminary investigations previously conducted [23]. For future research, we also plan to expand the investigation of emotional burstiness by integrating psychological frameworks as well as exploring the neural correlates of fan engagement in sports [36]. An even more interdisciplinary approach will provide deeper insights into how collective emotions manifest in sports communities and potentially influence team performance outcomes.

The quality of our results inevitably depends on the emotion classifier used to label comments, which is based on the *feel-it* corpus for Italian language [30]. Since this corpus is derived from social media data, it aligns well with the content of our dataset.

⁵<https://understat.com/>

However, the limitation to only four emotions may narrow the spectrum of possible engagement experienced in sport contexts. Although we mainly observed collective expressions of joy and anger, these could still be interpreted in different ways. For example, joy may arise not only from one's own team winning but also from a rival team losing. In the future, we plan to further explore these emotional shades, although doing so requires a more complex understanding of language.

To conclude, although online fandom is by no means comparable to the physical experience inside and outside stadiums, the digital traces left by fans on social media suggest potential avenues for research in sports analytics, social dynamics, and fan engagement studies.

List of Abbreviations

- *B*: Burstiness parameter
- *DTW*: Dynamic Time Warping
- *F*: F-statistics
- *GEO*: Geographical location
- *M*: Memory parameter
- *MV*: Market value
- *PCI*: Per Capita Income (City)
- *R*: Final rank
- *R_H*: Team heritage
- *R_{mv}*: Expected rank based on market value
- *R_{pci}*: Expected rank based on per capita income
- *RMSE*: Root Mean Squared Error
- *U*: Unemployment rate
- *W*: Welfare spending
- *xG*: Expected goals
- Δ_{mv} : Difference between expected rank (market value) and final rank
- Δ_{pci} : Difference between expected rank (PCI) and final rank

Declarations

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Availability of data and material

Data used in this work are available at the following link: <https://zenodo.org/records/15350266>

Competing interests

The authors declare that they have no competing interests.

Funding

This work is supported by SoBigData.it which receives funding from the European Union – NextGenerationEU – National Recovery and Resilience Plan (Piano Nazionale di Ripresa e Resilienza, PNRR) – Project: “SoBigData.it – Strengthening the Italian RI for Social Mining and Big Data Analytics” – Prot. IR0000013 – Avviso n. 3264 del 28/12/2021.

Authors’ contributions

Conceptualization: SC, GM; Data curation: SC; Formal analysis: SC, GM; Investigation: SC, GM, EF; Methodology: SC, GM; Supervision: EF; Validation: SC; GM, EF; Visualization: SC, GM; Writing, original draft: SC, GM, EF.

Acknowledgments

This work is supported by SoBigData.it which receives funding from the European Union – NextGenerationEU – National Recovery and Resilience Plan (Piano Nazionale di Ripresa e Resilienza, PNRR) – Project: “SoBigData.it – Strengthening the Italian RI for Social Mining and Big Data Analytics” – Prot. IR0000013 – Avviso n. 3264 del 28/12/2021.

We dedicate this work to Giorgio Vignando, a quiet sentinel of the midfield. He taught us that football’s joy does not simply live in the roar of goals, but in the unspoken beauty of a battle fought in shadow.

References

- [1] Girvan, M., Newman, M.E.: Community structure in social and biological networks. *Proceedings of the national academy of sciences* **99**(12), 7821–7826 (2002)
- [2] Lazer, D., Pentland, A., Adamic, L., Aral, S., Barabási, A.-L., Brewer, D., Christakis, N., Contractor, N., Fowler, J., Gutmann, M., *et al.*: Computational social science. *Science* **323**(5915), 721–723 (2009)
- [3] Edelman, A., Wolff, T., Montagne, D., Bail, C.A.: Computational social science and sociology. *Annual review of sociology* **46**(1), 61–81 (2020)
- [4] Tardelli, S., Nizzoli, L., Tesconi, M., Conti, M., Nakov, P., Da San Martino, G., Cresci, S.: Temporal dynamics of coordinated online behavior: Stability, archetypes, and influence. *Proceedings of the National Academy of Sciences* **121**(20), 2307038121 (2024)

- [5] Hutto, C., Gilbert, E.: Vader: A parsimonious rule-based model for sentiment analysis of social media text. In: Proceedings of the International AAAI Conference on Web and Social Media, vol. 8, pp. 216–225 (2014)
- [6] Zhao, J., Liu, K., Xu, L.: Sentiment analysis: Mining opinions, sentiments, and emotions. MIT Press One Rogers Street (2016)
- [7] Blei, D.M., Ng, A.Y., Jordan, M.I.: Latent dirichlet allocation. *Journal of machine Learning research* **3**(Jan), 993–1022 (2003)
- [8] Grootendorst, M.: Bertopic: Neural topic modeling with a class-based tf-idf procedure. arXiv preprint arXiv:2203.05794 (2022)
- [9] Goldenberg, A., Halperin, E., Van Zomeren, M., Gross, J.J.: The process model of group-based emotion: Integrating intergroup emotion and emotion regulation perspectives. *Personality and social psychology review* **20**(2), 118–141 (2016)
- [10] Goldenberg, A., Garcia, D., Halperin, E., Gross, J.J.: Collective emotions. *Current directions in psychological science* **29**(2), 154–160 (2020)
- [11] Cintia, P., Giannotti, F., Pappalardo, L., Pedreschi, D., Malvaldi, M.: The harsh rule of the goals: Data-driven performance indicators for football teams. In: 2015 IEEE International Conference on Data Science and Advanced Analytics (DSAA), pp. 1–10 (2015). IEEE
- [12] Rossi, A., Pappalardo, L., Cintia, P., Iaia, F.M., Fernández, J., Medina, D.: Effective injury forecasting in soccer with gps training data and machine learning. *PloS one* **13**(7), 0201264 (2018)
- [13] Pappalardo, L., Cintia, P., Ferragina, P., Massucco, E., Pedreschi, D., Giannotti, F.: Playerank: data-driven performance evaluation and player ranking in soccer via a machine learning approach. *ACM Transactions on Intelligent Systems and Technology (TIST)* **10**(5), 1–27 (2019)
- [14] Cintia, P., Mauro, G., Pappalardo, L., Ferragina, P.: An interactive dashboard for searching and comparing soccer performance scores. arXiv preprint arXiv:2105.04293 (2021)
- [15] Schumaker, R.P., Jarmoszko, A.T., Labeledz Jr, C.S.: Predicting wins and spread in the premier league using a sentiment analysis of twitter. *Decision Support Systems* **88**, 76–84 (2016)
- [16] Pacheco, D.F., Pinheiro, D., Lima Neto, F.B., Ribeiro, E., Menezes, R.: Characterization of football supporters from twitter conversations. In: *WI*, pp. 169–176 (2016)
- [17] Lucas, G.M., Gratch, J., Malandrakis, N., Szablowski, E., Fessler, E., Nichols, J.:

- Goaalll!: Using sentiment in the world cup to explore theories of emotion. *Image and Vision Computing* **65**, 58–65 (2017)
- [18] Thonhauser, G., Wetzels, M.: Emotional sharing in football audiences, 110–129 (2020)
- [19] Guzmán, E.M., Zhang, Z., Ahmed, W.: Towards understanding a football club’s social media network: an exploratory case study of manchester united. *Information Discovery and Delivery* **49**(1), 71–83 (2021)
- [20] Buongiovanni, C., Candusso, R., Cerretini, G., Febbe, D., Morini, V., Rossetti, G.: Will you take the knee? italian twitter echo chambers’ genesis during euro 2020. In: *International Conference on Complex Networks and Their Applications*, pp. 29–40 (2022). Springer
- [21] Attié, M., Pacheco, D., Oliveira, M.: Getting the boot? predicting the dismissal of managers in football. In: *International Workshop on Complex Networks*, pp. 132–140 (2023). Springer
- [22] Wang, Y., Lu, Z.: Making sense of post-match fan behaviors in the online football communities. In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pp. 1–17 (2023)
- [23] Citraro, S., Mauro, G., Ferragina, E.: Burstiness in emotions: A case study on collective affective responses in italian soccer fandoms. In: *International Conference on Advances in Social Networks Analysis and Mining*, pp. 56–69 (2024). Springer
- [24] Woods, J., Ludvigsen, J.A.L.: The changing faces of fandom? exploring emerging ‘online’and ‘offline’fandom spaces in the english premier league. *Sport in Society* **25**(11), 2234–2249 (2022)
- [25] Lawrence, S., Crawford, G.: The hyperdigitalization of football cultures, 1–16 (2018)
- [26] Mainwaring, E., Clark, T.: ‘we’re shit and we know we are’: identity, place and ontological security in lower league football in england. *Soccer & Society* **13**(1), 107–123 (2012)
- [27] Biscaia, R., Correia, A., Rosado, A., Maroco, J., Ross, S.: The effects of emotions on football spectators’ satisfaction and behavioural intentions. *European Sport Management Quarterly* **12**(3), 227–242 (2012)
- [28] Shumway, R.H., Stoffer, D.S., Stoffer, D.S.: Time series analysis and its applications. Springer **3** (2000)
- [29] Goh, K.-I., Barabási, A.-L.: Burstiness and memory in complex systems. *Europhysics Letters* **81**(4), 48002 (2008)

- [30] Bianchi, F., Nozza, D., Hovy, D., *et al.*: Feel-it: Emotion and sentiment classification for the italian language. In: Proceedings of the Eleventh Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis (2021). Association for Computational Linguistics
- [31] Monreale, A., Andrienko, G.L., Andrienko, N.V., Giannotti, F., Pedreschi, D., Rinzivillo, S., Wrobel, S., *et al.*: Movement data anonymity through generalization. *Trans. Data Priv.* **3**(2), 91–121 (2010)
- [32] Kim, E.-K., Jo, H.-H.: Measuring burstiness for finite event sequences. *Physical Review E* **94**(3), 032311 (2016)
- [33] Schleiss, M., Smith, J.A.: Two simple metrics for quantifying rainfall intermittency: The burstiness and memory of interamount times. *Journal of Hydrometeorology* **17**(1), 421–436 (2016)
- [34] Karsai, M., Jo, H.-H., Kaski, K., *et al.*: *Bursty human dynamics*. Springer (2018)
- [35] Addesa, F., Bond, A.J.: Determinants of stadium attendance in italian serie a: New evidence based on fan expectations. *PLoS one* **16**(12), 0261419 (2021)
- [36] Duarte, I.C., Afonso, S., Jorge, H., Cayolla, R., Ferreira, C., Castelo-Branco, M.: Tribal love: the neural correlates of passionate engagement in football fans. *Social cognitive and affective neuroscience* **12**(5), 718–728 (2017)

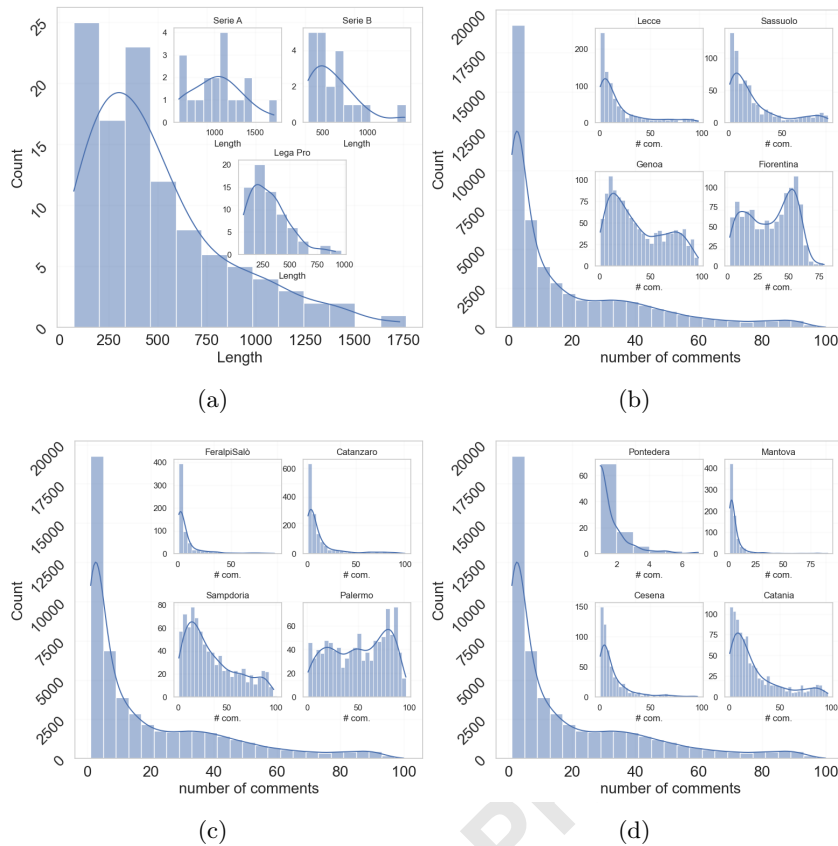


Fig. A1: (a) Distribution of time series lengths, where the length is given by the number of teams' posts over the 2023/24 seasons. The inner plots show the distributions of each league; (b-d) Distribution of user comments within posts, where the inner plots show the individual distributions of randomly selected teams across all leagues to show the strong heterogeneity between them.

Appendix A Notes on Materials and Methods

A.1 Posts and Comments

The frequency of posting activity varies significantly between teams. For instance, one team might post multiple times a day, while another might post only once or twice a week. This disparity is illustrated in Fig. A1 (a). The inner plots also indicate that, on average, teams from Serie A tend to post more frequently than teams from other leagues. In any case, in our analyses, we keep these teams' lengths when calculating the burstiness and memory coefficients. Instead, another pre-process step is needed to reduce noise when performing clustering, that is the temporal alignment of time series. This involves aggregating posts on a **daily** basis from the start to the end of the




Fig. A2: From real-value based time series (a) to event-based time series according to different emotions, i.e., joy (b) and anger (c).

considered seasons. Hence, when computing the percentage of emotions, we further average these percentages on the daily basis.

Another relevant distribution to consider is the number of user comments on posts. This distribution is also long-tailed, as shown in Fig. A1 (b-d). Since a post can have only one or just a few comments, it is necessary to reduce this noise. In other words, detecting 100% joy in a post should not be meaningful if the post contains only a single comment. However, as indicated by the inner plots, the overall distribution does not accurately represent the individual distributions. For example, the fan bases of Catania and Palermo engage more than those of FeralpiSalò and Pontedera, because their user pools differ a lot. Therefore, to account for this heterogeneity, we apply individual thresholds to select the posts that are retained: we consider only those posts that have a number of comments equal to or greater than the median of the distribution for each individual account.

A.2 Event-based Time Series

Fig. A2 highlights the event-based time series representation. We consider three teams from Serie B as a toy example, but the following holds true on average for all the fandoms in the dataset. Notably, the two dominant emotions expressed throughout the season are joy and anger, which is why we focus exclusively on these two emotions. The resulting event-based time series emphasize the occurrences of joy and anger events. Parma, for instance, consistently at the top of the rankings, exhibits a joy-based time series dominated by events, and the anger-based one with occasional bursts events. Palermo shows very important inhomogeneous patterns for both emotions, while Bari is consistently inclined towards anger-based events.

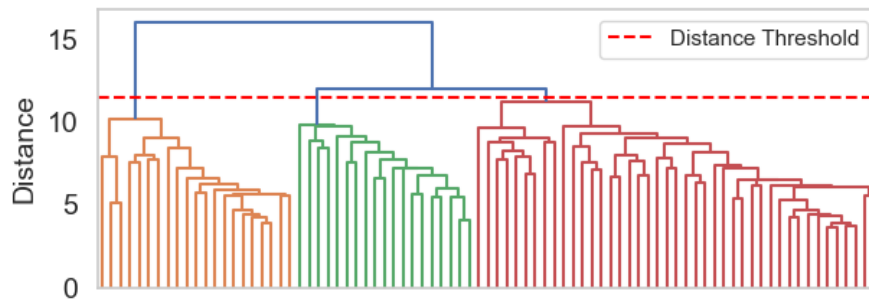


Fig. A3: The dendrogram obtained by performing hierarchical clustering on the real-value based time series dataset. The red line indicates the distance at which we cut to obtain $k = 3$ clusters.




Fig. A4: Distribution of categorical features across the three clusters identified by the k-means clustering. The labels S , C , N in the GEO legend represent *South*, *Center*, and *North Italy*, while C , B , A in the League legend stands for Lega Pro, Serie B, and Serie A, and $Q1$, $Q2$, $Q3$ correspond to the tertiles of the binned average per capita incomes (PCI) and market values (MV), with $Q1$ representing the lowest category and $Q3$ the highest.




Fig. A5: Distribution of continuous features across the three clusters identified by the k-means clustering.

A.3 Notes on Clustering

Figure A3 shows the dendrogram resulting from hierarchical clustering. The plot clearly indicates a division into $k = 3$ clusters, with the optimal cut made at a distance of $d = 10.5$.

This section also presents the results of k-means clustering, where we again chose $k = 3$ as the optimal number of clusters. As shown in Figures A4 and A5, there are notable differences compared to the hierarchical clustering results described in the main article. However, the differences do not indicate contradictory analyses, but rather highlight the ability of the two data partitioning methods to identify different patterns. With kmeans, indeed, a highly homogeneous cluster emerges in terms of the league — cluster id, $C2$. It contains only one team fandom from Serie A, *Empoli*, and one from Serie B, *FeralpiSalò*, out of 19 fandoms. Interestingly, most of the fans within this cluster fall into the lowest categories for both PCI and MV. For instance, the three teams that do not fall into Q1 tertile according to MV are *Empoli*, *FeralpiSalò*, and *Crotone*. The $C2$ cluster predominantly contains teams ranked in the lower part of the standings.

The two clustering methods provide complementary insights: while hierarchical clustering emphasizes geographic patterns, k-means tends to capture league-based distinctions more effectively. Both methods succeed in identifying meaningful characteristics — notably, their ability to detect patterns in external features by analyzing similarities in emotional time series alone.