

Cognitive Benefits of a Simple Visual Metrics Architecture

John King^(✉), Kathy Sonderer, and Kevin Lynch

Raytheon Missile Systems, Tucson, AZ, USA

{john.king, Kathy_M_Sonderer, Kevin_J_Lynch}@raytheon.com

Abstract. Many organizations produce metrics dashboards that take a long time to develop, are visually inconsistent, require specialized staff and skills, and (of most concern) don't clearly and rapidly identify actions or draw focus for further analysis. We addressed these issues, and realized unanticipated benefits as well, by creating a strong design and development architecture. Our results include: a templated set of metric visualizations, a radical decrease in cycle time, and realizing "self-service" business intelligence capabilities, empowering business users with expert domain knowledge to own and develop metrics. In this presentation, we discuss the visual architecture and design, the small set of templates, and the cognitive benefits of the visualizations now in use. The approach has garnered success at the company, program, directorate, and department levels, in large part due to the low cognitive burden for visualization understanding and development.

Keywords: Cognition · Visual architecture · Metrics

1 Introduction and Motivation

It is increasingly popular to display sets of disparate metric information as individual tiles laid out in grid fashion on the screen [1]. Our approach utilizes a library of standard templates in two categories, key performance indicators (KPI) and dashboards [2, 3]. Our metric development instructions are themselves visual, allowing them to be constructed easily by novices [4]. Our most utilized template, and the core visual element of our design, is our "metric display tile." Each tile delivers a discrete and easily discernible quantum of status. We designed a tile that visually encodes a substantial amount of status and trend data while boldly highlighting the primary messages, providing an economy of information interaction [5]. The library contains additional templates that provide more detailed charting capability for further analysis. The system works together so that tiles can link to other metrics displays, comment threads, and metric documentation, in snap-together fashion [6]. Finally, the templates are built to receive data from any source as long as the data are structured compatibly. Business analysts, data developers, and dashboard builders have a clear target for their data-related work; they simply get the template, set parameters and produce a functional, highly usable visualization. Constructed templates are simple to modify, highly reusable, and quick to deploy.

Cost and elapsed time are significant considerations for metrics development inside the enterprise. Different representations for metrics pose a challenge for consistent

interpretation and intuitive action across departments, directorates and programs. A framework was needed in which metrics could be developed and maintained at low cost, to produce metrics that were readily understood. The framework needed to scale to a large number of metrics, and produce metrics that required little or no customization, with an extremely low cognitive burden. To accomplish this, a small set of reusable templates with strong user experience principles needed to be developed and supported. The templates had to accommodate a wide range of data sources, and be deployable by subject matter experts, not requiring technical experts. The framework had to support reusable templates that produced consistent, high-quality metrics visualizations, both to facilitate standardization and to keep cost low and elapsed development time short.

This paper is organized as follows. Section 2 presents the metrics architecture. Section 3 describes the composition of the visual metrics. Section 4 outlines business value. Section 5 describes the cognitive benefits of this approach. The conclusion is provided in Sect. 6.

2 Metrics Architecture

Layered software architecture is a very common design for client-server configurations [7]. A multi-layered metrics architecture facilitates the creation of the visualizations by encapsulating and segregating functionality. Each layer has a well-defined interface to communicate with the connecting layer (Fig. 1). The architecture has several advantages; (1) it enables the ability to make changes in any layer without affecting the other layers, as long as the communication interface remains unchanged, (2) each layer can be independently worked, (3) work can be completed in parallel or asynchronously and (4) different teams can work on different layers.




Fig. 1. Overview of the metrics architecture

2.1 Data Sources

The foundational data layer can be traditional data warehouse tables, Excel or text files (Fig. 2). Connectivity to Excel or text files provides the option for metric development prior to data automation or when automation is not an option.



Fig. 2. Data can be consumed from most common sources

2.2 Data Encapsulation

The data source is connected to a pre-defined template which generates the required number of data points to render a visualization (Fig. 3). A single web service for each metric visualization is produced. Any data calculations or transformations can be encapsulated within the pre-defined template or performed prior to connecting to the data source.



Fig. 3. Data are manipulated for delivery to visualization templates

2.3 Visualization Templates

The single web service is connected to the template via an external interface connector which provides the data to the template in the correct format with the correct number of data points. The template contains a map that aligns each data point to the appropriate visual component (Fig. 4).




Fig. 4. Visualization templates

2.4 Design & Deploy

The completed metrics are deployed in a highly configurable web-based framework. Supplied parameters drive individual page configurations. Pages can be configured by role (executive level, mid-level management, program lead), and by product lines or

locations (plant, warehouse, factory). The standard visualizations plug into the framework via a parametrized URL. The visualization can be arranged in any order and integrated with other types of objects such as PowerPoint presentations. The plug and play components make it easy to create, customize and change pages (Fig. 5).



Fig. 5. Standardized visualizations are arranged for final delivery

3 Visual Metrics

Metric information is presented as a grid or series of tiles. Each tile delivers a discrete, digestible, quantum of status information “at-a-glance,” while more focused attention reveals further detail and additional features. At-a-glance, each tile highlights 3 performance parameters:

- *Current Metric Health Status* - performance relative to established expectations, e.g., exceeds, meets, or fails.
- *Latest Performance Delta* - performance improved or degraded over the latest measurement period.
- *Performance Trending*. The changes in performance over a series of measurement periods.

Additional details on the tile include the *Current Numeric Metric Value* and the *Maximum and Minimum Values* recorded over the entire measurement period. Moreover, users can interact with tiles by clicking or tapping to reveal additional features, including a *Link to Further Analysis*, a *Link to Threaded Notes* so that users can contribute comments regarding performance, and a *Link to Metric Definition Documentation* for additional information about how the metric is calculated.

We can present a great deal of information without confusing users because each performance parameter is communicated via a distinct visual display parameter. These include hue, intensity, and form (Fig. 6).

- Hue: *Current Metric Health Status* is indicated by one of three contrasting background-colors.
- Intensity: *Latest Performance Delta* is indicated by the background-color intensity and saturation.
- Form: *Performance Trending* is shown as a sparkline that is colored to indicate the health of the overall trend.




Fig. 6. Each tile communicates high-level status information through visual properties: color (hue and intensity) and form. Further detail is available by reading text values and by interacting with the tile. (Color figure online)

While form and color are powerful properties for preattentive detection, conjunction objects are not perceived preattentively [8]. In our implementation, we separate these properties. When viewing a set of tiles, any red background-color indicates an unhealthy status, while light- or dark-red indicates that the performance over that last measurement period improved or got worse respectively. Meanwhile, a green sparkline would indicate improvement over time, even in a metric that is currently unhealthy.

Because each mode is free of noise, critical information can be identified among a large set of tiles (see e.g., Fig. 7), exploiting cognitive expectation and attention [9]. Background colors are prominent and “chart junk” is minimized. Additional information (e.g., min and max values) contextualizes the metric without diluting the signal.

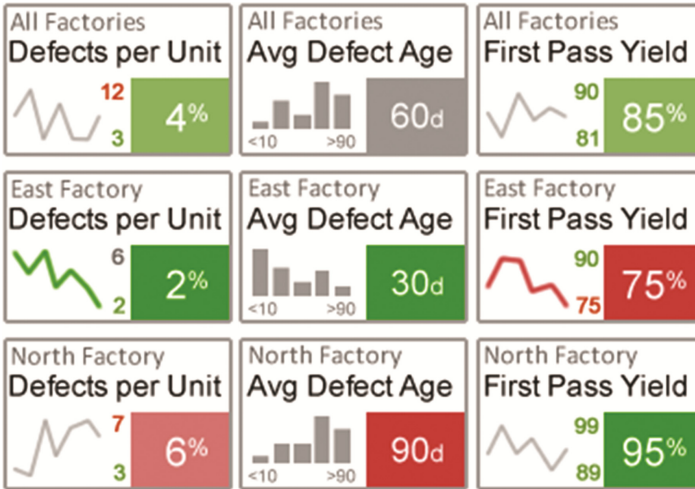


Fig. 7. A set of tiles grouped into a dashboard (Color figure online)

4 Business Value

During a 12-month period, we implemented our metric display architecture on three major projects. We compared these projects to equivalent work done using the old development process. Our results are as follows.

4.1 Cycle Time

Cycle times were reduced by 50–70 % for the equivalent work as compared to our previous development process. The overall architectural approach of having more atomized visualizations saves time as well since we are making files that are more simple and straightforward rather than building more complex displays into a single file.

4.2 Productivity

- The “fill out the form” nature of the development expands the number and types of staff that can produce metrics. File development can be accomplished by non-Developers, e.g., Business Analysts can set parameters and create a production ready visualization. This can also be extended to represent a Self Service BI capability [9].
- Rework involved in the development process was reduced for the equivalent work as compared to our previous development process.
- The templates are built to receive data from any source as long as the data are structured compatibly. This means that Business Analysts and Data Developers and Dashboard builders have a clear target for their data-related work.

- Having a standard set of visualizations streamlines the process of defining and refining customer requirements by allowing us to rapidly recycle through mockups and prototypes with customer input.

4.3 Quality

- The common visual vocabulary can be used by every organizational unit and functional area from Finance to Quality whether the intent is strategic or operational.
- Overall quality and consistency are improved.

4.4 Sustainability

- The visualizations are easier to maintain because they are simpler and the standardization makes the changes easier to locate. The changes that have been experienced so far have been in the data encapsulation layer.

5 Cognitive Benefits

The cognitive benefits of information visualizations are well-known [11, 12]. Visual imagery can have an important role in cognitive tasks [13, 14]. In this section we identify specific benefits for metrics display using the tile format based on our current work.

5.1 Scanning Dashboards or Scorecards Is Cognitively Costly


Tiles that are arranged into dashboards are useful for identifying metrics that may require intervention. A typical dashboard would lay out a set of tiles in positions specified during development, or possibly positioned by an end-user at run-time. The metrics consumer is required to scan the dashboard for metrics that may require further attention. Alternatively the user may formulate an *a priori* condition for which to scan, e.g., metrics that are out of control, or perhaps metrics that are performing exceptionally well.

These activities incur a cognitive burden for the user. Activities like formulating an *a priori* query or scanning the display for conditions that meet the query, or scanning the display for exceptions requires cognition and is subject to a host of external contextual demands and preexisting biases. There has been directed interest in identifying visualization techniques that reduce this cognitive burden [15, 16].

5.2 Sorting and Ordering the Display Structures and Prioritizes Attention

If a display is populated by tiles that are arranged dynamically at run-time, tiles can be arranged to highlight a higher-order level of insight. For example, in Fig. 10 the dashboard is split in half. Tiles representing Healthy metrics are one side while tiles for Unhealthy metrics are on the other (see Fig. 8 for details of how categories are assigned).

Evaluating Metrics into Healthy and Unhealthy Categories



Midpoint of the middle range divides Healthy and Unhealthy

Fig. 8. Assuming three performance categories, *healthy* metrics are those with values on the desirable side of the midpoint of the middle range.

Healthy and *Unhealthy* metrics are further divided into those that improved over the latest measurement period and those that did not improve. This produces four cases. Each case is presented as a column in a grid layout.

- Unhealthy metrics that got even worse.
- Unhealthy metrics that improved.
- Healthy metrics that got worse.
- Healthy metrics that improved.

Within each column, metric tiles are sorted in decreasing order of the *Latest Performance Delta* (i.e., the normalized magnitude of change over the latest measurement period). See Fig. 9 for an illustration.

5.3 Performance Signals Parsed into Distinct Channels

One implication of arranging tiles in this way (Figs. 9 and 10) is that four distinct performance signals are resolved with greater clarity as compared to statically positioned tiles. These performance signals are: (1) failing performance, (2) poor performance that may be turning around, (3) good performance that may be slipping, and (4) good performance that is getting even better, a.k.a., superstars.

5.4 Dynamically Sorted Dashboard Example

Figure 10 is an example of how metrics might be arranged to focus on metric improvement.

- The metric at position (a) is unhealthy and has gotten worse by the largest margin among all the unhealthy metrics that have gotten worse over the last reporting period. In this case, Alpha Factory has experienced an alarming spike in injuries.
- The metric at position (b) is unhealthy, but has improved by the largest margin among all the unhealthy metrics that have improved. In this case, Echo Factory has improved their ability to meet their schedule and they have made it into the middle (grey) range.

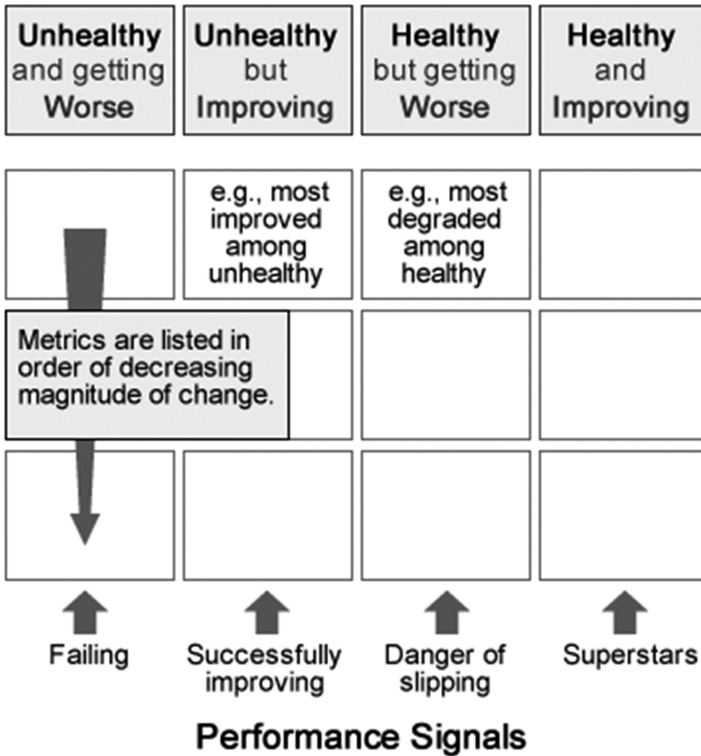


Fig. 9. Tiles are arranged dynamically within a structure that focuses on metric improvement. Tiles are divided into healthy and unhealthy metrics that have either improved or declined (i.e., gotten worse). Within each column, metrics are sorted in decreasing order of the *Latest Performance Delta* (i.e., the normalized magnitude of change over the latest measurement period).

- The metric at position (c) is healthy. But over the latest reporting period it has gotten worse by the largest margin among all the healthy metrics that have gotten worse over the same period. In this case, Alpha Factory has experienced a substantial slip in their ability to meet their schedule.

6 Conclusion

We present a simple visual metrics architecture composed of reducible, reusable, and serializable elements. These elements represent an inclusive approach that is concerned with source data at one end, and extends to user perception at the other end. Our data visualization architecture is Data are provided to the visualization templates in a platform agnostic way. These processed data are then provided to the visualization part of the architecture. Visualization elements are arranged in a purposeful way to facilitate orthogonal exploration of data by directing attention. That is, a viewer can scan among elements on the surface for high-level information but can also drill into further detail

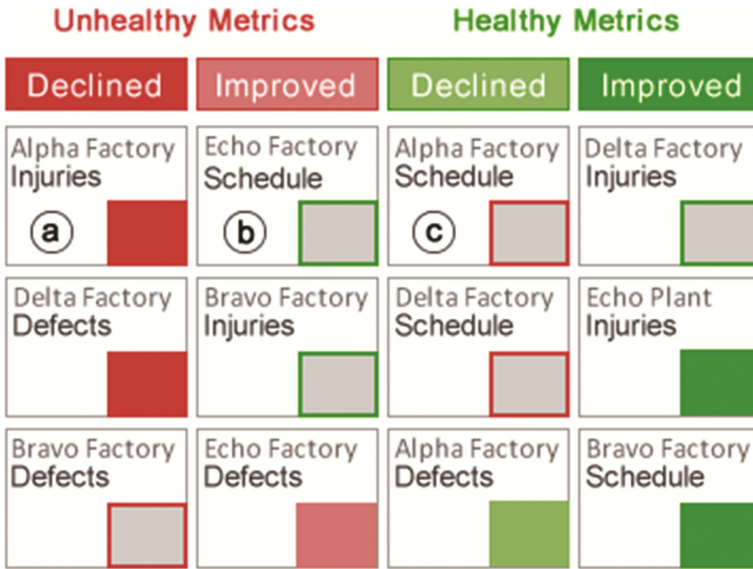


Fig. 10. Example of a dashboard with dynamically sorted tiles. We have omitted some details (sparklines and text) in order to illustrate this dashboard.

by focusing attention on a single element. These two levels of attention may be referred to as “orientation” and “engagement” [8]. Our visualization architecture is effective because any part of the visualization can “get out of the way” to allow the user to efficiently transition between orientation and engagement, and to transition among elements (e.g. among tiles or among elements within a tile).

References

1. Lamantia, J.: Enhancing the portal experience. *Enhancing Enterp. Serv.-Oriented Archit. Adv. Web Portal Technol.* 245 (2012)
2. Parmenter, D.: *Key Performance Indicators: Developing, Implementing, and Using Winning KPIs.* Wiley, Hoboken (2015)
3. Palpanas, T., Chowdhary, P., Mihaila, G., Pinel, F.: Integrated model-driven dashboard development. *Inf. Syst. Front.* 9(2–3), 195–208 (2007)
4. Grammel, L., Tory, M., Storey, M.A.: How information visualization novices construct visualizations. *IEEE Trans. Vis. Comput. Graph.* 16(6), 943–952 (2010)
5. Tudoreanu, M.E.: Designing effective program visualization tools for reducing user’s cognitive effort. In: *Proceedings of the 2003 ACM Symposium on Software Visualization*, June 2003, pp. 105–ff. ACM (2003)
6. North, C., Shneiderman, B.: Snap-together visualization: can users construct and operate coordinated visualizations? *Int. J. Hum.-Comput. Stud.* 53(5), 715–739 (2000)
7. Sharma, V.S., Jalote, P., Trivedi, K.S.: Evaluating performance attributes of layered software architecture. In: Heineman, G.T., Crnković, I., Schmidt, H.W., Stafford, J.A., Ren, X.-M., Wallnau, K. (eds.) *CBSE 2005. LNCS*, vol. 3489, pp. 66–81. Springer, Heidelberg (2005)

8. Healey, C.G., Enns, J.T.: Attention and visual memory in visualization and computer graphics. *IEEE Trans. Vis. Comput. Graph.* **18**(7), 1170–1188 (2012)
9. Sonderer, K., Lynch, K.: Metrics visualization templates case study in aerospace product development. In: International Symposium on Electronic Imaging Visualization and Data Analysis Conference, 16–18 February 2016, San Francisco, California (2016)
10. Summerfield, C., Egnér, T.: Expectation (and attention) in visual cognition. *Trends Cogn. Sci.* **13**(9), 403–409 (2009)
11. Fekete, J.-D., van Wijk, J.J., Stasko, J.T., North, C.: The value of information visualization. In: Kerren, A., Stasko, J.T., Fekete, J.-D., North, C. (eds.) *Information Visualization*. LNCS, vol. 4950, pp. 1–18. Springer, Heidelberg (2008)
12. Ware, C.: *Information Visualization: Perception for Design*. Elsevier, Amsterdam (2012)
13. Kosslyn, S.M., Thompson, W.L., Ganis, G.: *The Case for Mental Imagery*. Oxford University Press, New York (2006)
14. Chen, C.: Top 10 unsolved information visualization problems. *IEEE Comput. Graph. Appl.* **25**(4), 12–16 (2006)
15. Melcher, J., Seese, D.: Visualization and clustering of business process collections based on process metric values. In: 10th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing, SYNASC 2008, pp. 572–575 (2008)
16. Peng, W., Ward, M.O., Rundensteiner, E.: Clutter reduction in multi-dimensional data visualization using dimension reordering. In: *IEEE Symposium on Information Visualization, INFOVIS 2004*, pp. 89–96 (2004)