Visualisation and Metavisualisations: Helping a User with Colour Gamut Mapping*

Philip K . Robertson
Canon Information Systems Research, Australia

Abstract

Why do so many visualisations make poor or ineffective use of the colour gamut of a device? The trend with colour management systems is to increasingly automate gamut mapping operations so that the user does not need to know about gamut constraints of devices being used. Is this well-advised, bearing in mind the complex set of decisions needed for any given set of data, task, and gamut characteristics? An alternative approach is to stimulate implicit learning about device gamut capabilities, so that users gain the awareness and skills required to make sensible gamut mapping decisions that are appropriate for the task being undertaken. This paper describes the use of metavisualisations - interactive visualisations of perceptual colour gamuts that describe the colour mappings applied to generate a visualisation—to help users select appropriate colour mappings. The scope for applying incremental intelligence to these metavisualisations to support automation of gamut mapping will also be described. Such an approach is consistent with the overall process of visualisation, being a strategy of helping a user recognise salient characteristics, and develop and apply expertise, rather than supplanting the expertise of the user with automated techniques.

Introduction: The Need for User Support for Gamut Mapping Decisions

Users of visualisation systems need some form of help or guidance in the selection of colour schemes for visualisation, or in the selection of gamut mapping strategies. This is because few users are explicitly aware of the constraints of particular colour device gamut shapes, or of perceptual factors that influence effective interpretation of colour visualisations. Existing systems offer default based guidance at best, although some work on rule-based approaches have been taken. Methods of introducing guidance in choosing colour sequences, and gamut mapping, are therefore needed. More broadly, approaches to introducing guidance in presentation or visualisation design have ranged from sets of guiding principles,[4,12] compositional systems,[5,13,3,1,2] rule-based systems,[15,11] case-based examples and methodologies based around types of information content.[7] Rules and recipes do not provide sufficient support; they require some form of check-list, and have the drawback that the number of possible criteria to evaluate becomes too large for easy operation. Rule-based systems seem to offer the most viable framework for incorporating guidance but also run the risk of becoming unwieldy or requiring complex formulations. The problem is that the search space can be very large, and it can be difficult to reach a useful subset of possibilities easily. In this paper we explore an approach for incorporating guidance that is based on metavisualisations—graphical representations of operations associated with constructing a visualisation.

Metavisualisations provide a visualisation that tells the user something about the visualisation, in the same way that metadata provides information about data. Autonomous agents can then perform well-defined tasks associated with a metavisualisation. This allows an incremental approach to incorporating intelligence, avoiding the large search space of a more generalised rule-based system. The approach is based on the experience of building a system for interactive choice of colour mappings within a perceptual colour space, exploiting metavisualisations for specific well-defined aspects of the mapping.[8,9] Perceptual colour spaces offer a framework for specifying and controlling colour based on human colour perception.[6] They arise from the combination of colour order systems based on major perceptual dimensions of colour and experiments or predictions of the perceived size of colour differences. They are used increasingly as a basis for colour selection and enhancement across a range of disciplines

Metavisualisations in Gamut Mapping

A typical metavisualisation shows a graphical representation of the perceptual colour gamut of a colour device, for example using 2D cross-sections of constant lightness and constant hue through the perceptual colour gamut of the device. Such representations have been fairly widely used for visualising colour gamuts, and colour sequence paths through colour gamuts for some time now. Interactive controls can allow exploration and interrogation of such gamut representations. The gamut representation relies on having modelled the target display device in perceptual terms, either physically or numerically, so that perceptually specified colours can be generated from the device controls. The gamut boundary thus reflects the true gamut, in perceptual terms, for the target device. Generally a colour management system underpins this capability. Such a representation, used as an interactive interface for mapping data into colour, encourages implicit learning because the user develops an understanding of different display capabilities in terms of their perceptual colour gamuts.[8] The gamut boundary in the representation depends on the particular display; it hence determines scaling and other factors that must be applied to chosen data mappings. These actual display gamut representations also show clearly why gamut mapping, the process of mapping data from the gamut of one display to that of another with different characteristics, is a non-trivial problem whose solution may well be application dependent.[16] Therefore flexible, task-dependent strategies for gamut mapping are necessary, as are flexible approaches to mapping data points that fall outside the gamut in a particular specified mapping.

A key element of this approach is that each metavisualisations encapsulates a core operation necessary for choosing and performing the mappings of data into colour space. Associated with each metavisualisation is a set of parameters that must be set (by default or interactively) for the operation to be completed. Interaction must be constrained to be perceptually sensible for the operation. Recognising the data type and dimensionality, and offering data mappings that make sense for such data types, results in an operational system that provides colour mapping and control capabilities significantly beyond those of standard visualisation tools. Allowing the user to interact with mappings directly, in terms of the colour gamut of the target device, levers an understanding of devices and their limitations, and provides implicit guidance in making data mapping choices.

Metavisualisations and Intelligence

The metavisualisations provide a complete and encapsulated description of methods of mapping data into colour gamuts. These methods are designed along perceptual lines, and therefore offer constrained interaction, where the constraints ensure that arbitrary mappings do not result. Thus if appropriate metavisualisations are designed to represent schematically any required colour mapping, or more generally visualisation technique, a set of encapsulated mechanisms for performing mappings results. The encapsulation is important; it means that there is a limited set of requirements for invoking a metavisualisation.

Metavisualisations may hierarchically invoke lower level metavisualisations, corresponding to decisions about the mapping at a lower level in the overall decision tree. All the information necessary to complete the task for a metavisualisation can be obtained from one or more of the device, data, user, or project models. This includes, for example, the actual gamut definition for a display device, the information carrying capability of specific data mappings, or suitable default tie points for determining data mapping such as percentile points on a data histogram.

The metavisualisation approach can also be used to encourage a work pattern that is result-driven. The design starts with an unfinished "result", a specification of a scene to be visualised, which needs to be filled in with certain decisions. As these decisions are made they lead to more decisions, but each one can be made in the context of what went before it, helping the user to understand why each of the decisions has to be made, which in turn gives guidance on how to make it. The work pattern can also be top-down. The major decisions are made first, followed by increasingly less important ones. This means that if the less important decisions can be made automatically, they may never need to be presented to the user, which reduces the amount of detail that needs to be understood.

In this approach, a metavisualisation can have associated with it a set of autonomous agents whose task it is to find, or deduce, the required information to complete the task for a metavisualisation. When a metavisualisation is first instantiated, its agents work to establish sensible values for as many of its parameters as possible. This information is gathered from the system models, from other active metavisualisations, or deduced in a localised manner. Agents can monitor the current state and update their decisions accordingly. Thus intelligence can be added in-crementally around the operation of any particular metavisualisation. One of the key benefits of this approach is that the user remains in control of the level of automation provided by the system. The actions taken by the agents are clearly presented at the user interface, and can be accepted or changed at any time by the user. Locally made decisions can be reversed, while retaining full automation in other areas.

Binding Framework—Reference Models

Underpinning the approach must be a formalised reference model for the visualisation process, including models for critical components. The reference model need not include all possible operations initially; rather it must have sufficient flexibility to allow extensions as required. But the core tenet - that knowledge about devices, data, interpretation and tasks can be sufficiently described by the models-requires a degree of formalisation beyond that of most current systems.

A reference model needs at least the following:

Data Model: the data in the system needs to be stored using a rich, flexible, data model, that provides sufficient metadata to allow sensible decisions to be made about choice of visualisation techniques based on interpretation aims. The data model also needs to allow efficient and flexible access to all stored data.

User Model: the user model provides information about conventions and preferences of the current user of the visualisation system at several levels, ranging from domain knowledge, to project knowledge, to an individual user profile.

Visualisation Model: the visualisation model needs to describe the range of visualisation techniques that are available, and the interpretation potential of each individually, and in combinations. Each available technique defines a set of parameters that must be provided for the technique to be used. The visualisation model may also include information on the computational expense of various techniques.

Device Model: the device model must describe the capabilities and peculiarities of each available visualisation device. The use of gamut boundary information has already been described. Spatial, spectral and temporal resolution are also important aspects of device behaviour. The device model must allow devices to be addressed in perceptual terms for effective visualisations to be possible.

Task Model: the task model identifies the aims of the visualisation being designed, in terms of what elements from the data model need to be displayed, and what types of information are needed from them.

Defining a full reference model is a substantial task. The component models, however, form the basis for knowledge representation for the entire approach, and effective operation of the agents associated with metavisualisations depends on this.

Summary

This paper has described the development and use of metavisualisations—graphical representations of gamut mappings used in generating a visualisation—to help a user learn about colour gamut constraints and make suitable gamut mapping decisions. Perceptually sensible, well defined interactions with metavisualisations offer a basis for introducing incremental guidance into the use of colour in visualisations. This may be a more promising approach to incorporating intelligence than using rule-based systems covering the entire range of design decisions, although rule-based decision making can be incorporated at a level local to a particular metavisualisation. In other words, the decision tree is determined in part by pre-designed mapping strategies embedded in metavisualisations, rather than being fully contained in an expert system knowledge base. The advantage of this is that a user can see, and interact with, the steps that have been invoked by progressively traversing the decision tree for choosing colours or gamut mapping schemes.

References

1. C. Beshers, (1992) "Automated Design of Virtual Worlds for Visualising Multivariate Relations", *IEEE Visualization'92*, proc., Boston, MA, 283-290.
2. C. Beshers and S. Feiner (1994), "Automated Design of Data Visualizations", *Proc. ONR Workshop on Data Visualisation*, Darmstadt, 1993, in *Scientific Visualization: Advances and Challenges*, 87-102, Academic.
3. S. Casner, (1992) "A Task-Analytic Approach to the Automated Design of Graphic Presentations" *ACM Transactions on Graphics*, **10** (2), 111-151.
4. R. N. Haber and L. Wilkinson, (1982), "Perceptual Components of Computer Displays", *IEEE Computer Graphics and Applications*, pp.23-35.
5. J. Mackinlay (1986) "Automating the Design of Graphical Presentations of Relational Information", *ACM Transactions on Graphics*, **5** (2), 110-141.
6. G. W. Meyer and D. P. Greenberg, (1980), "Perceptual Color Spaces for Computer Graphics", *ACM Computer Graphics* (Siggraph), Vol. **14**, pp.254-261.
7. P. Robertson, (1991) "A Methodology for Choosing Data Representations", *IEEE Computer Graphics and Applications*, **11** (3), 56-67.
8. P. K. Robertson, M. Hutchins, D. R. Stevenson, S. Barrass, C. Gunn and D. Smith, (1994), "Mapping Data into Colour Gamuts—Using Interaction to Increase Usability and Reduce Complexity", Special Issue on Advanced Interaction (Felger & Boem, eds), *Computer & Graphics,* Vol. **18**, No. 5, pp.653-665, Sept '94
9. P. K. Robertson and M. Hutchins, (1995), Guided Color Representation Using Perceptual Color Spaces, *SPIE Symp on Electronic Imaging Science and Technology: Human Vision, Visual Processing and Digital Display VI*, Vol. **2411**, 44-50, San Jose, 5-10 Feb.
10. B. E. Rogowitz and L. A. Treinish, (1993) "Data Structures and Perceptual Structures", SPIE Symp. *Human Vision, Visual Processing and Digital Display IV*, proc., **1913**.
11. B. E. Rogowitz and L. A. Treinish, (1993) "An Architecture for Rule-Based Visualization", *Proc. of IEEE Visualization* '93, San Jose, 236-244.
12. B. E. Rogowitz, D. T. Ling and W. A. Kellog, (1992), "Task Dependence, Veridicality and Pre-Attentive Vision: Taking Advantage of Perceptually Rich Computer Environments", Proc. SPIE **1666**, *Human Vision, Visual Processing, and Digital Display III*.
13. S. Roth and J. Mattis, (1990) "Data Characterization for Intelligent Graphics Presentation" *ACM-CHI'90*. proc.
14. H. Senay and E. Ignatius, (1991) "Compositional Analysis and Synthesis of Scientific Data Visualization Techniques", in N. M. Patrikalakis (ed) *Scientific Visualization of Physical Phenomena,* 269-281. Springer-Verlag.
15. H. Senay and E. Ignatius, (1994), "A knowledge based system for visualisation design", *IEEE Computer Graphics and Applications*, Vol. **14** No. 6, 36-47.
16. M. C. Stone, W. C. Cowan and J. C. Beatty, (1987), "Colour Gamut Mapping and the Printing of Digital Colour Images," in *Colour Perception Tutorial Notes*, *SIGCHI and GI* 87, W. B. Cowan et al., eds., ACM, New York, pp. 183-220.