Developing Colour Sequences for High Dynamic Range Data

Matthew Tobiasz, Amanda Henderson, Sheelagh Carpendale, Alan Dunning, and Paul Woodrow

Abstract—We consider the problem of developing colour sequences that offer the best dynamic range when used for pseudo-colouring. We began with nine colour sequences currently suggested for visualization and informally compared them to discover which was subjectively seen to have the most pleasing visual aesthetic. This comparison indicated that different colour sequences are more appropriate for different data ranges and led to the development of three new colour sequences. Again consulting our large group of collaborators we consider which of the twelve colour sequences had the greatest dynamic colour range.

1 Introduction

For many types of data choosing an appropriate colour range is important for effective visualization [2, 4, 6]. We present an exploration of several existing colour sequences [3] and introduce three new sequences with the goal of most clearly revealing the greatest dynamic range of pseudo-coloured astrophysics data. This work is part of a research initiative into the visualization of large astronomical data sets, conducted in collaboration with computer scientists, astronomers, and media artists. Here we focus on our initial findings into pseudo-colouring this data. Our research concentrated on visualizing data from galactic spectral line emissions, gathered as part of the Canadian Galactic Plane Survey (CGPS) [5]. We used a relatively small subset of the entire CGPS data, measuring 879 megabytes in size.

Book Rdx is the software we developed to visualize this CGPS data set. This spectral line emission data is a 3D data set where spectral intensity, measured as velocity, is a function of galactic longitude, galactic latitude, and frequency range. Since all three axis contain contiguous values, the complete data set can be conceptually modelled as a single 3D volume. By filling the 3D volume with spectral intensity values. The contents of the data cube can be seen by cutting along a plane, aligned with either of the data cubes two primary axes and rendering only the area on the cutting plane. As in "Interactive Video Cubism" by Fels et al. [1], by sequentially moving this cutting plane up and down through the data cube and interpolating between contiguous values, a fluid animation can be made.

2 Pseudo-Colour

During visualization, scalar CGPS spectral intensity values are first normalized and then transformed to colour vectors, using the GPU for efficiency. Our investigation into the use of pseudo-colouring began with an assessment of this transformation. Here we considered globally and locally optimized colour, using a set of nine colour sequences as suggested by Levkowitz et al.[3]: nonlinearized gray, linearized gray, heated object, optimal colour scale, linearized optimal colour scale, rainbow, blue to cyan, blue to yellow, and magenta.

These sequences were first compared to discover which was subjectively seen to have the most pleasing visual aesthetic. We then introduced three additional colour sequences and this time considered which sequence revealed the greatest dynamic range.

2.1 Globally and Locally Optimized Colour

In globally optimized colour the complete data set is normalized relative to all values in the data set. This allows for a direct visual comparison of all resulting renderings since the scaling parameters to the normalization function are common to all values, providing a consistent colourization transformation. An assigned colour always maps to a constant scalar magnitude, no matter were this value resides in the data cube.

When the CGPS data set is normalized globally, the fundamental limitation of the resulting visuals is a pronounced lack of image contrast. This situation is due to the drastic difference in the mean spectral intensity values between cutting planes versus that of the global mean. The global colour would be appropriate if the complete CGPS data set were to be rendered in a single instant and presented in its entirely, however this is not the case since only a small slice of the data set is seen at a given instant of time. The fluctuations in the mean value of a cutting plane slice result in scalar to colour vector mappings which produce a poor utilization of a rich colour sequence.

In locally optimized colour, normalization of scalar spectral intensity values occurs on a per cutting plane basis, relative only to the values on the current cutting plane. By normalizing exclusively to those values to be visualized, the mapping of scalar to colour vector changes between positions of the cutting plane in the data cube. This means that direct visual comparison cannot be done between images created at discrepant discrete locations in the data set. However the result of the colourization process yields a vastly superior dynamic colour range for a given plane. Figure 1 illustrates the prominent difference between the global and local optimization. Local optimization was chosen for visualizing the CGPS data as it accentuated the subtle changes that occur within individual subsets of the data.

2.2 Considering Colour Aesthetics

Particularly for the artists in our collaboration, attention to the visual aesthetic of the final rendering was of concern. To evaluate the aesthetics of our visualization, we consulted with our larger group of collaborators who had not been part of the colour choices up to this point in time. The people's declared backgrounds were: four from art, seven from computer science, and one undeclared. We asked which were the three most aesthetically pleasing colour sequences, rated from highest to lowest, and which was the single worst colour sequence aesthetic.

The results indicated an interesting dichotomy between the preferences of people with art and computer science backgrounds. Monochromatic colour sequences where chosen exclusively by those with an artistic background. This same group also unanimously chose the rainbow colour sequence as the most displeasing aesthetic. In contrast, none of the people with a computer science background though the rainbow colour sequence was the single worst. The linear grey colour sequence was most favourable by all, with heated object and magenta as the next runners up.

2.3 High Dynamic Range Colour Sequences

Since different colour scales applied to identical data sets accentuate discrepant ranges within the data, one can imagine an ideal generic colour scale would have no bias towards any particular range within its gamut yet offer the greatest ordinal dynamic colour range. With this goal in mind, we developed three additional colour sequences to augment the existing nine. These three new sequences are shown in Fig-

- *Tobiasz, Carpendale: Department of Computer Science, University of Calgary: tobiasz@ucalgary.ca and sheelagh@cpsc.ucalgary.ca.*
- *Henderson, Dunning: Department of Media Arts and Digital Technology, Alberta College of Art and Design: amanda.henderson@acad.ca and einsteins-brain-project@shaw.ca.*
- *Woodrow: Department of Art, University of Calgary: ebp@shaw.ca.*

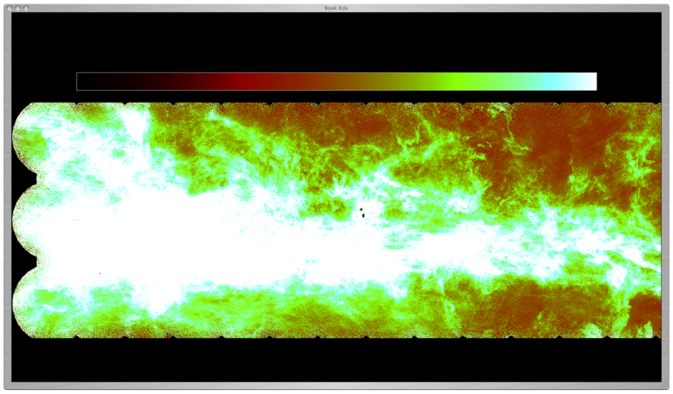

(a) Globally optimized colour.

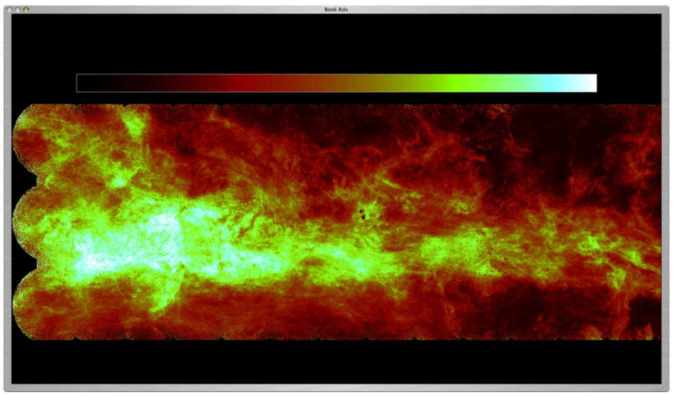

(b) Locally optimized colour.

Fig. 1. Comparison between the globally and locally optimized colour on identical sets of scalar values, both using the linearized optimal colour scale.

ure 2. The first two sequences, AM1 (an example of its usage shown in Figure 3) and AM2, were artistically inspired. The third sequence, Flow, was created based on colour sequence design recommendations by Ware, in which each colour in the scale "provides a kind of upward spiral in color space; each color is lighter than the preceding one" [6, page 132].

2.4 Improving the Colour Sequence's Dynamic Range

With the additional colour sequences, we once again consulted our larger group of collaborators, with the intent to evaluate the three best colour sequences from the group of twelve (again rated from highest to lowest) that have the greatest dynamic colour range; and again also selecting the colour sequence perceived to have the single worst dynamic colour range. Thirteen people where involved, all of whom had a background in computer science. The newly introduced colour sequences AM1 and Flow were highly favoured for the produced dynamic colour range with the CGPS data set, which was encouraging positive feedback. The feedback also indicated that the optimized colour scale was seen to perform most poorly for this task of revealing dynamic range, which was a surprise.

(a) AM1.

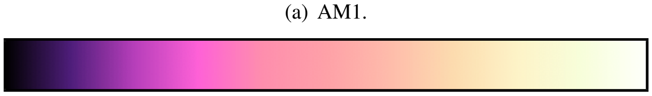

(b) AM2.

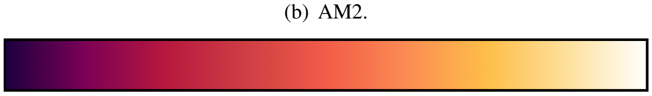

(c) Flow.

Fig. 2. Additional colour sequences.

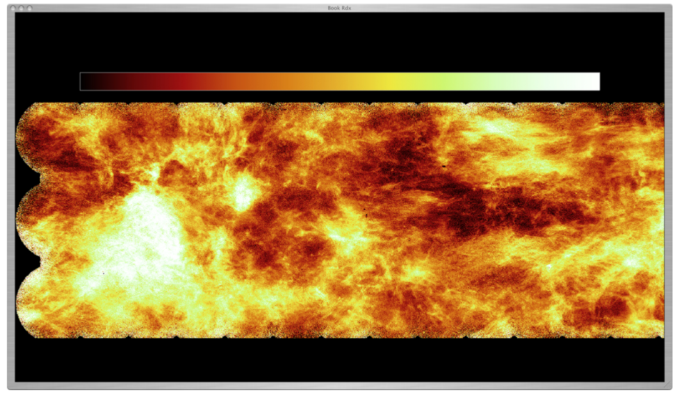

Fig. 3. Example of AM1 colour sequence.

3 Future Work

The next step in this research will be to formally evaluate the colour sequences, possibly a survey consisting of a large and diverse population sample. Interesting aspects include considering characteristics of colour sequence that best reveal the dynamic range locally, within a given sequence as well as globally across a large data set. Furthermore applying and evaluating our developed colour sequences to visualizations outside of the astrophysics is of interest.

Acknowledgements

The authors would like to thank Maureen Stone, and all the instructors and students of the 2006-07 *Art Science and Technology* course, offered conjointly by the University of Calgary, Alberta College of Art and Design, and The Banff Centre.

References

[1] S. Fels and K. Mase. Interactive video cubism. In *Proceedings of the 1999 Workshop on New Paradigms for Interactive Visualization and Manipulation (NPIVM '99)*, pages 78–82, New York, NY, USA, Computer Science 1999. ACM Press.

[2] C. G. Healey. Choosing effective colours for data visualization. In *VIS '96: Proceedings of the 7th conference on Visualization '96*, pages 263–ff., Los Alamitos, CA, USA, 1996. IEEE Computer Society Press.

[3] H. Levkowitz and G. T. Herman. Color scales for image data. *IEEE Comput. Graph. Appl.*, 12(1):72–80, 1992.

[4] M. Stone. *Field Guide to Digital Color*. A. K. Peters, Ltd., Natick, MA, USA, 2002.

[5] A. R. Taylor, S. J. Gibson, M. Peracaula, P. G. Martin, T. L. Landecker, C. M. Brunt, P. E. Dewdney, S. M. Dougherty, A. D. Gray, L. A. Higgs, C. R. Kerton, L. B. G. Knee, R. Kothes, C. R. Purton, B. Uyaniker, B. J. Wallace, A. G. Willis, and D. Durand. The canadian galactic plane survey. *The Astronomical Journal*, 125:3145–3164, 2003.

[6] C. Ware. *Information Visualization: Perception for Design 2nd Ed.* Morgan Kaufmann, 2004.