

SHARE:

[Join Our Email List](#)

The KIT – Knowledge & Information Technology

No. 259 - 2 March 2020

Was this forwarded to you?



In This Issue

[Graph APIs vs. RESTful APIs](#)

[Labelbox and Mursion](#)

[CoViD-19 Data Visualization](#)

[Fairness in Machine Learning](#)

[Seen Recently](#)



Consulting Services

- IT Strategy
- Enterprise Architecture Roadmap
- Business Process Modeling & Analysis
- Enterprise Software Selection
- IT Innovation Briefings
- IT Due Diligence
- Executive IT Seminars
- Cloud Computing
- Security Maturity
- Software Process
- Knowledge Strategy
- Technical Communities
- Knowledge Capture
- Taxonomy development
- Enterprise Social Media

Contact Us:



→ **Graph APIs vs. RESTful APIs**

This is probably a bit more geeky than the concerns of our average readers, but we find it interesting that the representation of information as graphs (a set of connected nodes and edges) rather than as databases (loosely speaking, a set of related spreadsheet-like tables, at least speaking of relational databases) is gaining increasing popularity. In an article at [zapier.com](#), [Introduction to Graph APIs](#), Brian Cooksey elaborates on the introduction of "graph APIs" and gives examples (including code fragments) from APIs provided by Facebook and GitHub.

If you are interested in knowledge graphs and are in the Washington, DC, area, please come to the Object Management Group's meet-and-greet session, [Knowledge Graphs and Ontologies](#), on Monday 23 March from 5 to 7 pm. OMG is monitoring the CoViD-19 situation, of course, and at this point expects the meeting to take place as planned, but please revisit the [event page](#) between now and March 23 for any updates.

→ **Two Interesting Companies**

Since the last issue (that is, in the last two weeks) we met two interesting companies, based in San Francisco, with offerings related to machine learning and virtual reality.

[Labelbox](#) provides a user interface for the labeling of images and videos, so that the manual identification of specific objects can be used to train an image recognition algorithm. For example, human users may sift through thousands of training images and "label" the areas of the images that contain cars; this will help train the software to recognize cars in new images. While the application to self-driving cars is obvious, there are also use cases in policing, agriculture, defense, and many more.

[Mursion](#) (a play on the word "immersion," as far as we can tell) is a virtual-reality based training system for soft skills ("essential workplace skills" is the phrase the company uses). For example, it places a manager in a VR discussion in which she has to tell an employee that his performance is not satisfactory. We assume that there is, in addition to VR, a combination of speech recognition, natural language processing, and AI-based reasoning. If true, this makes the offering quite a powerful combination of multiple advanced technologies.

→ **Data Visualization Revisited**

What makes data visualizations good or bad is a periodically visited subject, made popular by Yale Professor Edward Tufte's 1983 book, [The Visual Display of Quantitative Information](#). "Courtesy" of the CoVid-19 virus, we can now see new examples of good and a bad visual displays:

- The World Health Organization (WHO) [situation reports](#) mostly consist of two separate tables (China and rest of the world), no historical trend visualization except an "epidemic curve" that's hard to interpret, and a map of the world where the country name labels are illegible. This is presented as a static PDF produced once a day. Numbers are sometimes inconsistent across reports -- for example, the total number of cases by day D, plus the number of new cases on day D+1, may not equal the total number of cases by day D+1.

www.cebe-itkm.com
info@cebe-itkm.com
+1 415 870 ITKM
Twitter: @cbaudoin

[Archive:](#)
[Previous KIT Issues](#)

Forward this issue to colleagues and friends: use the "forward email" link below at left, rather than "Forward" in your email software, to preserve your privacy, give the recipient more options (their own unsubscribe link, etc.) and to give us better click-through data. Thanks!

- The [worldometers report](#) shows key numbers, followed by details that the user can click on to switch between numbers and charts. When mousing over the curves in the charts, you can see the numbers (of cases, deaths, etc.) on each day since the onset. The numbers add up: if a previous day's numbers were updated after the fact (probably the reason why WHO's reports don't compute), worldometers shows the corrected data for the previous day. And the whole site is updated every 10 minutes as reports come in from various countries.

Of course, this is not really surprising: WHO's specialty is health, and worldometers' specialty is statistics and information displays. But then, why does WHO consume a fraction of its resources to produce an inferior product instead of outsourcing this to an organization with better skills? There is a lesson not only in visual display in this, but also on outsourcing.

→ [Fairness in Machine Learning](#)

The [ACM TechTalk](#) on this subject given on February 26 by Tulsee Doshi, Product Lead at Google, was excellent. Ms. Doshi demonstrated how training data from social media can lead an ML algorithm to become biased based on race, ethnicity, gender, religion or sexual orientation. She discussed the various metrics for fairness, and how one can try to correct for bias in the training data. It was well explained, and the replay is worth watching.

→ [Seen Recently...](#)

"Step 1: Acknowledge the possibility that all or part of your workforce may need to work remotely."

-- [Cali Williams Yost](#), CEO and Founder, Flex+ Strategy Group
(retweeted by [Dion Hinchliffe](#))