Eurographics Conference on Visualization (EuroVis) (2014)
N. Elmqvist, M. Hlawitschka, and J. Kennedy (Editors)

Comparing Three Designs of Macro-Glyphs for Poetry Visualization

Alfie Abdul-Rahman, Eamonn Maguire, and Min Chen

Oxford e-Research Centre, University of Oxford, United Kingdom

Abstract

Glyphs have been successfully used in poetry visualization for depicting the characteristics and positions of each phonetic articulation in relation to the human vocal system. While existing glyph designs provide visual representations for detailed observation and external memorization of the dynamics throughout a poem, they are less effective for observing the relationship and variance between different lines in a poem and in comparing different poems. In this short paper, we present three designs of macro-glyphs for summarizing the spatio-temporal dynamics at the level of poetic lines. In particular, we use statistics from a collection of poems to guide and optimize the designs. We report our comparative study on the effectiveness of these three designs.

1 Introduction

In poetry, close reading is a literary form of "data exploration", in which scholars pay concerted attention collaboratively to various linguistic, literary and sociological features, as well as to the interplay and relationships among these features [ARLC*13]. Poetry visualization can support close reading by enabling more effective observations and external memorization of a collection of features, allowing scholars to devote more cognitive capacity to explore different interpretations and hypotheses. This work focuses on the visualization of phonemic features, the details of which are difficult to see and remember because such features are by nature auditory and temporal, but not visual; and in English, they are indirectly encoded in the original text. Some poetry scholars refer the process of studying phonemic features in close reading as "chewing the sound".

In [ARLC*13], mini-glyphs were used to represent three phonemic variables of vowels and consonants, and the transition between two consecutive vowels or consonants. While these mini-glyphs are particularly useful for detailed observation and external memorization of the phonetic dynamics of a poem at the level of phonemes and words, they are less effective in observing the relationships between the different lines in a poem, e.g., the level of dynamics of a poem and individual lines, similarity, diversity and changes of phonetic structure at the level of lines. Hence it is desirable to provide visual representations at a level higher than the mini-glyphs in [ARLC*13].

In this paper, we propose to introduce *macro-glyphs* for encoding multivariate time series associated individual lines in a poem. We considered three main glyph designs, namely *static radial glyph*, *animated transitions*, and *static transitions with temporal highlight*. We used various statistical indicators of 35 poems (range from sonnets, nursery rhymes and free-verse) to guide our designs, and we evaluated the three designs by consulting humanities scholars.

2 Related Work

There is a large collection of previous work on text and document visualization. The dominant techniques have been statistics graphics (e.g., [VCPK09, AC07]), network visualization (e.g., [Mil95, WV08, vHWV09, Pal02, CCP09]),) and pixel-based visualization (e.g., [KO07, OBK*08]).

There have been several pieces of existing work on poetry visualization. For example, in [CGM*12] and [CAT*12], pixel-based visualization was used to display the poetic forms and variables, such as meter and rhyming patterns. In [MFM13], radial-based visualization was used to guide poem composition in a collaborative visualization. In [ARLC*13], a web-based interactive system was developed allowing scholars to display some 26 variables using a combination of visual channels. In terms of visualizing phonemic variables, apart from the mini-glyph used in [ARLC*13], line graphs, density plots, and animated trajectories were used to display two phonemic variables of vowels for visualizing sung vowels in music [FAW11]. As discussed earlier,

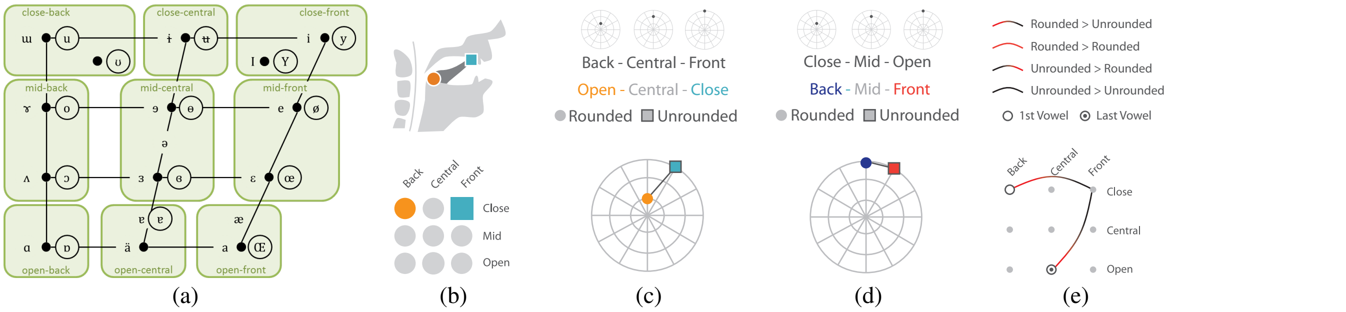

Figure 1: *(a) The vowel chart: (from left to right)* back, central *and* front *(which part of the tongue is raised) and (from top to bottom)* close, mid *and* open *(how far the tongue is raised). (b) A vowel position transition from a close back rounded to a close front unrounded. (c)-(e) Three different mappings of a vowel transition to the macro-glyph designs.*

mini-glyphs are ineffective in conveying phonetic dynamics at the level of poetic lines. Meanwhile, with line graphs, each poetic line would be displayed as a few time-series corresponding to different variables. Comparing different poetic lines would mean comparing different sets of multivariate time series. Density plots are effective for observing statistical patterns for a large number of vowels, for instance, in a poem or a piece of reasonably long text. Hence they are not suitable for visualizing phonetic dynamics at the level of poetic lines. In [FAW11], animated trajectories, which were referred to as “vowel worm”, were considered to be promising. We adapted this approach in one of our three designs.

The existing literature on visual encoding has offered us ample guidance in narrowing down our design options. Due to the constraints of the space, here we highlight those which have influenced our work the most. Bertin [Ber83] examined a set of basic glyph designs from a cartographer’s point of view. Ware [War12] discussed many perceptual considerations in visualization designs. Ward [War08] provided a technical framework for glyph-based visualization, covering aspects of visual mapping and layout methods, as well as addressing important issues such as bias in mapping and interpretation. Much of the existing work in the area of glyph-based visualization can be found in the recent survey by Borgo *et al.* [BKC*13]. Radial-based projection has been used to depict temporal data [DBS*11, WAM01, DH02] and the visualization of human movements [ZFAQ13]. Curved flow symbols [WSD11] have been used to visualize traffic movements. Time-series visualizations, such as ThemeRiver [HWN02], have been deployed to depict temporal patters in a large collection of messages or documents. Maguire *et al.* used statistics in the source data to guide the visual design [MRSS*12, MRSS*13]. Several authors, e.g., [RFF*08, Fis10], have discussed the use of animation in visualization.

3 Visualization and Design Approach

In this work, we consider only poems in English, and use the International Phonetic Alphabet (IPA) [Int99] to represent the sound and pronunciation of words in poems. We focus on vowel sounds, though the proposed designs can be easily adapted for consonants. Figure 1(a) illustrates three phonemic features of different vowels annotated using the IPA. The x-direction encodes the variable *frontness* or *backness*. The vowels displayed along the line on the left are the *back vowels*, in the middle the *central vowels* and on the right the *front vowels*. The positions correspond to the back, central, and front parts of the tongue respectively. Different vowels are produced when different parts of the tongue move up or down [Hou98, LD12]. The y-direction encodes the variable *vowel height* that is defined by both the mandible and the tongue. The vowels displayed along the top line are the *close vowels*, in the middle are the *close-mid vowels* and *open-mid vowels*, and at the bottom the *open vowels*. A vowel is said to be closed if the tongue is raised high and the mandibular is closed, and said to be open if the tongue is rested at the floor of the mouth and mandibular is open [Hou98, LD12]. The third variable is *lip rounding*, which refers to the shape of the lips during the production of a vowel [Hou98, LD12]. In Figure 1(a), all rounded vowels are shown in a circle.

Consider a line in a poem. Let L denote a multivariate time series, which is an ordered set of phonemic features of all vowels in the line, i.e., $L = \{(f_i, h_i, r_i) \,|\, i = 1, 2, \ldots\}$, where f_i defines the frontness (front, central, back), h_i defines vowel height (open, mid, close), and r_i defines lip rounding (rounded, unrounded). The common goal of the three designs to be described in the following sections is to encode L using a macro-glyph. Figure 2 shows two variations for each of the three designs.

Static Radial Glyph. In this design, we use clockwise angular positions to depict temporal ordering of vowels in a poetic line. In addition, we utilize three other visual channels, namely radial position, color and shape. Our first design decision is the number of radial lines in each glyph. With a circle, we may consider 4, 8, 12, 24, 30, and 60. Using our collection of poems, we compiled the statistics about the number of vowel phonemes per line. Among the 465 lines in the collection, 84.3% contain ≤ 12 vowel phonemes, while the maximum number of vowel phonemes is 22. Based on the statistics, we decided that 12 radial lines per circle is

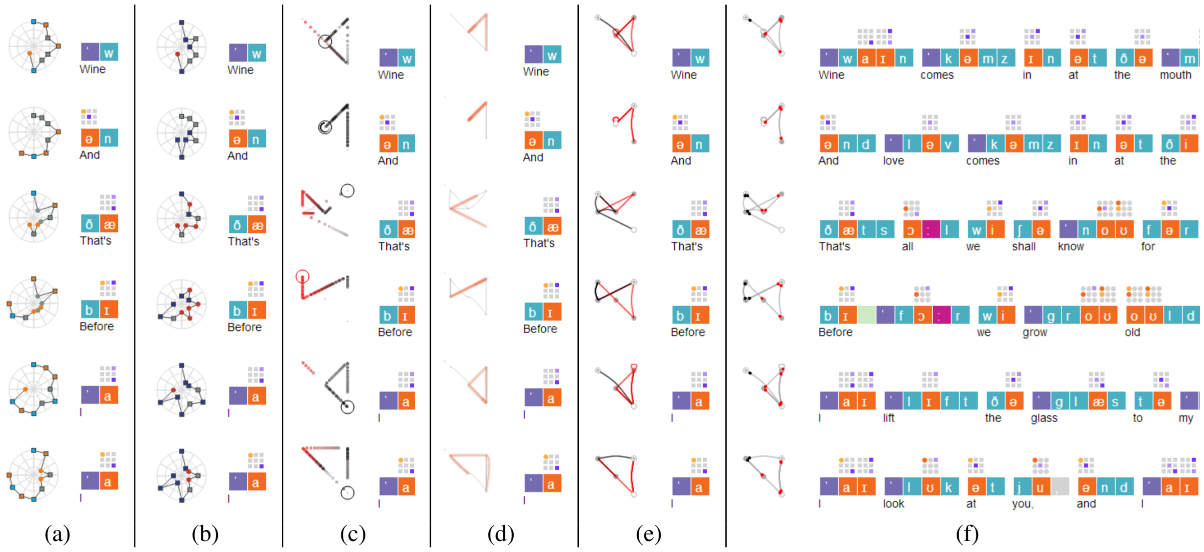

Figure 2: *Visualization of a short poem "A Drinking Song" by* W. B. Yeats *with macro-glyphs. (a) Static radial macro-glyph with a frontness layout while (b) static radial macro-glyph with a vowel height layout. (c)-(d) Animated transitions macro-glyphs with its variations (e)-(f) Static transitions with temporal highlight macro-glyphs with its variations.*

the most suitable option, because (i) it is consistent with the clock metaphor, (ii) it offers an appropriate density of radial lines, and (iii) most poetic lines (84.3%) need only 1 macro-glyph and 15.7% need 2. If 8 were chosen, 47.1% of the lines need 2 macro-glyphs and 2.4% need 3. If 24 were chosen, the radial lines would be too densely packed, while 84.3% would use only half of the circle.

Our second design decision is on the pairing of phonemic variables and visual channels. A few poetry scholars advised us during the design stage that vowel height and frontness may be more interesting than lip rounding. To help us decide the mapping, we calculated the average difference and standard deviation for vowel height (avg. diff.: 0.8704, std. dev.: 0.6699) and frontness (avg. diff.: 0.7709, std. dev.: 0.7073). According to the average difference, vowel height may have a slightly higher priority, while according to the standard deviation, frontness may be more favorable. As the statistics are inconclusive, we created two variations of the static radial glyph as illustrated in Figures 1(c,d) and 2(a,b). As rounding is considered least important by the poetry scholars, we assigned the color channel to frontness in Figure 1(c) and vowel height in Figure 1(d), while mapping lip rounding to the shape channel in both variations. We created a logological link between phonemic features and colors to help memorization, i.e., O for Open/Orange, C for Close/Cyan, F for Front/Fire and B for Back/Blue. Naturally, we used a circle for rounded and square for unrounded vowels.

Animated Transitions. We explored two variations of animated transitions, turbulence and beat, as shown in Figure 2(c,d). We used the abstract representation of 3×3 positions in Figure 1(a) as the 9 main reference points. When a vowel moves from one position to another, a transition line is drawn. For the turbulence variation, we trace out the movements using trailing circles as it moves across the line (similar to [FAW11]). We used two different colors to encode lip rounding, black (unrounded vowels) and red (rounded vowels). The color and gradient of a trace line encodes the relationship between the two consecutive vowels, e.g., a gradient line from red to black indicates the change from a rounded vowel to an unrounded vowel. For the beat variation, the transition lines are represented by a worm-like shape. We used line opacity to encode the frequency of a transition, i.e., the more transitions taking place between two positions, the darker the residual line becomes.

The advantage of this design is that only one macro-glyph is needed for any line. We anticipated some disadvantages of using animation. For example, it is hard to detect if there is no change in vowel positions. Nevertheless, as some scholars were enthusiastic about animation, we decided that it is better to include both variations in our evaluation.

Static Transitions with Temporal Highlight. This design uses the same 3×3 layout as the animated transitions, but place more emphasis on the transition lines and less on animation. All transition lines are static while a small marker moves along the transition lines to convey the temporal ordering. In Figures 1(e) and 2(e), three different circles indicate the first, last and intermediate vowel positions. A self-loop is used when there is a repetition of the same vowel position. Similar to the design of animated transitions, we used color and gradient of the transition lines to indicate four types of transitions. We also created a variation that encodes lip rounding by placing rounded or unrounded markers at the beginning and end of transition lines (Figure 2(f)).

	Static Radial Glyphs	Animated Transitions	Static Trans. w. Temp. Highlight
Helpfulness	Helpful in identifying similar phonetic patterns among poetic lines and differentiating dynamics and sound structures in poems.	Not easy to follow the movements. The residual line patterns do not convey temporal order or direction.	Helpful in observing similar phonetic patterns among poetic lines such as frequencies of certain phonemic features and transitions.
Helpfulness		Encoding the frequency of transitions using line opacities is helpful, but can be ambiguous.	Residual line patterns convey directional information and frequencies, but not temporal ordering.
Intuitiveness	Logological links between variables and colors help memorization.	Intuitive to observe the transitions in the flow of sound.	Intuitive to observe the transitions in the flow of sound.
Intuitiveness	Circles and squares are sufficient for differentiating lip shapes.	After a while, one tends to lose the tracking of the temporal ordering.	After a while, one tends to lose the tracking of the temporal ordering.

Figure 3: *Four humanities scholars compared the three designs of macro-glyphs.*

To help us determine the optimal paths to draw transition arcs and self-loops, we computed the statistics about the frequencies of transitions between different vowel positions, a copy of which is included in the supplementary materials. In our collection of poems, we found that there are no edges leaving or entering open-central and close-central vowel positions. This is likely because these vowel positions are more common in non-English languages, such as in Norwegian *butt* [bʉt] ‘blunt’ [Int99], or accented English dialect. Using the statistics, we created a lookup table for all 81 valid transitions, specifying the curve parameters for each transition.

4 Evaluation

Observing sound dynamics in poems requires a fair amount of general knowledge about linguistics and literature, specialized knowledge about poetry, and experience of close reading. A controlled user study with arbitrary participants may not reflect the presence of such knowledge and experience. We hence evaluated the three designs of macro-glyphs by consulting four humanities scholars who studied poetry in the past and had experience in close reading. The evaluation started with our explanation and live demonstration of the three designs and their variations. This was followed by a multiple-choice questionnaire with 43 questions. Among these, 6 were common questions that were repeated for each of the three designs and their variations, 9 questions were asked individually for specific designs, and 4 general questions were asked at the end, featuring comparisons across three designs. As all participants studied poetry at university level, we did not offer the “don’t know” option in the questionnaire based on the insight offered by [KHB*02, Kro91]. The structured survey was followed by free-form discussions. Participants were given opportunities to provide alternative answers to the questionnaire after the discussion. A qualitative summary of the scholars’ answers to the questionnaire is given in Figure 3, and their additional comments and suggestions are given below.

In general, domain experts found that macro-glyphs with animation are “fun to work with”, but preferred the static macro-glyphs as they provide a summary for each line while reading through the poem. Some commented that animating macro-glyphs simultaneously caused them to lose the tracking quickly.

The domain experts considered that no single phonemic feature is important by itself. Both designs of the static designs are considered to be equally important as different types of poems may exploit different parts of the vowel system. By having both variations, we can enable scholars to explore different hypotheses.

5 Conclusions and Future Work

This design study has shown that it is feasible to use macro-glyphs to encode a small amount of temporal information while depicting multivariate features. In visual analysis of phonetic dynamics of poems, it is important for a visual representation to support the observation and external memorization of temporal ordering, and directions of the movement. The domain experts prefer to have different visual representations for observing different types of dynamics (e.g., emphasizing vowel height or frontness), while having an essential requirement for observing different multivariate features at individual phonemes in a temporally connected manner. The study also confirms the disadvantages of animation at least from the perspective of analyzing phonetic dynamics in poems, while echoing the previous findings in [RFF*08, Fis10]. Our future work will focus on the visualization of temporal patterns of poetic lines while preserving the spatial context of phonetic features.

6 Acknowledgments

The authors wish to thank the humanities scholars participating in the user study: Franziska Kohlt, Laura Ludtke, David Robey, and Natasha Ryan, University of Oxford.

References

[AC07] ABBASI A., CHEN H.: Categorization and analysis of text in computer mediated communication archives using visualization. In *Proc. 7th ACM/IEEE-CS Joint Conf. Digital Libraries* (2007), pp. 11–18. 1

[ARLC*13] ABDUL-RAHMAN A., LEIN J., COLES K., MAGUIRE E., MEYER M., WYNNE M., JOHNSON C., TREFETHEN A., CHEN M.: Rule-based visual mappings – with a case study on poetry visualization. *Computer Graphics Forum 32*, 3 (2013), 381–390. 1

[Ber83] BERTIN J.: *Semiology of Graphics: Diagrams, Networks, Maps*. University of Wisconsin Press, 1983. 2

[BKC*13] BORGO R., KEHRER J., CHUNG D. H., MAGUIRE E., LARAMEE R. S., HAUSER H., WARD M., CHEN M.: Glyph-based visualization: Foundations, design guidelines, techniques and applications. In *Eurographics State of the Art Reports* (May 2013), EG STARs, pp. 39–63. 2

[CAT*12] CLEMENT T., AUVIL L., TCHENG D., CAPITANU B., MONROE M., GOEL A.: Sounding for Meaning: Analyzing Aural Patterns Across Large Digital Collections. In *Digital Humanities* (2012). 1

[CCP09] COLLINS C., CARPENDALE M. S. T., PENN G.: Docuburst: Visualizing document content using language structure. *Computer Graphic Forum 28*, 3 (2009), 1039–1046. 1

[CGM*12] CHATURVEDI M., GANNOD G., MANDELL L., ARMSTRONG H., HODGSON E.: Myopia: A Visualization Tool in Support of Close Reading. In *Digital Humanities* (2012). 1

[DBS*11] DROCOURT Y., BORGO R., SCHARRER K., MURRAY T., BEVAN S. I., CHEN M.: Temporal visualization of boundary-based geo-information using radial projection. *Computer Graphics Forum 30*, 3 (2011), 981–990. 2

[DH02] DRAGICEVIC P., HUOT S.: SpiraClock: A Continuous and Non-Intrusive Display for Upcoming Events. In *Extended Abstracts of CHI* (2002), pp. 604–605. 2

[FAW11] FROSTEL H., ARZT A., WIDMER G.: The vowel worm: Real-time mapping and visualisation of sung vowels in music. In *Proc. 8th Sound & Music Computing* (Jul 2011), pp. 214 – 219. 1, 2, 3

[Fis10] FISHER D.: *Animation for Visualization: Opportunities and Drawbacks*. Beautiful Visualization. O'Reilly Media, 2010, ch. 19, pp. 329 – 352. 2, 4

[Hou98] HOUSE L. I.: *Introductory Phonetics and Phonology: A Workbook Approach*. Psychology Press, 1998. 2

[HWN02] HAVRE S., WHITNEY P., NOWELL L.: ThemeRiver: Visualizing thematic changes in large document collections. *IEEE Trans. Visualization & Comp. Graphics 8* (2002), 9–20. 2

[Int99] INTERNATIONAL PHONETIC ASSOCIATION: *Handbook of the International Phonetic Association*. Cambridge University Press, 1999. 2, 4

[KHB*02] KROSNICK J. A., HOLBROOK A. L., BERENT M. K., CARSON R. T., HANEMANN W. M., KOPP R. J., MITCHELL R. C., RUUD P. A., SMITH V. K., MOODY W. R., GREEN M. C., CONAWAY M.: The impact of "no opinion" response options on data quality: Non-attitude reduction or an invitation to satisfice? *Public Opinion Quarterly 66*, 3 (2002), 371–403. 4

[KO07] KEIM D. A., OELKE D.: Literature fingerprinting: A new method for visual literary analysis. In *IEEE VAST* (2007), pp. 115–122. 1

[Kro91] KROSNICK J. A.: Response strategies for coping with the cognitive demands of attitude measures in surveys. *Applied Cognitive Psychology 5*, 3 (1991), 213–236. 4

[LD12] LADEFOGED P., DISNER S. F.: *Vowels and Consonants*, 3rd ed. Wiley-Blackwell, 2012. 2

[MFM13] MENESES L., FURUTA R., MANDELL L.: Ambiances: A Framework to Write and Visualize Poetry. In *Digital Humanities* (2013). 1

[Mil95] MILLER G. A.: WordNet: A lexical database for English. *Communications of the ACM 38*, 11 (Nov. 1995), 39–41. 1

[MRSS*12] MAGUIRE E., ROCCA-SERRA P., SANSONE S.-A., DAVIES J., CHEN M.: Taxonomy-Based Glyph Design - with a Case Study on Visualizing Workflows of Biological Experiments. *IEEE Trans. Visualization & Comp. Graphics 18*, 12 (2012), 2603 – 2612. 2

[MRSS*13] MAGUIRE E., ROCCA-SERRA P., SANSONE S.-A., DAVIES J., CHEN M.: Visual Compression of Workflow Visualizations with Automated Detection of Macro Motifs. *IEEE Trans. Visualization & Comp. Graphics 19*, 12 (2013), 2576–2585. 2

[OBK*08] OELKE D., BAK P., KEIM D., LAST M., DANON G.: Visual evaluation of text features for document summarization and analysis. In *IEEE VAST* (Oct. 2008), pp. 75 –82. 1

[Pal02] PALEY W. B.: TextArc, 2002. URL: `http://www.textarc.org/`. 1

[RFF*08] ROBERTSON G., FERNANDEZ R., FISHER D., LEE B., STASKO J.: Effectiveness of animation in trend visualization. *IEEE Trans. Visualization & Comp. Graphics 14*, 6 (Nov 2008), 1325–1332. 2, 4

[VCPK09] VUILLEMOT R., CLEMENT T., PLAISANT C., KUMAR A.: What's being said near "Martha"? Exploring name entities in literary text collections. In *IEEE VAST* (2009), pp. 107–114. 1

[vHWV09] VAN HAM F., WATTENBERG M., VIEGAS F. B.: Mapping text with phrase nets. *IEEE Trans. Visualization & Comp. Graphics 15*, 6 (Nov. 2009), 1169–1176. 1

[WAM01] WEBER M., ALEXA M., MULLER W.: Visualizing time-series on spirals. In *IEEE Symp. Information Visualization* (2001), pp. 7–13. 2

[War08] WARD M. O.: Multivariate data glyphs: Principles and practice. In *Handbook of Data Visualization*. Springer, 2008, pp. 179 – 198. 2

[War12] WARE C.: *Information Visualization: Perception for Design*, 3rd ed. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2012. 2

[WSD11] WOOD J., SLINGSBY A., DYKES J.: Visualizing the dynamics of london's bicycle-hire scheme. *Cartographica 46*, 4 (2011), 239–251. 2

[WV08] WATTENBERG M., VIÉGAS F. B.: The Word Tree, an interactive visual concordance. *IEEE Trans. Visualization & Comp. Graphics 14*, 6 (Nov. 2008), 1221–1228. 1

[ZFAQ13] ZENG W., FU C.-W., ARISONA S. M., QU H.: Visualizing interchange patterns in massive movement data. *Computer Graphics Forum 32* (2013), 271–280. 2