Interactive Visualisation of Debater Identification and Characteristics

Matt FOULIS [a], Jacky VISSER [a], Chris REED [a]

[a] *Centre for Argument Technology, University of Dundee, United Kingdom*

Abstract. We present an interactive visualisation for *Dialogical Fingerprinting*, a system for automatically identifying speakers and classifying dialogue characteristics. Presenting this information to a debate audience in real-time alongside debate playback provides an overview of how the speaker identity and role predictions change over the course of the debate and how the participants of the debate relate to each other in terms of emotionality and ideology. The touchscreen-based visualisation program allows users to control debate playback, view the classification results, and make comparisons between system configurations.

Keywords. BBC Moral Maze, Debate Visualisation, Debate Analytics, Speaker profiling

1. Introduction

Argument Analytics leverage the wealth of quantitative data available to provide deeper insights into the dynamics and content of argumentative texts and debates [1]. Recent research in this field includes work on providing members of the public with real-time visualisations of debates in order to enhance engagement with the topic under deliberation [2]. By providing additional information about a debate, the audience is better able to understand the content being discussed, the stances of the participants, and the arguments they advance. Related research has also been carried out into the use of 'second screens' in which an additional device, such as a smartphone, is used whilst observing a televised debate [3].

With this aim of providing further insight into a debate using Argument Analytics, we present an interactive visualisation program for Dialogical Fingerprinting [4], a system which uses the unique dialogical characteristics of participants within a debate to automatically identify who's speaking and what role they play in the debate. The interactive program allows listeners of the debate to explore the predictions made and see how the system's performance develops in real-time over the course of a debate. Displaying this information in an intuitive way aims to provide the audience of a debate with a better understanding of the proceedings and the dialogical characteristics of the participants.

2. Automatic Debater Identification and Characteristics

The Dialogical Fingerprinting system uses the unique dialogue characteristics of speakers in a debate, such as sentence length and vocabulary to build a dialogical fingerprint for each participant. These characteristics are used as features in a Machine Learning approach to perform speaker identification and role classification for each dialogue turn in a debate. The system was trained using transcribed episodes of the BBC Radio 4 programme The Moral Maze, in which a set of guest experts, debate panellists, and a debate moderator discuss a topical, divisive issue. The dataset contains an unbalanced distribution of 93 speakers, with some participants appearing in each episode, and others appearing only once.

Dialogical Fingerprinting (using a Support Vector Machine model) achieves a macro f1 of 1.0 for the task of role classification, correctly classifying the role of participants in the debate as either: Moderator, Panellist, or Witness. The system achieves a macro f1 of 0.52 for identifying the individual speaker on the basis of their debate contributions. The system also measures the relative emotionality of a speaker turn, using the subjectivity lexicon developed by Wilson et al [5]; and a relative ideological scaling, using the unsupervised methods developed by Glavas et al. [6]. Outputs for both the emotionality and ideological positions of speaker turns are scaled to a 0 to 1 interval.

3. Interactive Visualisation of Dialogical Fingerprinting

To demonstrate the Dialogical Fingerprinting system, we developed a touch-based graphical user interface (GUI) (Figure 1). The visualisation program aims to provide the audience of a debate with an intuitive, user-friendly, way of interacting with the results of the system, as well as allowing a comparison between the system configurations - such as, the Machine Learning model used by the system, or the specific set of features used by the model. To this end, a focus is placed upon the 'real-time' nature of the demonstration. By showing the audience members how the predictions change during episode playback, they can see how the performance of the speaker identification and role classification develops, as well as how the relative emotionality and ideological positions of the participants becomes clear as further speaker turns are conducted. To allow for this real-time report of the analytics, the Dialogical Fingerprinting system was adapted to output a performance report after each turn within a dialogue. This information is then matched to the episode transcript timestamps, with each speaker turn being associated with the model performance at that time. The audience can interact with the visualisation program to explore the results achieved.

The environment for the program is a large kiosk style device with a 54 inch touchscreen (Figure 2). The touchscreen interface allows the user to directly manipulate the onscreen elements. For example, by tapping on an element to make a selection, or pressing and dragging to adjust the audio playback position. To develop the interface and accommodate the touch interaction a suitable programming framework was required. The Qt Framework was chosen[1]. For lower level

[1]available at: https://www.qt.io

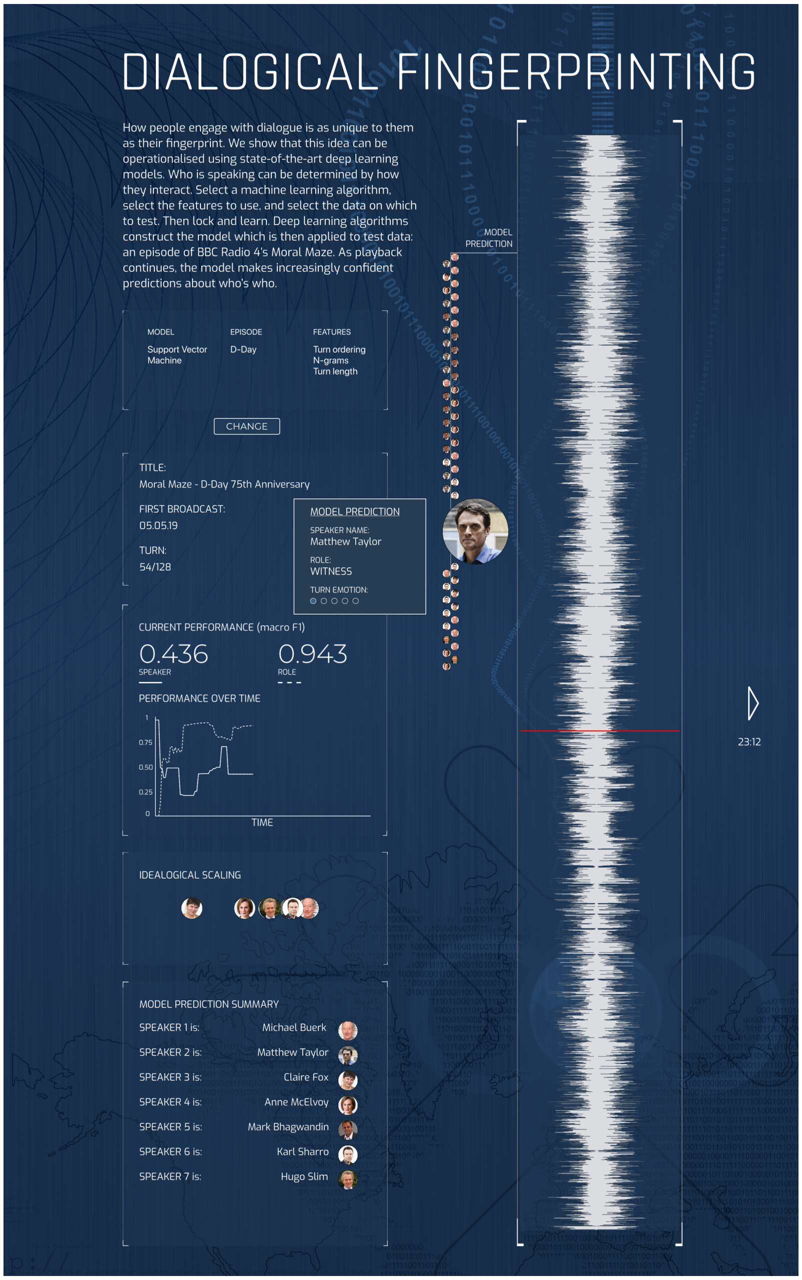

Figure 1. The interactive dialogical fingerprinting demonstration program. On the left, from top to bottom: a text introduction to the work; the user-selection of model, Moral Maze episode, and dialogical features; details of the selected Moral Maze episode; the current model performance and performance over time; the relative ideological scaling of the episode participants; and a summary of the participant predictions. In the middle section, a list of the previous and current speaker turns are represented with images of the speakers. Selecting a turn presents speaker and role predictions and turn emotionality. On the right, an audio waveform representation of the selected episode is shown which can be used to navigate to a specific point within an episode.

Figure 2. The interactive Dialogical Fingerprinting demonstrator in use.

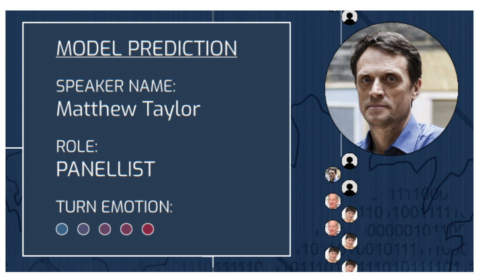

Figure 3. A selected speaker turn from the turn list showing the current predictions for speaker identity and role as well as the relative emotionality level of the turn.

processes, such as file manipulation and audio playback, the Qt C++ implementation is used. For visualisation aspects, the Qt QML language is used, exploiting the Qt concepts of *signals* and *slots* to facilitate communication between the two.

To initiate the program, the user is required to select: one of several pre-trained Machine Learning models, a Moral Maze episode to run the model on, and one or more dialogical characteristics (e.g. turn length or n-grams). Providing this choice allows a comparison to be made between configurations, for example, between the results obtained with different Machine Learning models, producing different results and hence measured performance during the course of an episode. Once selected, playback of the chosen episode begins.

As the episode progresses, the user is presented with information about the episode and the current performance of the selected model. A list of speaker turns is shown, allowing individual turns to be selected, presenting the model's

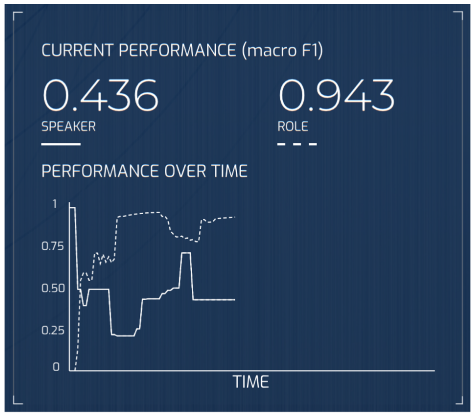

Figure 4. The macro f1 scores for the speaker identity, 0.436, and role predictions, 0.943, at the current point in the selected episode. The change in macro f1 score over the course of the episode is displayed on a line graph.

current best prediction for the speaker's identity and role (Figure 3). Also shown is the relative emotionality level of the selected turn. This is indicated using an intensity indicator style element, where a set of circle images are filled with colour relative to the emotionality level of the turn. If at a later point in the episode, when additional speaker turns have been performed, a speaker or role prediction is changed, the corresponding turn image and text are updated accordingly.

The model performance for the current point in the episode is presented as the macro F1 score for both the individual speaker identification and role identification (Figure 4). The change in macro F1 performance over time is displayed using a line graph, visualising how the model prediction performance improves as the episode progresses and more data becomes available for the model to base its predictions on.

The relative ideological position of each of the speakers is displayed by grouping speaker's pictures on a relative scale (top of Figure 5). When a speaker first enters the debate, their portrait is positioned at the centre of the scale. As the debate proceeds, the picture is moved left or right, relative to the other participants. This provides the audience with a visual representation of how the ideological positions of the participants relate to each other, with similar participants being grouped together, and opposing participants being further apart.

A summary of the speaker predictions shows for each participant within the debate, the current best prediction for their identity, with a corresponding picture (bottom of Figure 5). The speaker pictures for both the ideological scaling and summary elements are changed in the same manner as the turn list in cases of an updated prediction. When a new speaker identity classification is output from the model, the corresponding pictures are updated.

Using the waveform representation of the episode (shown on the right in Figure 1), the user is able to select a position which moves the audio playback to the corresponding point. This allows for quick comparisons to be made between different points in time; for example, comparing the initial predictions of role

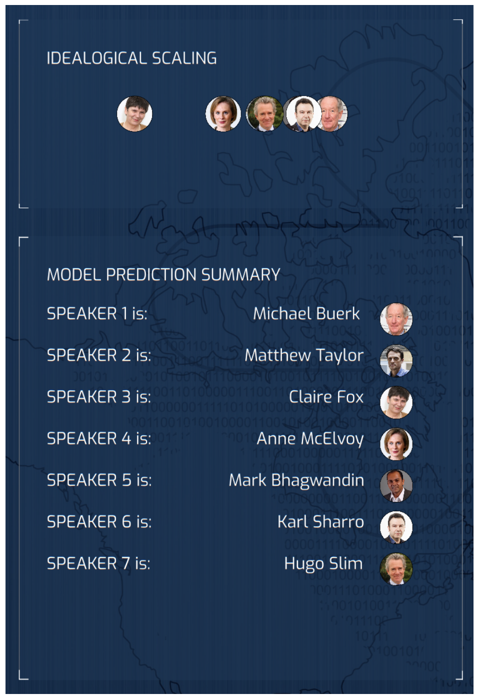

Figure 5. Two elements showing the relative ideological scaling of the participants within an episode, and a summary of the speaker identification predictions.

or speaker identity under the sparsity of data available at the beginning of an episode to those once more data has been gained towards the end of the episode.

4. Conclusion

Within a dialogue, participants exhibit particular communicative characteristics which are unique to them. With *Dialogical Fingerprinting*, these characteristics are used to identify participants and provide informative debate analytics. The interactive demonstration program for this system provides audiences with an intuitive way of engaging with the outputs, giving further insights into the debate and helping to increase the understandability of the work being carried out.

Despite the work being at a preliminary stage, and as yet without quantitative user evaluation, this prototype clearly demonstrates the potential for large format touch based interaction with debate analytics.

5. Acknowledgements

This work was supported in part by Dstl, the Defence Science & Technology Laboratory in the UK, under the Defence And Security Accelerator (DASA) grant ACC6005123.

References

[1] J. Lawrence, R. Duthie, K. Budzynska, and C. Reed, “Argument analytics,” in *Computational Models of Argument* (P. Baroni, T. Gordon, T. Scheffler, and M. Stede, eds.), vol. 287 of *Frontiers in Artificial Intelligence and Applications*, (Netherlands), pp. 371–378, IOS Press, 2016.

[2] B. Plüss, M. El-Assady, F. Sperrle, V. Gold, K. Budzynska, A. Hautli-Janisz, and C. Reed, “ADD-up: Visual analytics for augmented deliberative democracy,” in *Computational Models of Argument - Proceedings of COMMA 2018* (S. Modgil, K. Budzynska, J. Lawrence, and K. Budzynska, eds.), (Netherlands), pp. 471–472, IOS Press, 2018.

[3] K. Gorkovenko and N. Taylor, “Audience and expert perspectives on second screen engagement with political debates,” in *Proceedings of the 2019 ACM International Conference on Interactive Experiences for TV and Online Video*, TVX 19, (New York, NY, USA), p. 7082, Association for Computing Machinery, 2019.

[4] M. Foulis, J. Visser, and C. Reed, “Dialogical fingerprinting of debaters,” in *Computational Models of Argument. Proceedings of COMMA 2020*, in press.

[5] T. Wilson, J. Wiebe, and P. Hoffmann, “Recognizing contextual polarity in phrase-level sentiment analysis,” in *Proceedings of HLT and EMNLP*, p. 347354, ACL, 2005.

[6] G. Glavas, F. Nanni, and S. P. Ponzetto, “Unsupervised cross-lingual scaling of political texts,” in *Proceedings of the 15th EACL*, pp. 688–693, ACL, apr 2017.