

Enhancing Arabic Poetic Structure Analysis through Visualization

Abdelmalek Berkani*

Adrian Holzer†

Information Management Institute
University of Neuchâtel, Switzerland




Figure 1: Example of Arabic Poetry visualization.

ABSTRACT

This study explores the potential of visual representation in understanding the structural elements of Arabic poetry, a subject of significant educational and research interest. Our objective is to make Arabic poetic works more accessible to readers of both Arabic and non-Arabic linguistic backgrounds by employing visualization, exploration, and analytical techniques. We transformed poetry texts into syllables, identified their metrical structures, segmented verses into patterns, and then converted these patterns into visual representations. Following this, we computed and visualized the dissimilarities between these images, and overlaid their differences. Our findings suggest that the positional patterns across a poem play a pivotal role in effective poetry clustering, as demonstrated by our newly computed metrics. The results of our clustering experiments showed a marked improvement over previous attempts, thereby providing new insights into the composition and structure of Arabic poetry. This study underscored the value of visual representation in enhancing our understanding of Arabic poetry.

Index Terms: Arabic Poetry, Poetry visualization, Poetry Clustering

1 INTRODUCTION

Poetic compositions are fundamentally rooted in two aspects: content and structure. While translation can bridge language gaps and facilitate content comprehension, it falls short in fully conveying the original semantics, message, affect, and imagery as intended by the poet [28]. However, the structure of literary works is universally comprehensible, making visualization an invaluable tool for exploration and analysis. Information visualization provides a powerful means to explore and understand literary works. By translating textual data into a visual format, readers can get insights [23] that may not be immediately discernible from a cursory reading of the text. Visualization aids in identifying structure, patterns, and trends [18]. Moreover, it enables the comparison and contrast of different literary works, highlighting their similarities and differences.

*e-mail: abdelmalek.berkani@unine.ch

†e-mail: adrian.holzer@unine.ch

In the context of Arabic poetry, structure visualization offers two benefits. Firstly, it provides a visual medium that enables both Arabic and non-Arabic readers to understand the components of poetry and comprehend the poet's arrangement. Secondly, it facilitates the analysis and comparison of various compositions, aiding in the extraction of potential insights, whether the analysis is conducted on text or graphics. This necessitates an examination of the relationship between the presentation of poems and the resulting insights, and the impact of appearance on insight outcomes. This research is a component of a multi-phase project aimed at making Arabic poetry accessible through visualization, exploration, and analysis tools. The first phase focused on studying poetic structures. An initial article addressed meter detection [5], and the second one centered on analyzing poetic compositions and their structures using clustering techniques [4]. This research supplements the latter effort to enhance clustering results. Until now, only textual visualizations of the data have been employed. For the current article, we have chosen graphical visualization to extract insights that were previously unnoticed.

This article seeks to contribute to this objective by addressing the following research question:

RQ: How does visualization aid in identifying new features and enhancing the analysis of classical Arabic poetic structure?

To address this question, this article is structured as follows: Section 2 gives an overview of metered Arabic poetry, Section 3 presents previous work related to this research, Section 4 details the research method and approach, Section 5 presents patterns and dissimilarities visualizations, Section 6 details the clustering experiment, Section 7 summarizes results and insights and finally, Section 8 wraps up with a conclusion and an outlook on future work.

2 ARABIC POETRY STRUCTURE

The composition of a poem in metered Arabic poetry is characterized by its structure, which is organized into verses. Each verse can be delineated as a sequence of syllables, forming specific patterns associated with one or more meters. In metered Arabic poetry, the verse is divided in two parts or hemstitches. In some rare exceptions, the verse does not have a second part.

Syllables. In Arabic poetry, there are three types of syllables, based on the duration of their vocalization. A syllable is represented by the letter (*L*) for long, (*S*) for short and (*A*) in case of absence of vocalization. Syllabic quantity (*Q*) is the base metric to calculate

the weight of any sequence of syllables. A long syllable (L) is twice as long as a short syllable (S) [2]. The absence of vocalization (A) is assumed to have no weight. Thus, the syllable quantity is computed as $Q = 2n_L + 1n_S + 0n_A$. For example the syllable quantity of $LSLA$ is 5.

Patterns. A pattern denotes a designated sequence of syllables, such as the pattern $f\bar{a}'il\bar{a}n$, which consists of the four syllables $LSLA$. Arabic poetry encompasses a total of 43 unique patterns [5]. The syllable count within these patterns varies from 1 to 6, while the syllabic quantity spans from 2 to 9.

Meters. The meter is a set of ordered patterns in a verse. In most poems, the verses adhere to one consistent meter. The occurrence of multiple meters within a single poem is uncommon. The meter is related to an entire verse, and the number of pattern positions in the verse depends on the meter, ranging from 2 to 8 [5]. Since a pattern may exist in more than one meter at one or more positions, any change in verse patterns arrangement may have an impact on the verse's structure. Meters are assigned unique names and serve to establish the rhythm of the verse. There are 16 root meters and 29 different variants in metered Arabic poetry. A variant that retains the same number of patterns as the root meter is considered as *Complete* (C) while one that does not is considered as *Partial* (P).

3 RELATED WORK

3.1 Poetry visualization

Various methods for visualizing poetry structures have been explored in previous work, serving educational, analytical, and artistic purposes [11]. One such tool, Emily, represents poems through colored lines and bars, focusing on the 19th-century verse-letters and poetry of Emily Dickinson [17].

A suite of visualization tools proposed by Meneses et al. [20] assists scholars in conducting critical poetry analysis, synthesizing and highlighting key poem elements. These tools employ graphic shapes like squares, circles, and lines in different colors to create wave-like patterns.

The Poem Viewer [1] is a user-centered design interface for poetry visualization, allowing for the display of poetry text and its associated semantic and phonetic elements. Visualization is achieved using variously colored circles, rectangles, and lines.

Poemage [19] is an interactive tool for exploring a poem's sonic topology. It investigates the role and influence of technology on the close reading of a poem, representing elements like rhyme types, alliterations, and anagrams in various shapes and colors.

SPARSAR [12] is an application for poetry analysis and visualization, using variable shapes and colors to visualize 75 of Shakespeare's sonnets. It provides three views: Phonetic Relational, Poetic Relational, and Semantic Relational, each based on different features of the poem.

The paper on multi-way poetry visualizations [21] proposes a scheme to visualize poetry at four levels, with a case study based on Portuguese poetry.

Poetry Barcode [22] is an application that visualizes a poem's structure, including verb usage, emotional tone, alliterations, adjectives, and word length. The case study focused on the collection of Turkish poet Nâzım Hikmet Ran.

Soliloquy [24] is an interface that visualizes expert thoughts during poem reading and interpretation, aiming to provide insight into its design and potential learning impact.

3.2 Poetry clustering

Some researchers have investigated unsupervised machine learning classification techniques when the categories of poems are unknown. For instance, a study utilized a series of unsupervised classifications based on abstracted semantic features of poems to gain

insights into the relationships between poetry structure and text meaning [26]. This research provided formal evidence of the association between poetic meter and semantics in 18th and 19th century European literature, using collections from Czech, German, and Russian poetry, with additional data from English poetry and early modern Dutch songs.

A recent paper suggests a clustering approach for mining emotional patterns in Romanian poetry using lexicon-based emotion features [16]. That study employs a hierarchical clustering algorithm that yields 50 clusters, which are validated by a Silhouette Index of 0.79.

Another study [7] used a clustering approach for mining T.S. Eliot's poem "*The Waste Land*" using a slightly modified version of the K-means algorithm and achieved better results than baseline approaches.

Another study analyzed a collection of American poems and computed weighted metrics based on orthographic, syntactic, and phonetic features [14]. The findings show that these features are better suited for comparing poetry styles than traditional word-occurrence features.

In our first poetry clustering attempt [4] we explored the unique characteristics of metered Arabic poetry to allow for comparative poetry structures. The study conducted a clustering method based on six poem metrics (Diversity, Evenness, Variability, Repetition, Pattern Potential, and Pattern Usage) using two different algorithms. The findings reveal that Repetition, Variability, and Pattern Usage are the most significant metrics, and only two attributes are needed to characterize the poetry, instead of six.

This literature review shows that researchers have used various techniques to analyze poems using mainly textual features. Furthermore, the literature underscores a prevalent tendency where research in this field predominantly revolves around English and other languages, leaving Arabic—a widely spoken global language—largely unexplored, particularly in the domain of poetry. This study aims to bridge this gap and contribute to the enhancement of natural language processing techniques through the exploration of Arabic poetry. To the best of our knowledge, this paper represents a rare attempt to visually and analytically explore the intricate composition of metered Arabic poetry.

4 MATERIALS AND METHODS

The approach adopted in this article generally focused on five main areas: data preparation, data processing, visualization, analysis, and evaluation as shown in Figure 2.

A selection of 104 poems was drawn from the poetry corpus [3], which was previously utilized by Berkani et al. [5]. The corpus comprises works from 65 different poets and encompasses 26 distinct meter variants. The poems belong to various historical periods extending from pre-Islamic times to the present.

For each poem, the raw poetry text is converted into syllable sequences using the AMIS system [5], which detects the meter and outputs the patterns that compose each verse. Table 1 illustrates an example of a poem's structure with 4 patterns. This structure is particularly utilized after converting the patterns into their respective syllabic quantities, ensuring numerical values that are optimal for analysis and experimentation.

For visualization, analysis, and evaluation purposes, and in order to have comparable poem dimensions, we supplemented the poems' data with dummy syllables according to the maximum syllabic quantity per verse of 48, as found in the meter detection dataset [3]. As a result, all poems are structured in 8 columns, which is the maximum number of possible patterns.

Two visualization options were employed. The first option involved visualizing the patterns with their distinct colors. The second option utilized shared colors based on the syllabic quantity with varying opacity. Differences between images are computed in or-

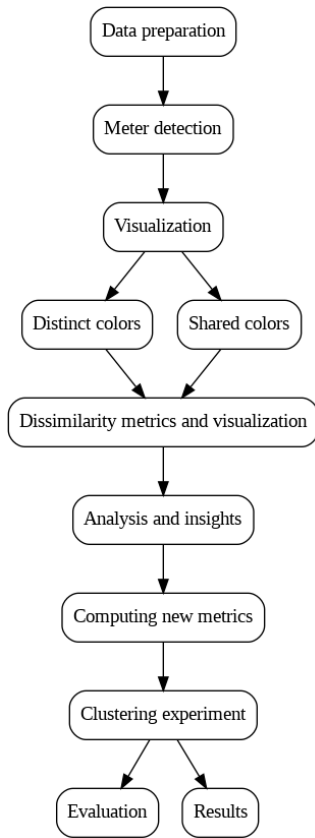


Figure 2: Research approach overview

der to overlay dissimilarities on one of the compared images. This allowed for the analysis of dissimilarities and the observation of potential insights that led to the computation of new metrics and conducting a clustering experiment for evaluation and for effectiveness of the clustering results.

All data used in this article are publicly available [3]. All Python code used for image processing will also be published at the same location.

Table 1: Example of a poem’s metrical structure represented as patterns

P1	P2	P3	P4	P5	P6	P7	P8
LSLL	LSL			LSLL	SSL		
LSLL	LSL			LSLL	SSL		
SSLL	LSL			LSLL	SSL		
LSLL	SSL			LSLL	LSL		

5 VISUALIZATION

Given the difference in the number of verses per poem and to ensure uniform height for all poems converted to PNG format (Portable Network Graphics), only the first 10 verses of each poem were used for visualization.

Patterns are represented by colored rectangles whose lengths correspond to their respective syllabic quantities. The difference between the verse’s syllabic quantity and the maximum width is shown as white rectangles to the right of each verse, as illustrated

in Figures 1, 3, 4, 5, 6.

Although the Arabic language is written from right to left, the patterns are represented from left to right to make the reading of structures more accessible.

5.1 Patterns with distinct colors

In this visualization, each pattern is assigned its own color as shown in the example of Figure 3. The sole criterion guiding the color selection in this visualization is the imperative to visually differentiate one pattern from another [27].



Figure 3: Distinct colors option: Umru al-Qays, Al-Mu’allaqah, At-Tawil C. (poet, poem, meter).

5.2 Patterns with shared colors

In this visualization, as shown in Figure 4, color selection is determined by the syllabic quantity. A unique root color is assigned to each syllabic quantity, and patterns sharing the same syllabic quantity are given colors with reduced opacity, distributed linearly among the patterns. The “viridis” color palette [13] was used for its ability to provide perceptual uniformity, ensuring that changes in data are represented uniformly by changes in color intensity. Additionally, this palette enhances accessibility for individuals with color vision deficiencies.



Figure 4: Shared colors option: Umru al-Qays, Al-Mu’allaqah, At-Tawil C. (poet, poem, meter).

5.3 Dissimilarities

To highlight differences and to quantify the disparities between pairs of images in both structure and color, a balanced dissimilarity approach is used to capture global and perceptual differences. Visualization masks are computed to indicate the regions of significant change, assisting in the interpretation and analysis of the differences between a pair of images. Both histogram and SSIM (Structural Similarity Index Measure) metrics are employed in the mask calculation process.

Histogram. Histogram Similarity measures the distribution of pixel intensities in an image [6]. Comparing the histograms of two images can reveal differences in color distribution, which may indicate changes in the scene or variations in image capture conditions. This metric is sensitive to global changes in color and intensity distribution, making it suitable for detecting significant variations in the overall appearance between two images.

The histogram mask highlights specific regions where the color or intensity differs significantly between the images. A pixel-wise mask is generated by computing the absolute difference between the two images and applying a threshold, marking pixels with differences greater than the threshold.

SSIM. This measure is a perceptual metric that considers changes in structural information, luminance, and contrast [29]. SSIM is calculated between the grayscale versions of the images, and the SSIM score is inverted to express dissimilarity. SSIM focuses on structural changes that are more relevant to human visual perception, making it a robust measure for detecting perceptible differences.

The SSIM mask highlights regions where structural changes are significant. It is generated from the SSIM difference map by marking pixels where the SSIM index is below a certain threshold. By focusing on structural dissimilarities, this mask complements the histogram mask, offering a different perspective on image differences that are more aligned with human visual perception.

Highlighting differences. A white and a contrasted hatching pattern with specific spacing and thickness is used to overlay the areas of difference between the images. Each mask, which identifies the regions of change, is converted to a 3-channel BGR (Blue, Green, Red) image if necessary. A temporary image is created by applying the hatched pattern to the designated areas, and this is then blended with a copy of the original image, preserving its details while highlighting the differences. This approach can accommodate multiple masks, highlighting diverse types of differences within a single image.




Figure 5: Example of dissimilarities between two poetry compositions where each pattern has its distinct color.

6 EXPERIMENT

When visualizing the results of image dissimilarities, we observed that differences are highlighted not only when patterns change but also when their positions change completely or partially. In Arabic poetry, the presence of dissimilarity is a constant, regardless of whether the poems being compared follow the same meter. However, in a few exceptional instances, the dissimilarity disappears



Figure 6: Example of dissimilarities between two poetry compositions with sharing patterns colors.

when the compared poems employ identical patterns at the same positions.

We computed five new metrics related to the distribution of patterns across the poem to determine if clustering results using these new metrics outperform results obtained in a prior research [4]. The experiment was conducted on the same data using the same number of clusters as in the previous study.

In the context of metered Arabic poetry, the columns correspond to the 8 positions of the patterns. The rows correspond to the verses.

Distribution across columns. This metric represents the average quantity per column. It is computed by dividing the total syllabic quantity by the number of columns. This gives an indication of how much quantity, on average, is distributed across each column.

$$D_{\text{cols}} = \frac{Q_{\text{total}}}{N_{\text{cols}}} \quad (1)$$

Q_{total} : Total syllabic quantity of the poem.
 N_{cols} : Number of columns.

Distribution across rows. This metric represents the average quantity per row. It is computed by dividing the total syllabic quantity by the number of rows. This gives an indication of how much quantity, on average, is distributed across each row.

$$D_{\text{rows}} = \frac{Q_{\text{total}}}{N_{\text{rows}}} \quad (2)$$

Q_{total} : Total syllabic quantity of the poem.
 N_{rows} : Number of rows.

Distribution across cells. This metric represents the average quantity per cell. It is computed by dividing the total syllabic quantity by the total number of cells. This gives an indication of how

much quantity, on average, is distributed across each cell in the poem's data.

$$D_{\text{cells}} = \frac{Q_{\text{total}}}{N_{\text{rows}} \times N_{\text{cols}}} \quad (3)$$

Q_{total} : Total syllabic quantity of the poem.
 N_{rows} : Number of rows.
 N_{cols} : Number of columns.

Positions across columns. This metric represents the mean of the column position ratios that are computed as the sum of each column divided by the difference between the number of rows and the sum of that column. The mean of these ratios provides an average measure of how the quantities are positioned across the columns relative to the number of rows.

$$P_{\text{cols}} = \frac{1}{N_{\text{cols}}} \sum_{j=1}^{N_{\text{cols}}} \frac{Q_{\text{col},j}}{Q_{\text{col},j} - N_{\text{rows}} + \varepsilon} \quad (4)$$

N_{cols} : Number of columns.
 $Q_{\text{col},j}$: Sum of syllabic quantities in the j -th column.
 ε : Small positive number to prevent division by zero.

Positions across rows. This metric represents the mean of the row position ratios that are computed as the sum of each row divided by the difference between the number of columns and the sum of that row. The mean of these ratios provides an average measure of how the quantities are positioned across the rows relative to the number of columns.

$$P_{\text{rows}} = \frac{1}{N_{\text{rows}}} \sum_{i=1}^{N_{\text{rows}}} \frac{Q_{\text{row},i}}{Q_{\text{row},i} - N_{\text{cols}} + \varepsilon} \quad (5)$$

N_{rows} : Number of rows.
 $Q_{\text{row},i}$: Sum of syllabic quantities in the i -th row.
 ε : Small positive number to prevent division by zero.

6.1 Clustering

Selecting the optimal set of features for clustering is crucial for achieving meaningful clusters. The quality of clustering results heavily depends on the chosen features, as irrelevant or redundant features can obscure the underlying structure in the data, leading to poor clustering performance. In this approach, we employ a systematic method to identify, among the 5 metrics (D_{cols} , D_{rows} , D_{cells} , P_{cols} , P_{rows}), the best metrics combination that give the best performance.

We start by iterating through all possible combinations of features. Each subset of features is standardized to ensure that all features contribute equally to the clustering process, mitigating the effects of varying scales among features.

For each combination of features, K-means clustering algorithm is performed. The number of clusters is set to 3 in order to compare with result found by the article [4]. The clustering process is repeated 10 times to ensure stability and reliability of the results.

6.2 Evaluation

The clustering performance of each feature combination is evaluated using three metrics. Silhouette coefficient [25] to measure how similar an object is to its own cluster compared to other clusters. A higher Silhouette score indicates better-defined clusters. Ranging between -1 and 1, a value approaching 1 signifies that a data point is significantly distant from neighboring clusters. A value of 0 indicates that the data point is in close proximity to neighboring clusters. Conversely, a negative value suggests that the data point may have been wrongly assigned to an incorrect cluster. According

to [9] and [15], when the Silhouette coefficient exceeds the threshold of 0.50, it indicates a good clustering quality. A value surpassing 0.70 provides strong evidence of excellent clustering quality. Davies-Bouldin index [10] evaluates the average similarity ratio of each cluster with respect to all other clusters. It theoretically ranges from 0 to infinity and should be minimized. Lower values signify better clustering performance. Calinski-Harabasz index [8] also known as the Variance Ratio Criterion considers the ratio of the sum of between-cluster dispersion to within-cluster dispersion. This index ranges from 0 to infinity and should be maximized. Higher values indicate better-defined clusters.

The best features combinations is the one that yields the highest Silhouette score, the lowest Davies-Bouldin score and the highest Calinski-Harabasz score.

The evaluation metrics process shows that the features combinations that contains the single metric "Positions across columns" is the best one for clustering poems data. Clustering scores obtained through this approach surpass those from previous attempts at clustering Arabic poetry compositions [4] as shown in Table 2.

Table 2: Comparison of clustering scores using 'Positions across columns' metric with past values.

Score	Found value	Past value
Silhouette	0.88	0.58
Davies-Bouldin	0.19	0.70
Calinski-Harabasz	1225.20	116.57

The increase of Silhouette coefficient from 0.58 to 0.88 suggests a significant improvement in cluster cohesion and separation. This indicates that the new clustering approach has created more distinct and well-defined clusters compared to the previous results. The decrease of Davies-Bouldin Index from 0.70 to 0.19 implies a substantial improvement in clustering quality. It means that the clusters are more compact and well-separated. The increase of Calinski-Harabasz Index from 116.57 to 1225.20 suggests a significant enhancement in the clustering structure. This indicates that the new clustering results have a much better separation between clusters and less variance within clusters.

7 RESULTS AND DISCUSSION

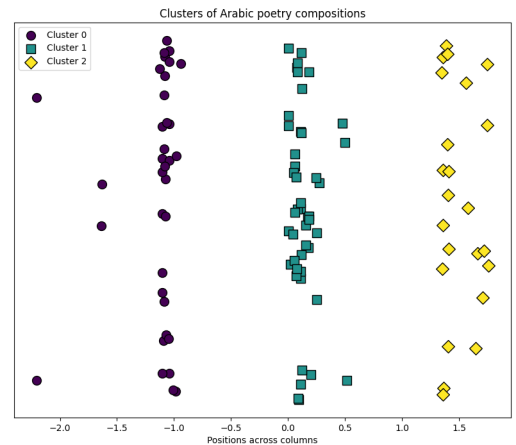


Figure 7: Clustering of Arabic poems by positions across columns

The 104 poems are divided into three clusters containing 35, 46, and 23 members, respectively as shown in Figure 7.

The cluster with 35 members includes all poems that employ meters with patterns arranged in 2, 3, or 4 positions. Additionally, it encompasses poems with meters of 6 positions where patterns have lower syllabic quantities. On the left in Figure 7, we observe four poems that are quite distant from the other poems in their cluster. These poems use rare meters in which all patterns occur in the first part of the verse, consequently reducing their respective syllabic quantities. Hence, this cluster represents poems with a low syllabic composition, averaging around 25 syllables per verse. This falls within the lower range of Arabic poetry's syllabic quantity, which varies between 10 and 48 syllables per verse [5].

The cluster with 46 members encompasses poems using meters with 6 positions. It also includes poems with meters of 8 positions having patterns with lower syllabic quantities. This cluster represents poems with a medium syllabic composition, averaging approximately 36 syllables per verse, placing it in the mid-range of the syllabic quantity spectrum.

The cluster with 23 members includes poems that use meters with 8 positions. This cluster signifies poems with higher syllabic compositions, averaging around 44 syllables per verse, which is close to the upper limit of the syllabic quantity range in Arabic poetry.

In the rare instances where a poem utilizes more than one meter, it is classified based on the dominant meter within the poem, determined by the number of verses adhering to that meter.

Standardizing all poems to 8 positions and filling with dummy syllables to reach a maximum syllabic quantity of 48 per verse has enabled the comparison of all poems, even when their meters differ in the number of pattern positions. Thus, the empty positions in poems that do not have all 8 positions filled provide important information, just as the filled positions do. This standardization allows for a consistent framework in which the structural characteristics of each poem can be analyzed and compared, highlighting both the presence and absence of poetry elements across different compositions.

8 CONCLUSION

This research has made three significant contributions. First, it demonstrated the capability of visualization to alert analysts of Arabic poetry to the existence of insights that warrant exploration through new metrics. By leveraging visual representations, we can uncover previously unseen insights and foster a more interactive and dynamic approach to studying and appreciating Arabic poetry. Second, the study found that the metric related to pattern positions across columns is a significant feature for clustering poetry compositions and highlighting the degree of similarity or dissimilarity among poetic structures. Finally, the creation and availability of a corpus of Arabic poetry in the form of images encourage the visual exploration of poetic compositions. This study addressed the research question and contributed to facilitating an understanding of the intricate patterns inherent in Arabic poetry.

In future research, we plan to investigate the role of structural Arabic poetry visualization in education through user studies. These studies will assess the usefulness of visualizations in enhancing understanding of poetic structures and gaining insights into various aspects of poem compositions. Potential areas of exploration include tracking the evolution of stylistic trends in poems sharing the same meter over time, exploring authorship recognition based on structural compositions, and investigating various structural metrics. Another area of investigation involves examining the connection between poetic structure and semantics by exploring the relationship between the choice of meter and the message conveyed by the poet.

REFERENCES

- [1] A. Abdul-Rahman, J. Lein, K. Coles, E. Maguire, M. Meyer, M. Wynne, C. R. Johnson, A. Trefethen, and M. Chen. Rule-based visual mappings - with a case study on poetry visualization. In *Computer Graphics Forum*. The Eurographics Association, 2013. doi: 10.1111/cgf.12125 2
- [2] I. Anis. *Al Aswat Al Loghawiya (Linguistic sounds)*. Maktabat Al Anglo Masriya, Cairo, 1975. 2
- [3] A. Berkani. Arabic poetry analysis datasets. <https://zenodo.org/doi/10.5281/zenodo.8256824>, 2024. Publisher: Zenodo, Version 1.0. 2, 3
- [4] A. Berkani and A. Holzer. Clustering analysis of metered arabic poetry compositions. In *2023 20th ACS/IEEE International Conference on Computer Systems and Applications (AICCSA)*, pp. 1–8. IEEE, 2023. 1, 2, 4, 5
- [5] A. Berkani, A. Holzer, and K. Stoffel. Pattern matching in meter detection of arabic classical poetry. In *2020 IEEE/ACS 17th International Conference on Computer Systems and Applications (AICCSA)*, pp. 1–8. IEEE, 2020. 1, 2, 6
- [6] M. A.-A. Bhuiyan and A. R. Khan. Image quality assessment employing rms contrast and histogram similarity. *Int. Arab J. Inf. Technol.*, 15(6):983–989, 2018. 3
- [7] J. Brooke, G. Hirst, and A. Hammond. Clustering voices in the waste land. In *Proceedings of the Workshop on Computational Linguistics for Literature*, pp. 41–46, 2013. 2
- [8] T. Caliński and J. Harabasz. A dendrite method for cluster analysis. *Communications in Statistics-theory and Methods*, 3(1):1–27, 1974. 5
- [9] E. S. Dalmajer, C. L. Nord, and D. E. Astle. Statistical power for cluster analysis. *BMC bioinformatics*, 23(1):1–28, 2022. 5
- [10] D. L. Davies and D. W. Bouldin. A cluster separation measure. *IEEE transactions on pattern analysis and machine intelligence*, 1(2):224–227, 1979. 5
- [11] M. De Sisto, L. Hernández-Lorenzo, J. De la Rosa, S. Ros, and E. González-Blanco. Understanding poetry using natural language processing tools: a survey. *Digital Scholarship in the Humanities*, 39(2):500–521, 2024. 2
- [12] R. Delmonte. Visualizing poetry with sparsar visual maps from poetic content. In *Proceedings of the Fourth Workshop on Computational Linguistics for Literature*, pp. 68–78, 2015. 2
- [13] J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007. doi: 10.1109/MCSE.2007.55 3
- [14] D. M. Kaplan and D. M. Blei. A computational approach to style in american poetry. In *Seventh IEEE International Conference on Data Mining (ICDM 2007)*, pp. 553–558. IEEE, 2007. 2
- [15] L. Kaufman and P. J. Rousseeuw. *Finding groups in data: an introduction to cluster analysis*. John Wiley & Sons, 2009. 5
- [16] M. Lupea, A. Briciu, and E. Bostenaru. Emotion-based hierarchical clustering of romanian poetry. *Stud. Inform. Control*, 30:109–118, 2021. 2
- [17] N. Madnani. Emily: A tool for visual poetry analysis. *University of Maryland, Tech. Rep*, 2005. 2
- [18] R. Mazza. *Introduction to information visualization*. Springer Science & Business Media, 2009. 1
- [19] N. McCurdy, J. Lein, K. Coles, and M. Meyer. Poemage: Visualizing the sonic topology of a poem. *IEEE transactions on visualization and computer graphics*, 22(1):439–448, 2015. 2
- [20] L. Meneses and R. Furuta. Visualizing poetry: creating tools for critical analysis. *paj: The Journal of the Initiative for Digital Humanities, Media, and Culture*, 3, 2012. 2
- [21] A. Mittmann, A. von Wangenheim, and A. L. dos Santos. A multi-level visualization scheme for poetry. In *2016 20th International Conference Information Visualisation (IV)*, pp. 312–317. IEEE, 2016. 2
- [22] O. Musaoglu, Ö. Dag, Ö. Küçükakça, A. A. A. Salah, and A. A. Salah. A generic tool for visualizing patterns in poetry. In *DH*, 2017. 2
- [23] D. Oelke, D. Kokkinakis, and M. Malm. Advanced visual analytics methods for literature analysis. In *Proceedings of the 6th Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities*, pp. 35–44, 2012. 1
- [24] Z. Risha, D. Sonmez Unal, and E. Walker. Soliloquy: Fostering poetry comprehension using an interactive think-aloud visualization. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pp. 1–16, 2023. 2
- [25] P. J. Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65, 1987. 5
- [26] A. Šeļa, P. Plecháč, and A. Lassche. Semantics of european poetry is shaped by conservative forces: The relationship between poetic meter and meaning in accentual-syllabic verse. *Plos one*, 17(4):e0266556, 2022. 2
- [27] M. Stone. Choosing colors for data visualization. *Business Intelligence Network*, 2, 2006. 3
- [28] L. Venuti. Introduction: Poetry and translation. *Translation studies*, 4(2):127–132, 2011. 1
- [29] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 4