

Immersive Analytics: Building Virtual Data Worlds for Collaborative Decision Support

Richard Hackathorn
Bolder Technology

Todd Margolis
Qlik

ABSTRACT

Immersive analytics is an emerging research area that blends analytical reasoning with immersive virtual space to enhance collaborative decision support. The intent of this position paper is to stimulate discussion and cooperation toward maturing immersive analytics. An open innovation community to build immersive data worlds has been established at ImmersiveAnalytics.com, to serve as a bridge and catalyst between academia and corporate communities. The paper outlines the objectives for analytical reasoning and immersive data spaces, followed by suggestions for the design and architecture of data worlds. Finally, current work for building data worlds is described.

Keywords: virtual reality, data analytics, decision support, visual analytics, data visualization, immersive spaces, collaborative decision making, machine learning, predictive analytics.

Index Terms: Information Systems: Decision Support Systems, Human-Centered Computing: Collaborative and Social Computing Systems and Tools, Human-Centered Computing: Visual Analytics, Human-Centered Computing: Visualization Theory, Concepts and Paradigms.

1. INTRODUCTION

Immersive analytics (IA) is an emerging research area that blends analytical reasoning with immersive virtual spaces to enhance collaborative decision support. IA is receiving attention from both the commercial and academic sectors. Hackathorn & Franks [1] in 2012 described its potential for transforming the integrated data warehouse into persistent, rich, interactive, social and engaging” virtual data world, to be used for “education, training, discovery and problem solving.” Chandler and others [2] have outlined the potential for “multi-sensory interfaces” and need for research on “higher-level usability and design issues in creating effective user interfaces for data analytics in immersive environments.” Also, scientific organizations have been pushing 3D data visualization to higher levels of sophistication, reaping scientific discoveries, such as the [Fold.It](#) serious game for protein folding. [3] Finally, the wide availability of consumer VR technology will enable business users to explore complex systems in amazing new ways.

1.1 Open Innovation Community

This position paper describes the objective and design possibilities of an open innovation community to build immersive data worlds, as facilitated by the [ImmersiveAnalytics.com](#) website. The goal of this IA community is to stimulate extensive discussion and prototyping via collaboration and sharing. To encourage wide spread adoption and commercial offerings, its intellectual property is openly shared using Creative Common with Attribute license on

*email: richardh@bolder.com, todd.margolis@qlik.com

its website and open-source licensing (MIT/BSD/Apache) on its GitHub. The innovation approach is to build, document, share, critique, and then repeat. For the coming years, the strategy flow is shown in Figure 1.

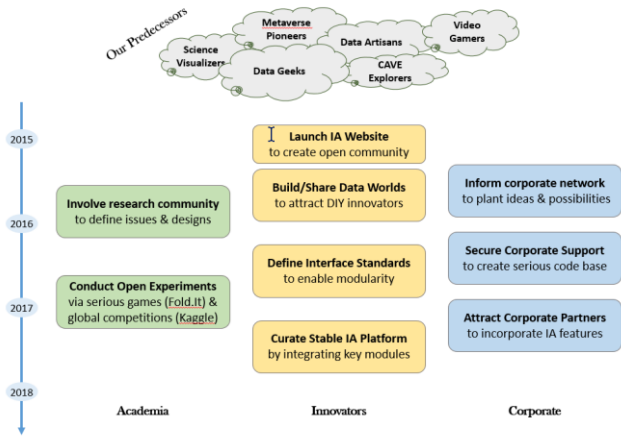


Figure 1- IA.com Strategy

1.2 Catalyst Between Academia and Corporate

An important aspect of this strategy is the role of the IA innovation community as a bridge and catalyst between the academic and corporate communities. For academia, the issues involved with IA are deep and significant, requiring quality academic peer-review research to guide long-term evolution. For corporate, the potential applications for IA are diverse and potentially game-changing, requiring corporate investments that result in stable commercial offerings. The intent is to focus interests toward building practical and well-designed applications.

1.3 Historical Perspective

The ideas surrounding immersive analytics have been fermenting for decades, as illustrated by Our Predecessors in Figure 1. The **Science Visualizers** have pushed the technology boundaries to generate amazing complex and valuable 3D visualizations. The **Data Geeks** have extended the scope and power of predictive analytics in countless directions. The **Metaverse Pioneers** demonstrated compelling examples of large scale virtual simulations shared by thousands. The **Data Artisans** expanded our awareness of the power of beauty having proper form and function. The **CAVE explorers** created large virtual spaces that pushed the limits of interactivity and tele-collaboration for large groups. The **Video Gamers** spawned a huge industry that drove virtual reality to the masses with cheap and reliable technology. Finally, an early IA effort can be found in a 1972 video of John Tukey demonstrating his PRIM-9 system for “discovering meaningful relationships hidden within 9 dimensions” via 3D visualization. [4]

2. OBJECTIVE

In the spirit of a data-driven organization, the objective of IA is to facilitate collaborative decision support for managing complex systems, while leveraging current analysis tools and architectures.

To achieve this objective, the following two essential components are blended.

2.1 Analytical Reasoning

Analytics are required to augment, not replace, human judgment, thereby enabling persons within a collaborative environment to generalize (infer) beyond known data and to identify situations where human judgment is required.

There are four levels of analytical reasoning. First, **descriptive** analytics summarizes and aggregates known data so that persons can sense the raw information being collected. Current visualization tools have greatly aided in this process. Second, **unsupervised** analytics clusters known data into categories and reduces the dimensionality of features, thus simplifying data for future analysis. Third, **supervised** analytics seeks an underlying model trained on known data and tests its predictive ability. Finally, **reinforced** analytics combines the previous levels into operational applications that improve through learning better models via successive train/test iterations. All of these levels play important roles in augmenting human judgment.

The good news is that machine learning (third and fourth levels above) is rapidly evolving, steadily enabling smarter applications, like autonomous vehicles. The bad news is that the complexity of machine learning is outstripping human ability to understand its methods and concepts, like sampling from petabyte datasets, separation of training and test data, tradeoffs of bias versus variance, tradeoffs of overfitting versus model complexity, cross-validation techniques, and optimization of model ensembles. [5]

To achieve a balanced marriage of human and computer, analytics should be *turned inside-out* implying that persons can sense key aspects of analytical processes as behaviors within the data world. For instance, the behavior of the gradient descent for a regression could be represented as a varying tone attached to the data object that predicts future sales.

Domingos argues that a master learning algorithm can be derived from a blending of the five tribes (approaches) to machine learning, resulting in “asymptotical growth to a perfect understanding of the real world.” [6] If true, ethical issues arise about how to turn such algorithms inside-out and thus enable humans to understand and govern their functioning.

2.2 Immersive Data Spaces

Virtual 3D spaces are highly malleable and extensible, with infinite possibilities. To support analytical reasoning, these spaces should be procedurally generated driven by the data and its analyses. The content of data spaces should be data objects, which in appearance and behavior represents various characteristics of the data.

There are four levels of data objects within immersive data spaces. First, the **data sculpture** is a static non-interacting 3D object whose appearance and behavior represents key characteristics of a dataset. [see No 3D Chartjunk] Second, the **data kiosk** is an interactive 3D object upon which a person can perform selection, filtering, highlighting, annotating and other operations, much like 2D coordinated multi-view visualizations. This is similar to interactive 2D visualizations. Third, the **data tour** is predefined sequence of data sculptures and kiosks that tells a story (problem, alternatives, solution). Finally, the **data world** brings the previous levels into a coherent whole. A data world integrates all the pertinent information (many datasets plus ongoing analytical processing) about a specific system into a single immersive space. In a corporate context, a data world is the human-facing virtual space reflecting the contents of an integrated data warehouse ecosystem.

The design of data worlds can benefit from the legacy of its physical counterparts. The work of Scott Lukas, an architect/designer, focuses on physical places where people have an immersive experience, such as amusement parks, resorts, shopping malls, and even cruise ships. He defines an immersive world as “a complete, diverse and consistent place, with a history and culture, inhabited by persons. It is ever-changing, evolving by relationships and forms of interconnections.” [7] These examples highlight the social aspects of immersion and should be thoughtful employed within IA designs.

Data spaces should invoke an immersive experience as a sense of *presence*, rather than the technology usage. Inexpensive and reliable virtual reality (VR) technology will be widely available driven by the video gaming and digital cinema industries for entertainment. However, we need use this technology to build immersive data worlds that enable us to experience analytics, in the spirit of the Flatland book by Edward Abbott [8]. For example, the point is not to witness raw matrix operations, but to sense the presence of those discerning eigenvectors that reveal simpler dimensions in a PCA.

A useful paradigm for thinking about the variety of immersive environments is the Virtuality Continuum (VC). [9] In Figure 2, VC spans from the fully physical world (on the left) to the fully virtual world (on the right).




Figure 2- Virtuality Continuum of Milgram & Kishino

IA designs can span most of this continuum, from the right with abstract 3D data objects toward the left with icons (or overlays) as anchors (or metaphors) for the real world. For instance, the “ground truth” plane could be configured as a schematic of a manufacturing plant so that the data spaces are located near the point of instrumentation in real-world processes. These compositional design alternatives can serve as mnemonic devices to aid participants in navigating complex data worlds. Additionally, augmented reality (with VR glasses) could superimpose selected parts of data world over the real-world machinery and sensors.

3. DESIGN POSSIBILITIES

To stimulate discussion, this section suggests various design possibilities for building immersive data worlds. Given available technology, the number and variety of design alternatives are immense, requiring creative experimentation and careful evaluation of their usability.

3.1 Mature Persistent Worlds

As data worlds mature, they will gradually become more persistent (always-on 24x7), like a botanical garden rather than a movie. A botanical garden is a good analogy because it is designed, curated and themed as a permanent facility, yet constantly changing. A botanical garden serves many purposes, from simple fun browsing to serious insightful study, from education to research. It is engaging, both emotionally and intellectually. It is information rich, offering many levels of resolution for the observer. It is interactive, by smelling the scents, feeling the breezes, and touching the leaves. It can be experienced solitary or with a collaborative team. Immersive data worlds should be similar.

3.2 The Physics of Data Worlds

A useful approach to data world design is to define the physics of data analytics to guide behaviors within an immersive data world.

In other words, what is the data-world analogy to gravity, momentum, or magnetic attraction/repulsion?

Information entropy (uncertainty or surprise in data) could be similar to gravity. The ground of the data world could represent the ‘ground truth’ of observational data, with below-the-ground as the unknowable complexity of reality. Near the ground would have high entropy since any new data would have high surprise value that lacks a model to generalize beyond known data. The ‘up’ dimension would be a measure of model prediction accuracy with less surprise and more certainty. This gravity analogy implies the growth of model trees above the ground truth. The taller the tree, the better the generalization. And, a learning algorithm would appear as a vigorous forest of trees, constantly striving upward.

Likewise, data worlds could have an attractive force between similar datasets. For example, parallel coordinate plots could have a force-graph dynamic that would move feature poles closer if highly correlated. Also, models that generate similar results would attract each other.

3.3 Sensing, Not Just Seeing

The term sensing is more appropriate for IA than visualizing for two reasons. First, the lesson from video gaming is that sound and haptic feedback is essential for an immersive experience. [10] Not only does it reinforce the visual experience, but it can add an ambient sense of what is happening in that locality of virtual space.

Second, there are deeper implications of sensing that involve the brain’s subconscious ability to detect patterns from new sensory signals. This is best illustrated by the work of David Eagleman with *Sensory Substitution*, which is a “non-invasive technique for circumventing the loss of one sense by feeding its information through another channel,” as was demonstrated in the TED talk *Can We Create New Senses for Humans*. [11]. By using brief sweeps of vibratory tactical stimulus embedded in a vest, hearing impaired persons were able to recognize spoken words. [12] By using similar visual or spatial audio stimuli within a data world (which is *Sensory Addition*), a person could sense subconsciously complex information, such as subtle harmony or discord in analytical processing.

A key concept is *umwelt* or perceived model of our physical environment by an organism (including humans), which is used to understand behavior based on differing sensory abilities across organisms. In effect, IA strives to expand the human *umwelt* by adding analytics as a set of new sensory inputs for understanding and managing complex systems.

3.4 No 3D Chartjunk

Immersive data worlds are not in competition with 2D descriptive data visualizations but are synergistic. 2D visualizations serve a critical role of providing drill-down information within the data world. Given current technology, techniques for managing 2D textures within a 3D space are quite robust.

For example, a data space generated from a simple dataframe could have a descriptive data object showing a box plot to describe the distribution of a feature. This data object could be 2D or 3D within the 3D space, depending on its usability. Hence, many data objects could be 2D within immersive data spaces, implying that IA can (and should) leverage current data visualization concepts and technology.

Research about usability of 2D data visualizations are relevant to building immersive data worlds. In particular, Munzner [13] recommends 3D visualization for natural 3D spatial formats like weather modeling. However, she cautions against “no unjustified 3D” where 3D visualizations have special problems with depth perception and occlusion. The take-away is that 3D data objects must provide unique capabilities and be synergistic to 2D

visualizations. If immersive data worlds are used solely to simply extend 2D visualizations by one dimension, effective use cases will be compromised.

3.5 Storytelling as Engaging and Collaborative

Effective storytelling should be engaging and collaborative for the listeners. [14] In most business intelligence, a misunderstanding about the implied audience (rather than the actual listener) leads to an over-simplified delivery of information at best or misleading interpretation at worst. IA has the opportunity to engage users with a high degree of participation, more like video gaming than watching a movie or reading a book. Sharing information via IA can be an active process that allows for transparent information and useful debate. Further, the infinite canvas of IA can eliminate the constraints of current dashboard paradigms, enabling scalable collaboration with thousands.

Promising work is highlighted by *Explorable Explanations*, a public-domain website by Nicky Case [15] and draws inspiration from the work of Bret Victor [16] and others. It is like storytelling via interactive infographics, but for active learning and knowledge construction, rather than for persuasion. The goal is to engage the user in exploring the behavior of a complex system, such as *Earth Primer* app by Chaim Gingold [17]. This work offers clues to extending IA designs for shared knowledge construction.

4. ARCHITECTURAL POSSIBILITIES

To stimulate discussion, this section suggests an architecture for building immersive data worlds, consisting of layers for design, render and display, as shown in Figure 3.

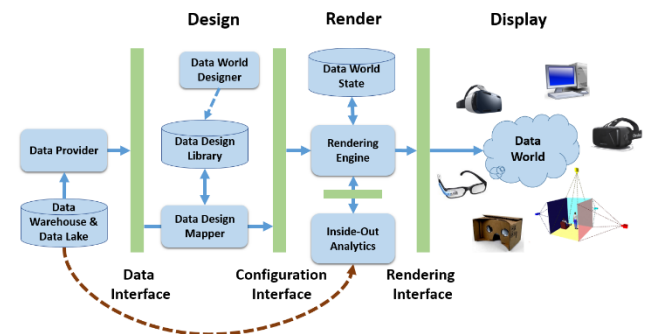


Figure 3 - IA Architecture

Data consists of a **Data Provider** (managed SQL-like access) and DW/DL combo. The data flow (in blue) is primarily metadata (attribute names, typing, counts) for multiple SQL views as tidy data (observations and variables in 3NF) used for the data world. Previously, the data world was designed with the **Data World Designer** that stores the configuration that maps each view to an in-world data object. The designer chooses from a library of various types of data objects to highlight aspects of the data. The view plus configuration flows to the **Rendering Engine**, which generates the 3D virtual space to be consumed by various display devices. The twist is the **Inside-Out Analytics**, which directly accesses the data stores (via curved arrow) and performs a specific analytic on a specific data view. Both the results and processing dynamics of the analytic are rendered in-world.

4.1 Elements

The elements for building data worlds could be as follows...

An IA support system consists of a catalog of one or more data worlds, called **dWorlds**. Each dWorld reflects a real-world complex system with a set of specific problems to investigate. For example, a dWorld could reflect all or a portion of an integrated data warehouse for a company.

Each dWorld consists of one or more data spaces, called **dSpaces**. Each dSpace has an associated external data frame (like an SQL view), called a **dFrame**. The dFrame is the ground-truth, represented by dSpace regions across landscape of dWorld.

Upon each dSpace, data objects, called **dObjects**, grow based on analytical generalizations, according to their level of analytical reasoning (descriptive, etc.).

Each dObject is instantiated (created) from a **dPrefab** in the dWorld design library. Each dPrefab is customized to render the results from a specific class of analytical algorithms. Initial work will examine the API design for the scikit-Learn project. [18] For each dObject with its dSpace, its API parameters are mapped to execution results via the dWorld design mapper.

4.2 Components

The components that could be used in this architecture are listed in Table 1.

Rendering Engines	Inside-Out Analytics	3D Displays
Gaming Engines, such as Unity3D, Unreal, CryEngine	Visual Analytics, such as Qlik, Tableau	Head-Mounted Display, such as Oculus Rift, Samsung GearVR, Google Glasses, Cardboard
HTML5 with WebGL, three.js	Analytic Platforms, such as Teradata Aster, Spark/DataBricks, AWS Analytics ML	Cave Rooms, such as CAVE2
Online Metaverse, such as High Fidelity, virBELA, OpenSimulator, Second Life	Analytic Flow IDE, such as RapidMiner, Alteryx, KMINE	Desktop with room projectors
Open Sandbox, such as MineCraft & Lego World	Custom R and Python Services	Input/Feedback Devices, such as joysticks, spatial sound, tactile gloves

Table 1- IA Architectural Components

5. CURRENT WORK

We are in the early stages of forming an open community by encouraging interested innovators with apt skills to contribute to IA projects. The authors are pursuing the following projects. Help!

Hackathorn is blending the technology from Python scikit-Learn with the gaming engine Unity3D. The scikit-learn community is an excellent beachhead with established bridges to other analytic tools (Jupyter notebooks, QlikView/Sense, RapidMiner), packages (NLP, Theano), platforms(Aster) and languages (R, Scala), along with scaling via cloud platforms (Spark MLlib). Further, the API structure of scikit-learn algorithms is open-source and accessible, [17] implying that data hooks could be placed within the algorithms to monitor intermediate stages of the analytic processing, like the behavior of gradient descent. The plan is to use scikit-learn as the template for exploring the combinations and flow of analytic methods. Examples rely upon Matplotlib, whose contents highlight key characteristics of analysis results. Latest Anaconda Python3 distribution enables interactive analytic development environment with Jupyter notebooks.

Likewise, the Unity3D community is an excellent beachhead with bridges into a wide range of 3D visualization and virtual reality technologies and communities, driven by a multi-billion-dollar gaming industry. Unity3D has become the preferred development environment for high-end games, having thousands of options for spawning, rendering and animating game-objects of any structure and shape. The Unity scene for a data world is procedurally generated (as opposed to manually built). C# is the implementation language, along with rapid prototyping tools.

Margolis is working on building 3D data objects and 3D data worlds powered by the Qlik Sense visual analytics platform rendered through Unity for Oculus Rift. The powerful Sense APIs enable connectivity to a wide range of data sources, advanced analytical capabilities through the Qlik associative engine and an infinite potential for visualization techniques. Qlik Sense has comprehensive Javascript APIs that allow developers to

procedurally create and interact with data objects within Unity and WebGL. The goal of this early work is to prototype 3D data objects that inherently provide users with new ways of understanding data that are unique to 3D visualization. We believe that certain classic 2D chart types such as Sankey, Sigma and Chord visualizations are more conducive to 3D or even 4D (adding a temporal dimension) representations. Also, the research is focused on taking an aesthetic approach harnessing natural human abilities for pattern seeking.

In conclusion, immersive analytics has the potential for realizing the ultimate user interface for understanding and managing complex systems. IA can be engaging, persistent, interactive, and collaborative, leveraging future VR and analytical technologies into integrated immersive data worlds. The challenges are not technology, but creative functional design that blends effective user interactions, collaborative decision support, and insightful learning algorithms. This is beautifully illustrated by Victor's talk "Media for Thinking the Unthinkable." [19] Success will be measured by the degree that human judgment is augmented, rather than replaced.

REFERENCES

- [1] Hackathorn, R.; Franks, B., "Create an Augmented Reality", Teradata Magazine, 2012. [link](#)
- [2] Chandler, T.; Cordeil, M.; Czauderna, T.; Dwyer, T.; Glowacki, J.; Goncu, C.; Klapperstueck, M.; Klein, K.; Marriott, K.; Schreiber, F.; Wilson, E., "Immersive Analytics," in Big Data Visual Analytics (BDVA), 2015, September 2015. [link](#)
- [3] S. Cooper, et al. "Predicting protein structures with multi-player online game," Nature, p. 756-760, Aug 5 2010. [link](#)
- [4] Hackathorn, R. "Beginning of Interactive 3D Data Visualization," ImmersiveAnalytics.com, Nov 18, 2015. [link](#)
- [5] Domingos, P. "A few useful things to know about machine learning." Comm of the ACM, 55.10, 2012, p78-87. [link](#)
- [6] Domingos, P., "The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World," Basic Books, 2015. [link](#)
- [7] Lukas, S., Immersive Worlds Handbook, Focal Press, 2012.
- [8] Abbott, E. Flatland: Romance of Many Dimensions, 1884 [link](#)
- [9] Milgram, P., Kishino, F. "A Taxonomy of Mixed Reality Visual Displays," IEICE Transactions on Information Systems, Vol E77-D, No.12, December 1994. [link](#)
- [10] GoBmann J., Hackbarth B., West R., Margolis T., Lewis J.P., & Mostafavi I. "Scalable Auditory Data Signatures for Discovery Oriented Browsing in Expressive Context," Proc. 14th International Conf on Auditory Display, Paris, 2008.
- [11] Eagleman, D. "Can We Create New Senses for Humans?" TED2015, March 2015. [link](#)
- [12] Novick, S.D., & Eagleman, D. "Using space and time to encode vibrotactile information: toward an estimate of the skin's achievable throughput." Experimental Brain Research, Oct 2015, v233, issue 10, pp 2777-2788. [link](#)
- [13] Munzner, T. Visualization Analysis and Design, A.K. Peters/CRC Press, December 2014, p.117-130.
- [14] Farmer, D. et al. "The Art of Persuasive Communication in the Workplace," Qlik Innovation and Design Research Digest, Winter 2015. [link](#)
- [15] Case, N. "Explorable Explanations" website. [link](#)
- [16] Victor, B. "Worry Dream" website. [link](#)
- [17] Gingold, C. "Earth: A Primer" website. [link](#)
- [18] Buitinck et al., "API design for machine learning software: experiences from the scikit-learn project," 2013. [link](#)
- [19] Victor, B. "Media for Thinking the Unthinkable," Vimeo, May 2013. [link](#)