Immersion Analytics

Simplify Complexity for AI Governance

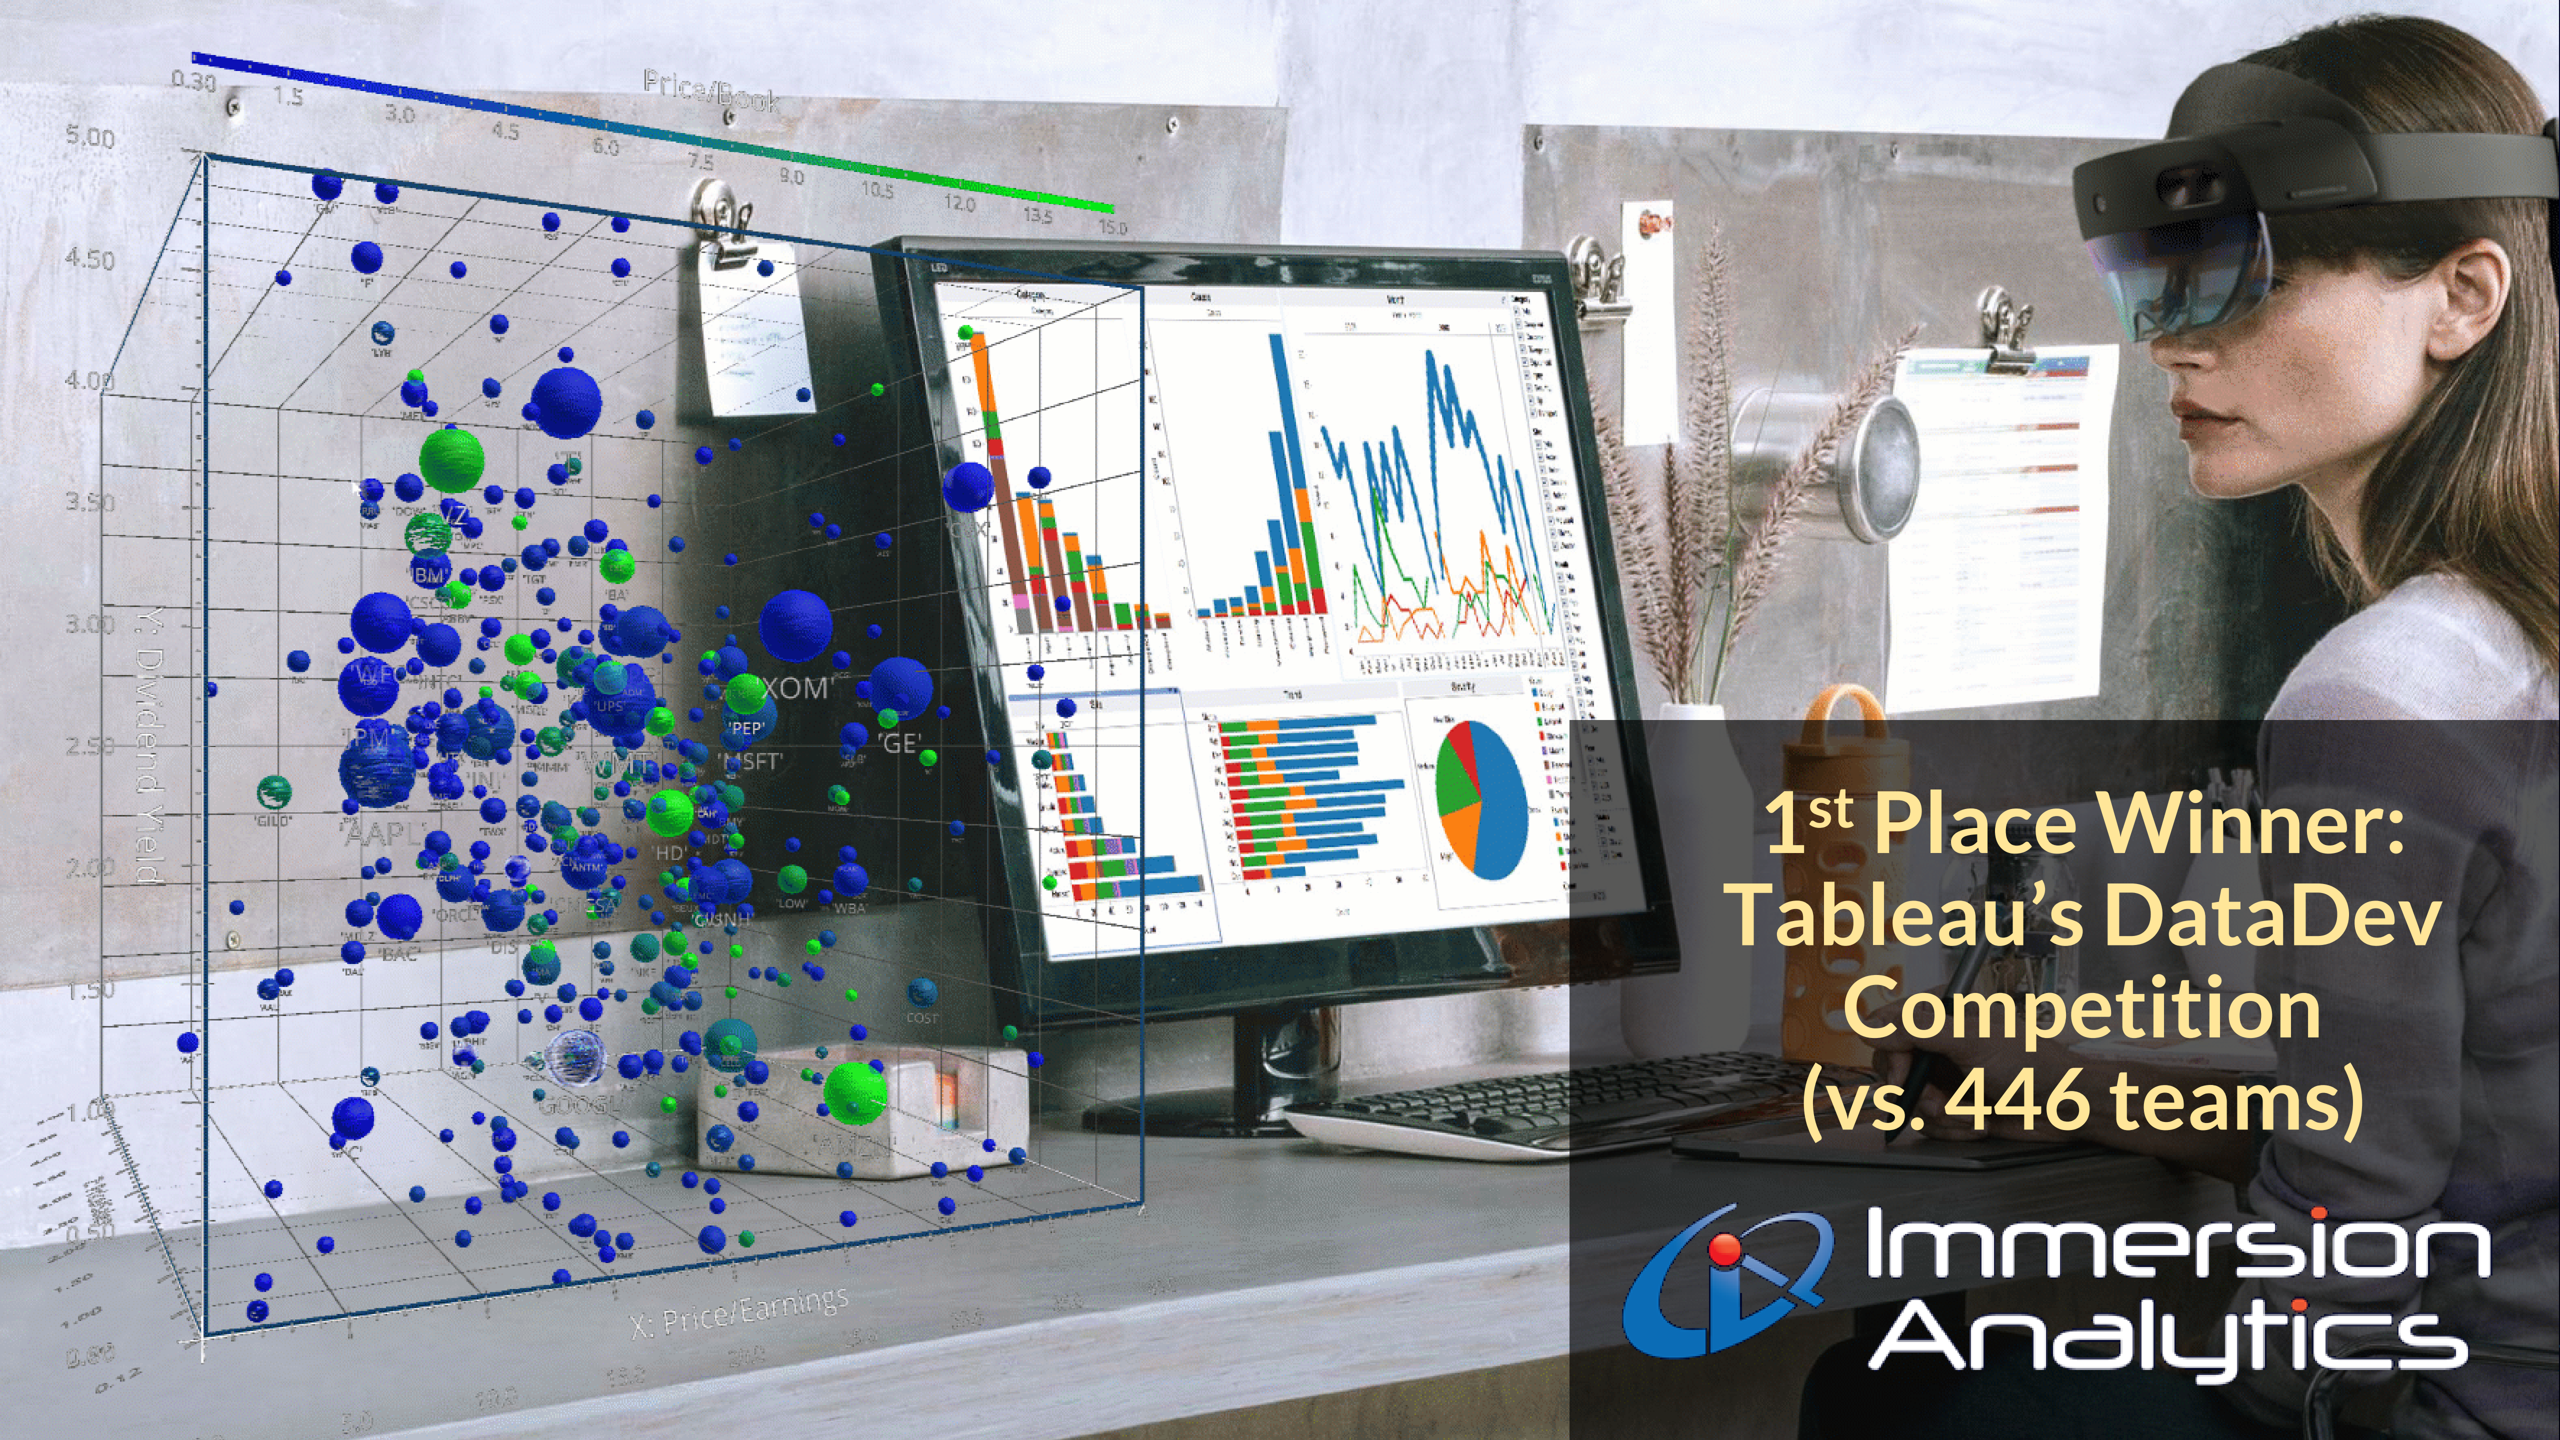

1st Place Winner: Tableau's DataDev Competition (vs. 446 teams)

Immersion Analytics

Immersion Analytics

Software Company focused on Explainable AI (XAI) including Large Language Models (LLMs)

Patented Immersive Visualization tech for Simplifying Data & Models

Web-embeddable Component

IA Runtime for Enterprise Integrations

Optionally, use Spatial Computing

Developer

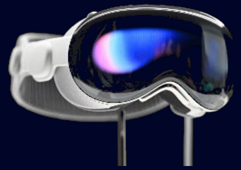

Microsoft

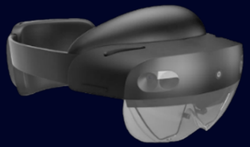

Meta

- 1st place winner, Tableau DataDev Hackathon (competing vs. 446 teams)
- 1st place winner, MIT Reality Virtually Hackathon (productivity category)
- Featured at MIT Technology Review EmTech Caribbean

mit media lab

Patent Issued

Dimensional Engine™

- *See up to 18D without distortion inherent in PCA or t-SNE.*
- *Achieved by layering special effects onto each data point, intensity of each encoding an additional numeric attribute.*

Stepwise Storytelling™

- *Simplifies by incrementally layering axes/effects one at a time.*

PDF blocks animations,
check out www.immersionanalytics.com

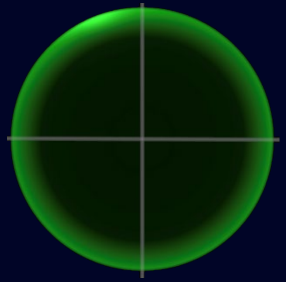

Opacity

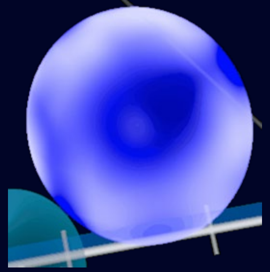

Shimmer

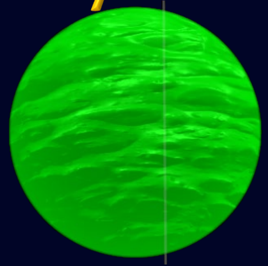

Bumpiness

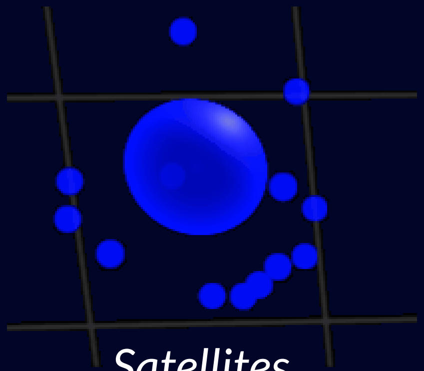

Satellites

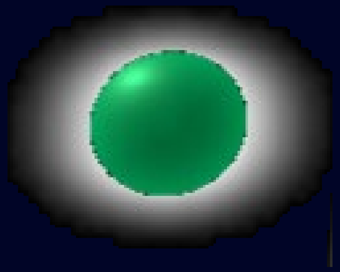

Glow

Pulsation

18 Axes include: X, Y, Z, Height, Width, Angles, Color, Shape, Opacity, Glow, Vectors, Bumpiness, Metallicity, Animated e.g., Shimmer, Satellites, Vibration, Pulsation, and sub-properties thereof

Patent Issued

Problem with AI (including LLMs)

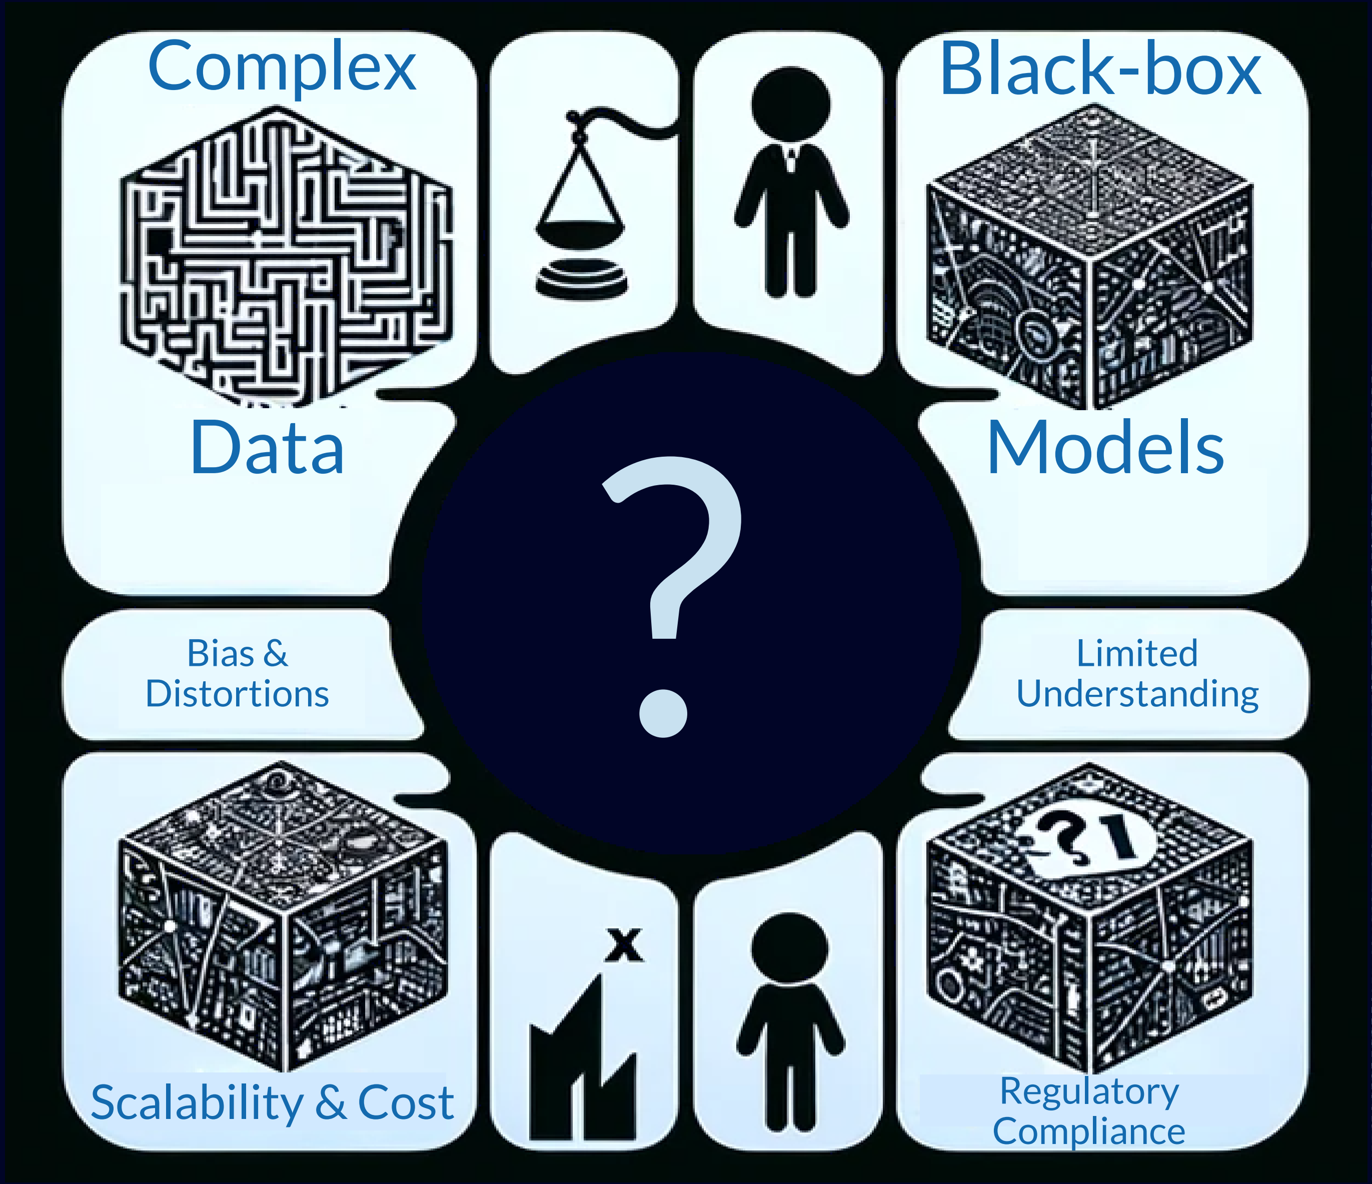

Complex Data

Black-box Models

?

Bias & Distortions

Limited Understanding

Scalability & Cost

Regulatory Compliance

Humans Need to Keep Up

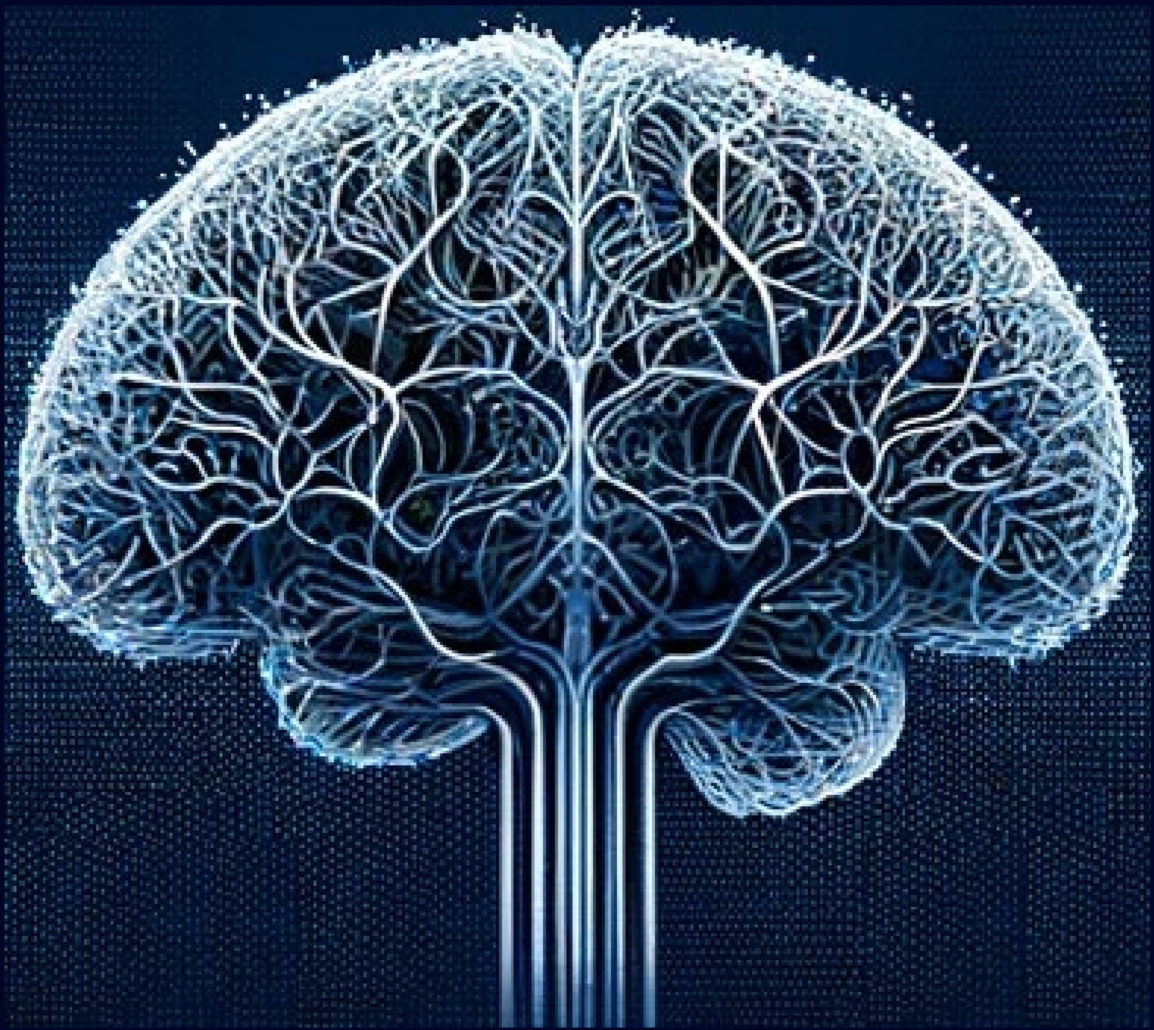

Solution

**PDF blocks animations,
check out www.immersionanalytics.com**

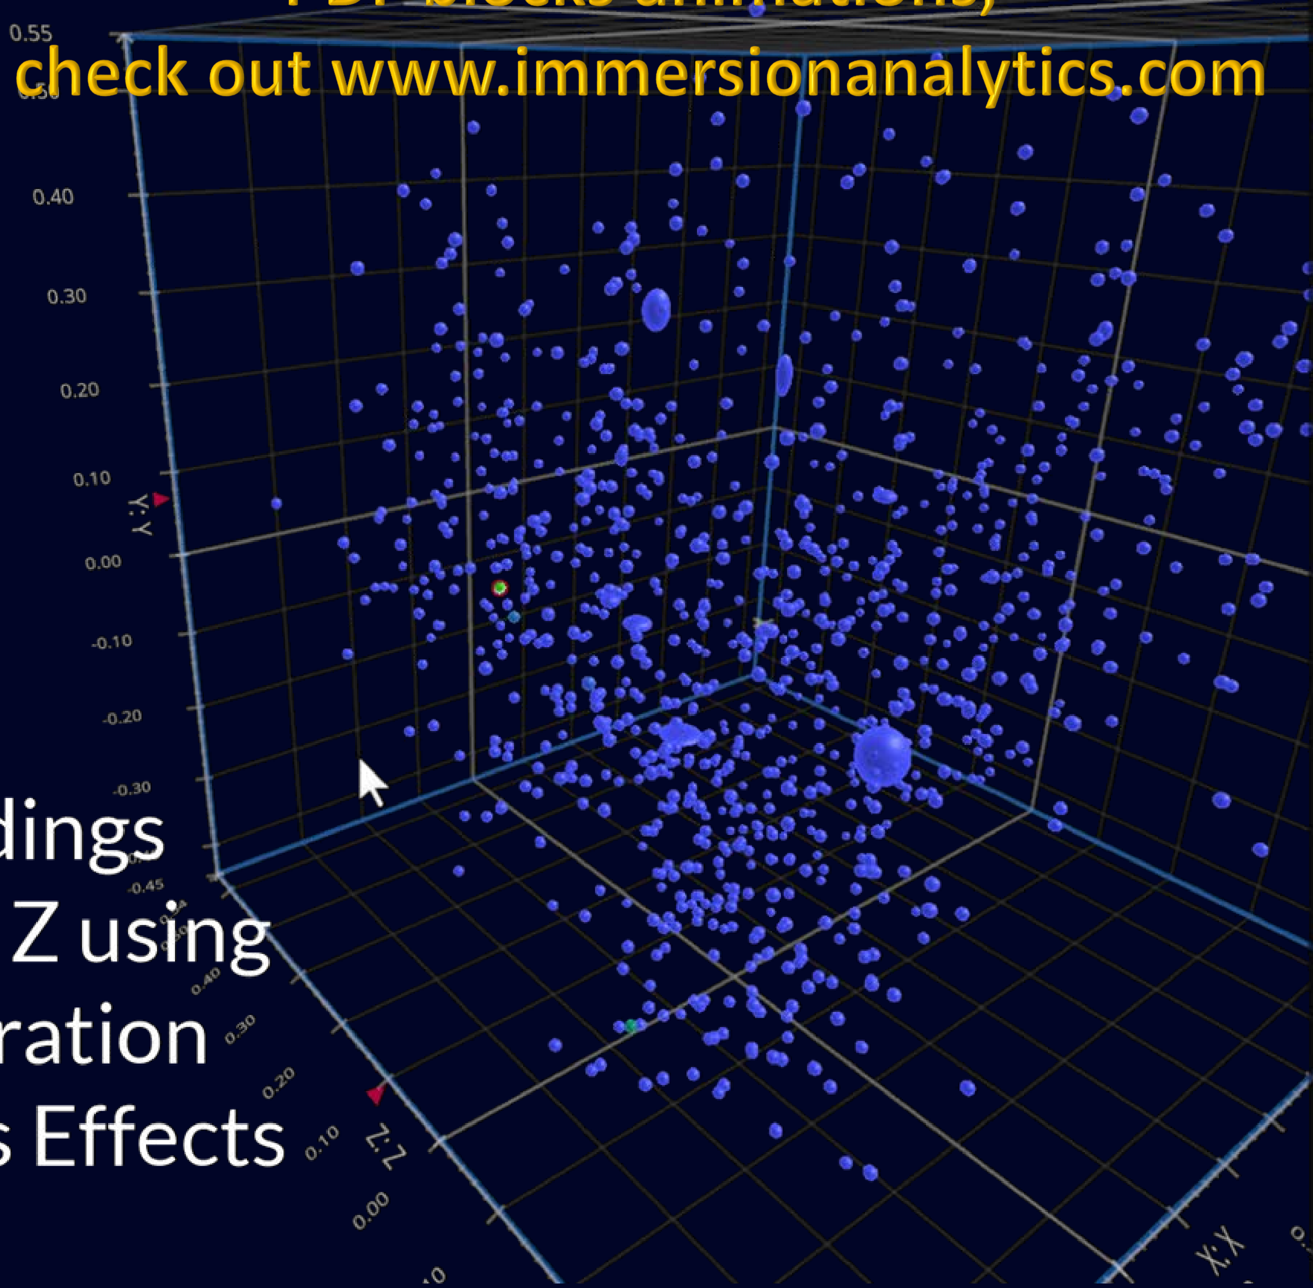

Vector Embeddings
Reduced to X, Y & Z using
PCA, plus Moderation
Factors Layered as Effects

Upgrading AI Safety & Governance

For: LLMs

Layer Moderation API Factors:

- Overlay vector embeddings from PCA on X, Y, and Z axes
- Then, layer moderation API factors e.g. odds response contains hate speech, profanity, self-harm, violence etc.
- This holistic view reveals underlying patterns in LLM behavior, including potential biases and areas requiring ethical guardrails.

For: Both

Hyperparameter Tuning: Visualize a random sample of hyperparameter combinations in the context of error rates to optimize model performance.

Exploratory Data Analysis: See complex data patterns to inform model design.

Data Quality: Notice subtle, context-specific data issues via domain expertise at scale.

Realtime Monitoring

Education

For: Machine Learning

Model Selection, Training & Debugging: See model behavior across all data points & epochs

Ensembling: Seeing above to assemble more effective model combinations.

Explainable AI (XAI):

- SHAP: see how features interact & influence predictions
- Better understand LIME
- Compare Explanations
- Complex Relationships: See beyond Partial Dependence Plots (PDPs).

Patent Issued

Upgrading AI Safety & Governance

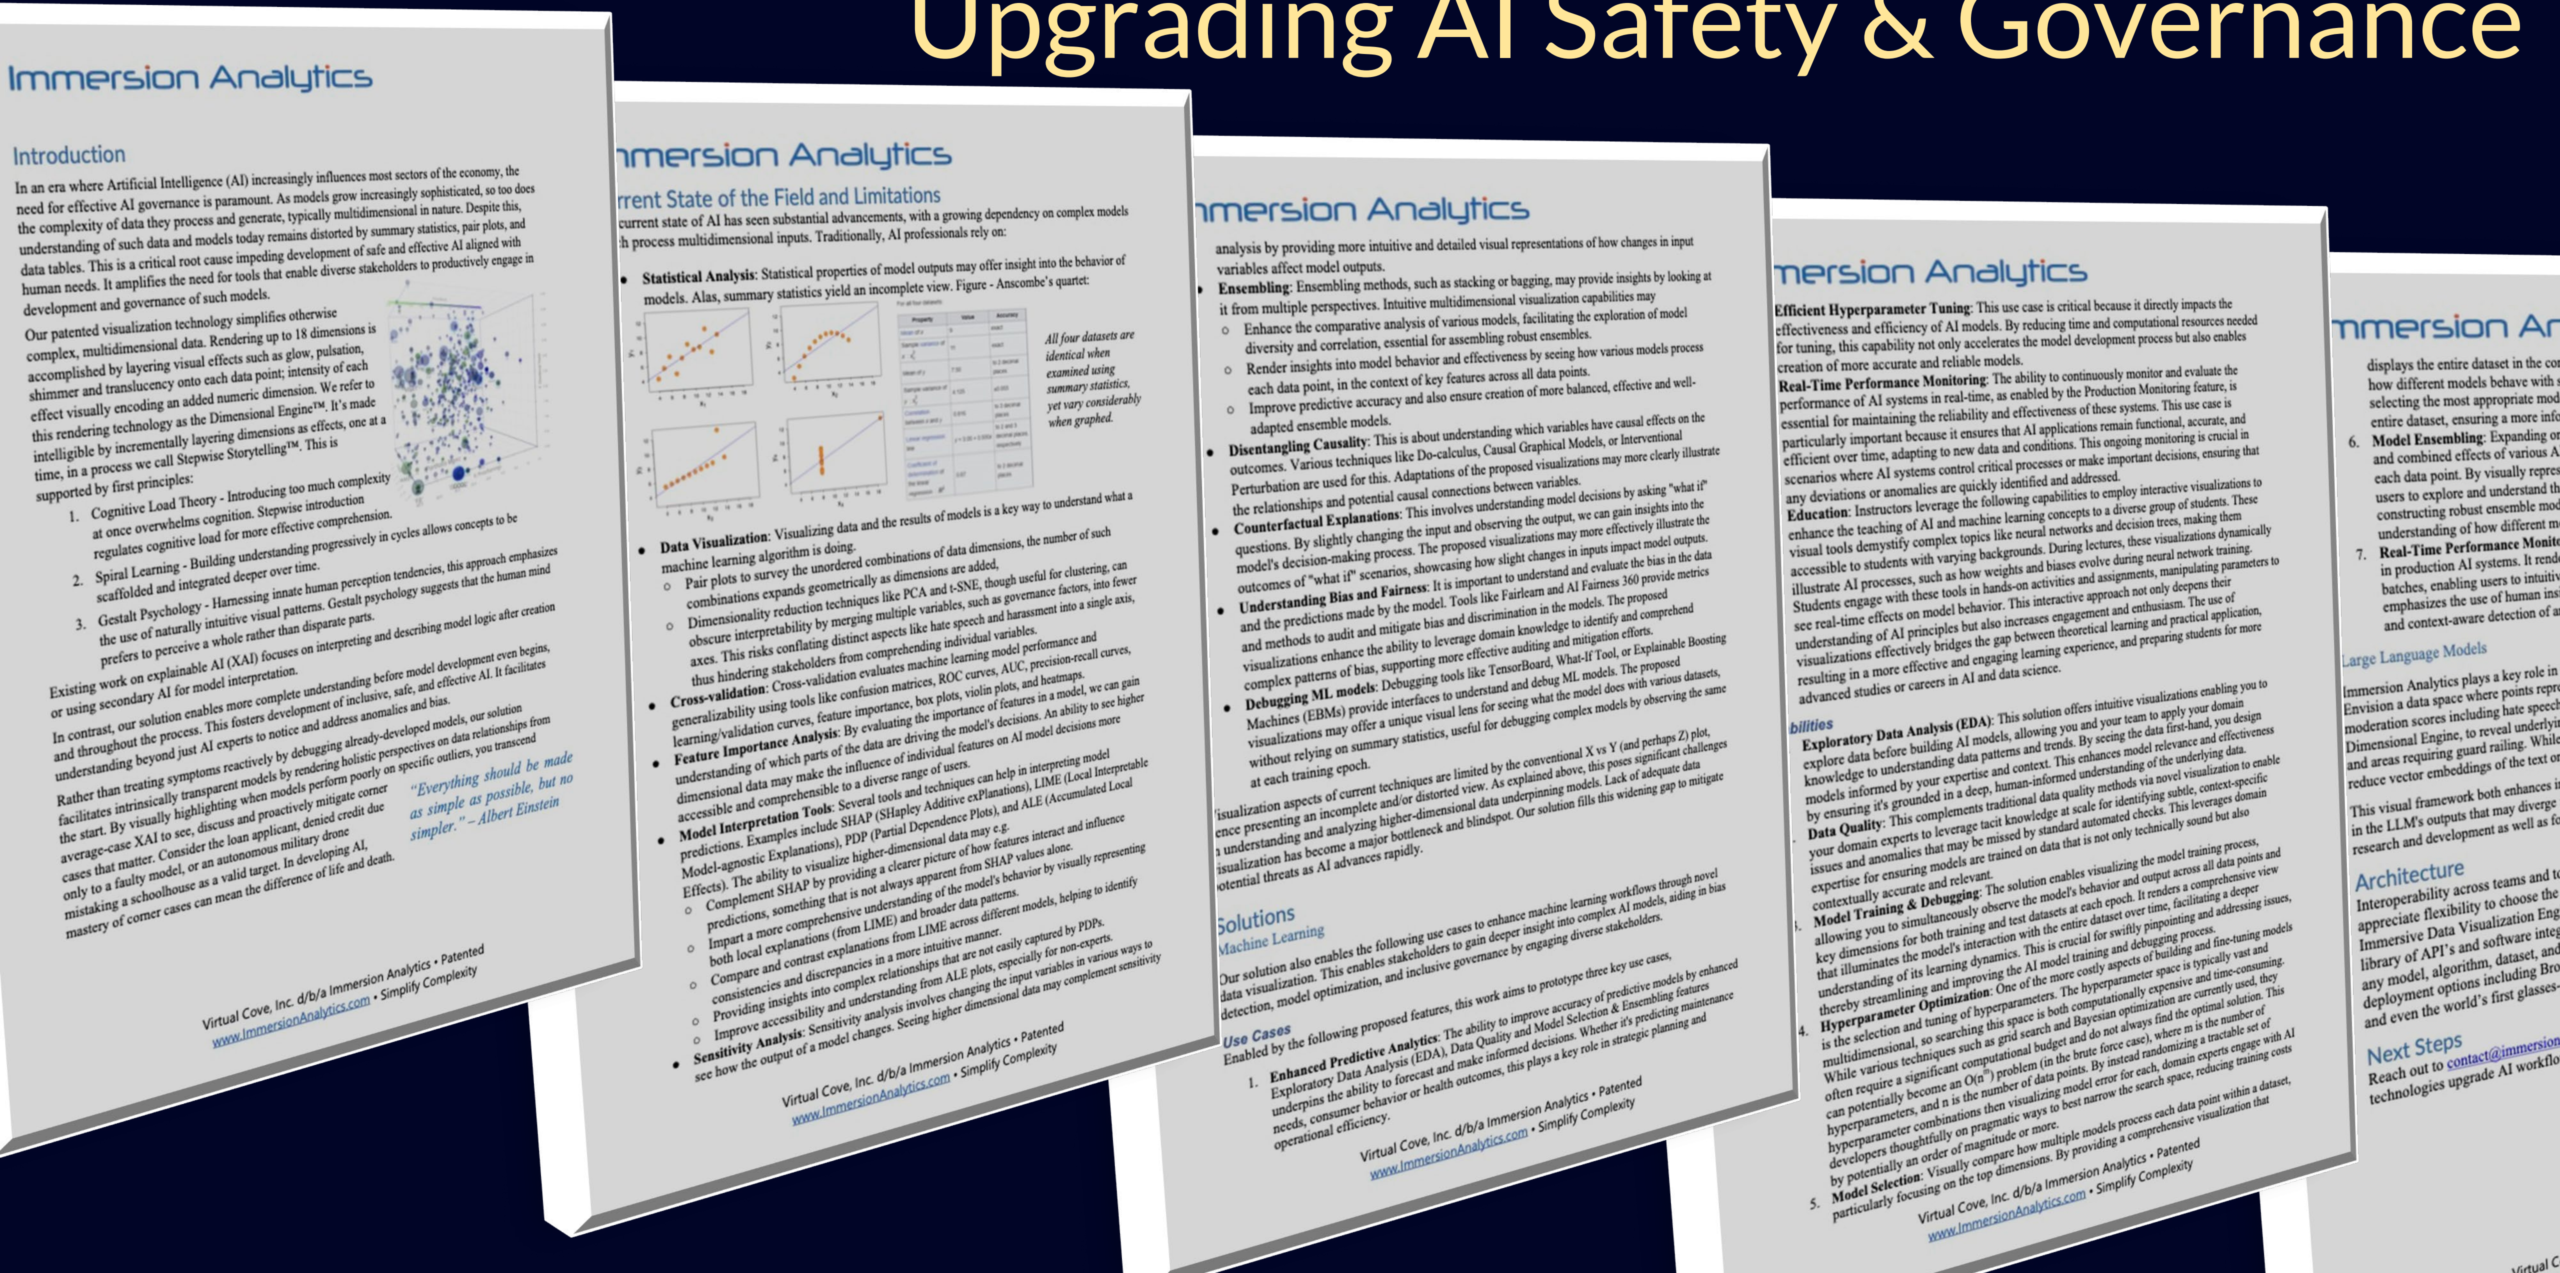

Immersion Analytics

Introduction

In an era where Artificial Intelligence (AI) increasingly influences most sectors of the economy, the need for effective AI governance is paramount. As models grow increasingly sophisticated, so too does the complexity of data they process and generate, typically multidimensional in nature. Despite this, understanding of such data and models today remains distorted by summary statistics, pair plots, and data tables. This is a critical root cause impeding development of safe and effective AI aligned with human needs. It amplifies the need for tools that enable diverse stakeholders to productively engage in development and governance of such models.

Our patented visualization technology simplifies otherwise complex, multidimensional data. Rendering up to 18 dimensions is accomplished by layering visual effects such as glow, pulsation, shimmer and translucency onto each data point; intensity of each effect visually encoding an added numeric dimension. We refer to this rendering technology as the Dimensional Engine™. It's made intelligible by incrementally layering dimensions as effects, one at a time, in a process we call Stepwise Storytelling™. This is supported by first principles:

1. Cognitive Load Theory - Introducing too much complexity at once overwhelms cognition. Stepwise introduction regulates cognitive load for more effective comprehension.
2. Spiral Learning - Building understanding progressively in cycles allows concepts to be scaffolded and integrated deeper over time.
3. Gestalt Psychology - Harnessing innate human perception tendencies, this approach emphasizes the use of naturally intuitive visual patterns. Gestalt psychology suggests that the human mind prefers to perceive a whole rather than disparate parts.

Existing work on explainable AI (XAI) focuses on interpreting and describing model logic after creation or using secondary AI for model interpretation.

In contrast, our solution enables more complete understanding before model development even begins, and throughout the process. This fosters development of inclusive, safe, and effective AI. It facilitates understanding beyond just AI experts to notice and address anomalies and bias.

Rather than treating symptoms reactively by debugging already-developed models, our solution facilitates intrinsically transparent models by rendering holistic perspectives on data relationships from the start. By visually highlighting when models perform poorly on specific outliers, you transcend average-case XAI to see, discuss and proactively mitigate corner cases that matter. Consider the loan applicant, denied credit due only to a faulty model, or an autonomous military drone mistaking a schoolhouse as a valid target. In developing AI, mastery of corner cases can mean the difference of life and death.

> *"Everything should be made as simple as possible, but no simpler." – Albert Einstein*

---

nmersion Analytics

rrent State of the Field and Limitations

current state of AI has seen substantial advancements, with a growing dependency on complex models
h process multidimensional inputs. Traditionally, AI professionals rely on:

- **Statistical Analysis**: Statistical properties of model outputs may offer insight into the behavior of models. Alas, summary statistics yield an incomplete view. Figure - Anscombe's quartet:

All four datasets are identical when examined using summary statistics, yet vary considerably when graphed.

- **Data Visualization**: Visualizing data and the results of models is a key way to understand what a machine learning algorithm is doing.
  - Pair plots to survey the unordered combinations of data dimensions, the number of such combinations expands geometrically as dimensions are added,
  - Dimensionality reduction techniques like PCA and t-SNE, though useful for clustering, can obscure interpretability by merging multiple variables, such as governance factors, into fewer axes. This risks conflating distinct aspects like hate speech and harassment into a single axis, thus hindering stakeholders from comprehending individual variables.
- **Cross-validation**: Cross-validation evaluates machine learning model performance and generalizability using tools like confusion matrices, ROC curves, AUC, precision-recall curves, learning/validation curves, feature importance, box plots, violin plots, and heatmaps.
- **Feature Importance Analysis**: By evaluating the importance of features in a model, we can gain understanding of which parts of the data are driving the model's decisions. An ability to see higher dimensional data may make the influence of individual features on AI model decisions more accessible and comprehensible to a diverse range of users.
- **Model Interpretation Tools**: Several tools and techniques can help in interpreting model predictions. Examples include SHAP (SHapley Additive exPlanations), LIME (Local Interpretable Model-agnostic Explanations), PDP (Partial Dependence Plots), and ALE (Accumulated Local Effects). The ability to visualize higher-dimensional data may e.g.
  - Complement SHAP by providing a clearer picture of how features interact and influence predictions, something that is not always apparent from SHAP values alone.
  - Impart a more comprehensive understanding of the model's behavior by visually representing both local explanations (from LIME) and broader data patterns.
  - Compare and contrast explanations from LIME across different models, helping to identify consistencies and discrepancies in a more intuitive manner.
  - Providing insights into complex relationships that are not easily captured by PDPs.
  - Improve accessibility and understanding from ALE plots, especially for non-experts.
- **Sensitivity Analysis**: Sensitivity analysis involves changing the input variables in various ways to see how the output of a model changes. Seeing higher dimensional data may complement sensitivity

---

nmersion Analytics

analysis by providing more intuitive and detailed visual representations of how changes in input variables affect model outputs.

- **Ensembling**: Ensembling methods, such as stacking or bagging, may provide insights by looking at it from multiple perspectives. Intuitive multidimensional visualization capabilities may
  - Enhance the comparative analysis of various models, facilitating the exploration of model diversity and correlation, essential for assembling robust ensembles.
  - Render insights into model behavior and effectiveness by seeing how various models process each data point, in the context of key features across all data points.
  - Improve predictive accuracy and also ensure creation of more balanced, effective and well-adapted ensemble models.
- **Disentangling Causality**: This is about understanding which variables have causal effects on the outcomes. Various techniques like Do-calculus, Causal Graphical Models, or Interventional Perturbation are used for this. Adaptations of the proposed visualizations may more clearly illustrate the relationships and potential causal connections between variables.
- **Counterfactual Explanations**: This involves understanding model decisions by asking "what if" questions. By slightly changing the input and observing the output, we can gain insights into the model's decision-making process. The proposed visualizations may more effectively illustrate the outcomes of "what if" scenarios, showcasing how slight changes in inputs impact model outputs.
- **Understanding Bias and Fairness**: It is important to understand and evaluate the bias in the data and the predictions made by the model. Tools like Fairlearn and AI Fairness 360 provide metrics and methods to audit and mitigate bias and discrimination in the models. The proposed visualizations enhance the ability to leverage domain knowledge to identify and comprehend complex patterns of bias, supporting more effective auditing and mitigation efforts.
- **Debugging ML models**: Debugging tools like TensorBoard, What-If Tool, or Explainable Boosting Machines (EBMs) provide interfaces to understand and debug ML models. The proposed visualizations may offer a unique visual lens for seeing what the model does with various datasets, without relying on summary statistics, useful for debugging complex models by observing the same at each training epoch.

isualization aspects of current techniques are limited by the conventional X vs Y (and perhaps Z) plot,
ence presenting an incomplete and/or distorted view. As explained above, this poses significant challenges
n understanding and analyzing higher-dimensional data underpinning models. Lack of adequate data
isualization has become a major bottleneck and blindspot. Our solution fills this widening gap to mitigate
otential threats as AI advances rapidly.

Solutions

Machine Learning

Our solution also enables the following use cases to enhance machine learning workflows through novel data visualization. This enables stakeholders to gain deeper insight into complex AI models, aiding in bias detection, model optimization, and inclusive governance by engaging diverse stakeholders.

Use Cases

Enabled by the following proposed features, this work aims to prototype three key use cases,

1. **Enhanced Predictive Analytics**: The ability to improve accuracy of predictive models by enhanced Exploratory Data Analysis (EDA), Data Quality and Model Selection & Ensembling features underpins the ability to forecast and make informed decisions. Whether it's predicting maintenance needs, consumer behavior or health outcomes, this plays a key role in strategic planning and operational efficiency.

---

mersion Analytics

Efficient Hyperparameter Tuning: This use case is critical because it directly impacts the effectiveness and efficiency of AI models. By reducing time and computational resources needed for tuning, this capability not only accelerates the model development process but also enables creation of more accurate and reliable models.

Real-Time Performance Monitoring: The ability to continuously monitor and evaluate the performance of AI systems in real-time, as enabled by the Production Monitoring feature, is essential for maintaining the reliability and effectiveness of these systems. This use case is particularly important because it ensures that AI applications remain functional, accurate, and efficient over time, adapting to new data and conditions. This ongoing monitoring is crucial in scenarios where AI systems control critical processes or make important decisions, ensuring that any deviations or anomalies are quickly identified and addressed.

Education: Instructors leverage the following capabilities to employ interactive visualizations to enhance the teaching of AI and machine learning concepts to a diverse group of students. These visual tools demystify complex topics like neural networks and decision trees, making them accessible to students with varying backgrounds. During lectures, these visualizations dynamically illustrate AI processes, such as how weights and biases evolve during neural network training. Students engage with these tools in hands-on activities and assignments, manipulating parameters to see real-time effects on model behavior. This interactive approach not only deepens their understanding of AI principles but also increases engagement and enthusiasm. The use of visualizations effectively bridges the gap between theoretical learning and practical application, resulting in a more effective and engaging learning experience, and preparing students for more advanced studies or careers in AI and data science.

bilities

Exploratory Data Analysis (EDA): This solution offers intuitive visualizations enabling you to explore data before building AI models, allowing you and your team to apply your domain knowledge to understanding data patterns and trends. By seeing the data first-hand, you design models informed by your expertise and context. This enhances model relevance and effectiveness by ensuring it's grounded in a deep, human-informed understanding of the underlying data.

Data Quality: This complements traditional data quality methods via novel visualization to enable your domain experts to leverage tacit knowledge at scale for identifying subtle, context-specific issues and anomalies that may be missed by standard automated checks. This leverages domain expertise for ensuring models are trained on data that is not only technically sound but also contextually accurate and relevant.

3. **Model Training & Debugging**: The solution enables visualizing the model training process, allowing you to simultaneously observe the model's behavior and output across all data points and key dimensions for both training and test datasets at each epoch. It renders a comprehensive view that illuminates the model's interaction with the entire dataset over time, facilitating a deeper understanding of its learning dynamics. This is crucial for swiftly pinpointing and addressing issues, thereby streamlining and improving the AI model training and debugging process.
4. **Hyperparameter Optimization**: One of the more costly aspects of building and fine-tuning models is the selection and tuning of hyperparameters. The hyperparameter space is typically vast and multidimensional, so searching this space is both computationally expensive and time-consuming. While various techniques such as grid search and Bayesian optimization are currently used, they often require a significant computational budget and do not always find the optimal solution. This can potentially become an $O(n^m)$ problem (in the brute force case), where m is the number of hyperparameters, and n is the number of data points. By instead randomizing a tractable set of hyperparameter combinations then visualizing model error for each, domain experts engage with AI developers thoughtfully on pragmatic ways to best narrow the search space, reducing training costs by potentially an order of magnitude or more.
5. **Model Selection**: Visually compare how multiple models process each data point within a dataset, particularly focusing on the top dimensions. By providing a comprehensive visualization that

---

mmersion An

displays the entire dataset in the co
how different models behave with s
selecting the most appropriate mod
entire dataset, ensuring a more info

6. **Model Ensembling**: Expanding o
and combined effects of various A
each data point. By visually repres
users to explore and understand th
constructing robust ensemble mod
understanding of how different m
7. **Real-Time Performance Monit**
in production AI systems. It rend
batches, enabling users to intuitiv
emphasizes the use of human ins
and context-aware detection of a

Large Language Models

Immersion Analytics plays a key role in
Envision a data space where points repr
moderation scores including hate speech
Dimensional Engine, to reveal underlyi
and areas requiring guard railing. While
reduce vector embeddings of the text or

This visual framework both enhances i
in the LLM's outputs that may diverge
research and development as well as f

Architecture

Interoperability across teams and t
appreciate flexibility to choose the
Immersive Data Visualization Eng
library of API's and software inte
any model, algorithm, dataset, and
deployment options including Br
and even the world's first glasses-

Next Steps

Reach out to contact@immersion
technologies upgrade AI workflo

Virtual C

Patent Issued

Immersive Data Visualization Types

All capable of animating real-time streaming data

Dimensional Scatterplot

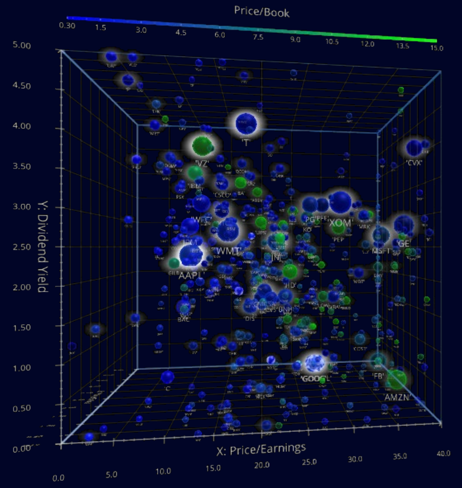

Radial Time Series

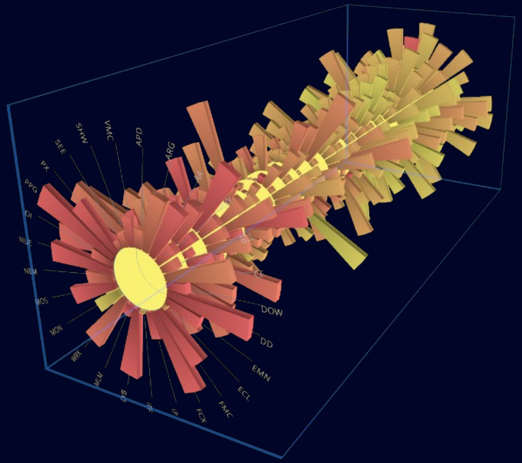

Weights & Biases

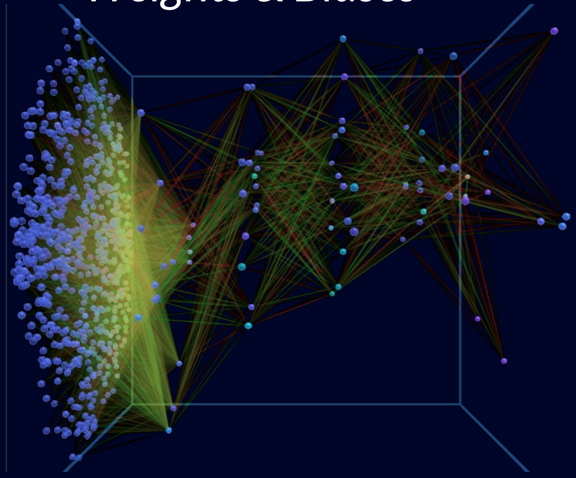

Dimensional Bar Chart

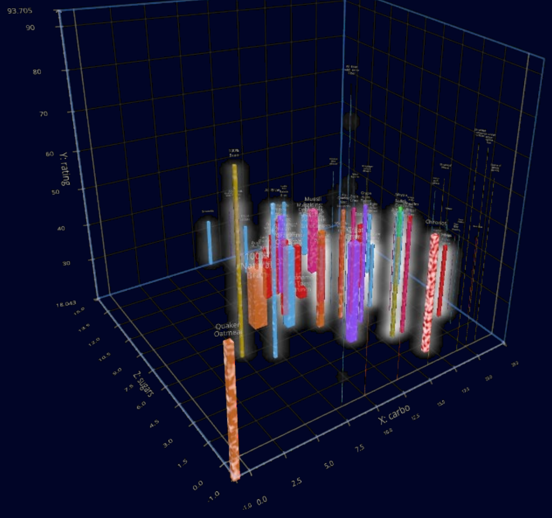

Dimensional Surface

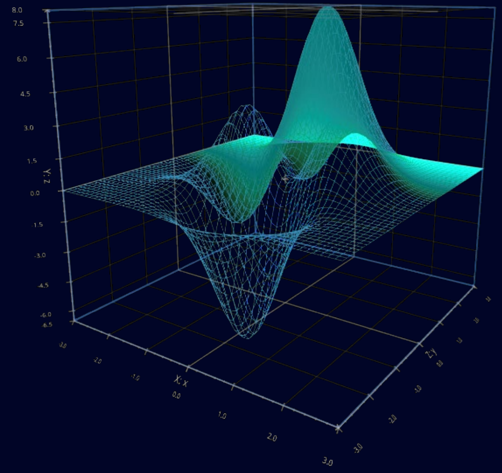

Dimensional Time Series

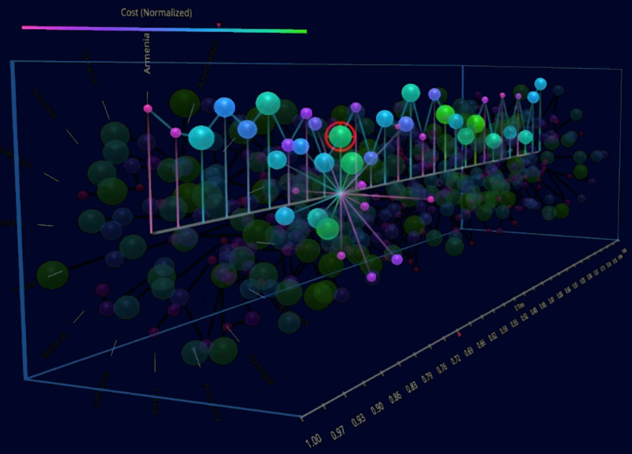

Dimensional Graph

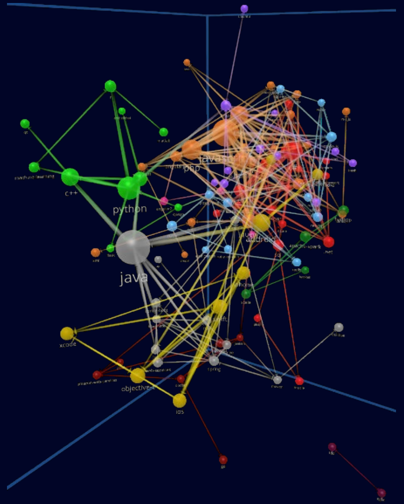

Dimensional Line Chart

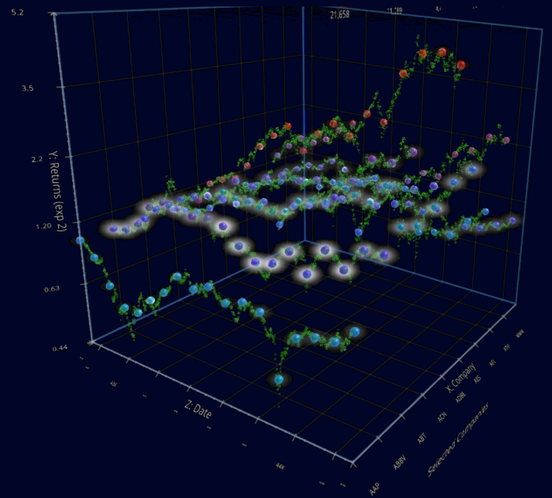

Immersion Analytics

Patent Issued

https://www.immersionanalytics.com/viz/demos/dashboard

<-- Try It!

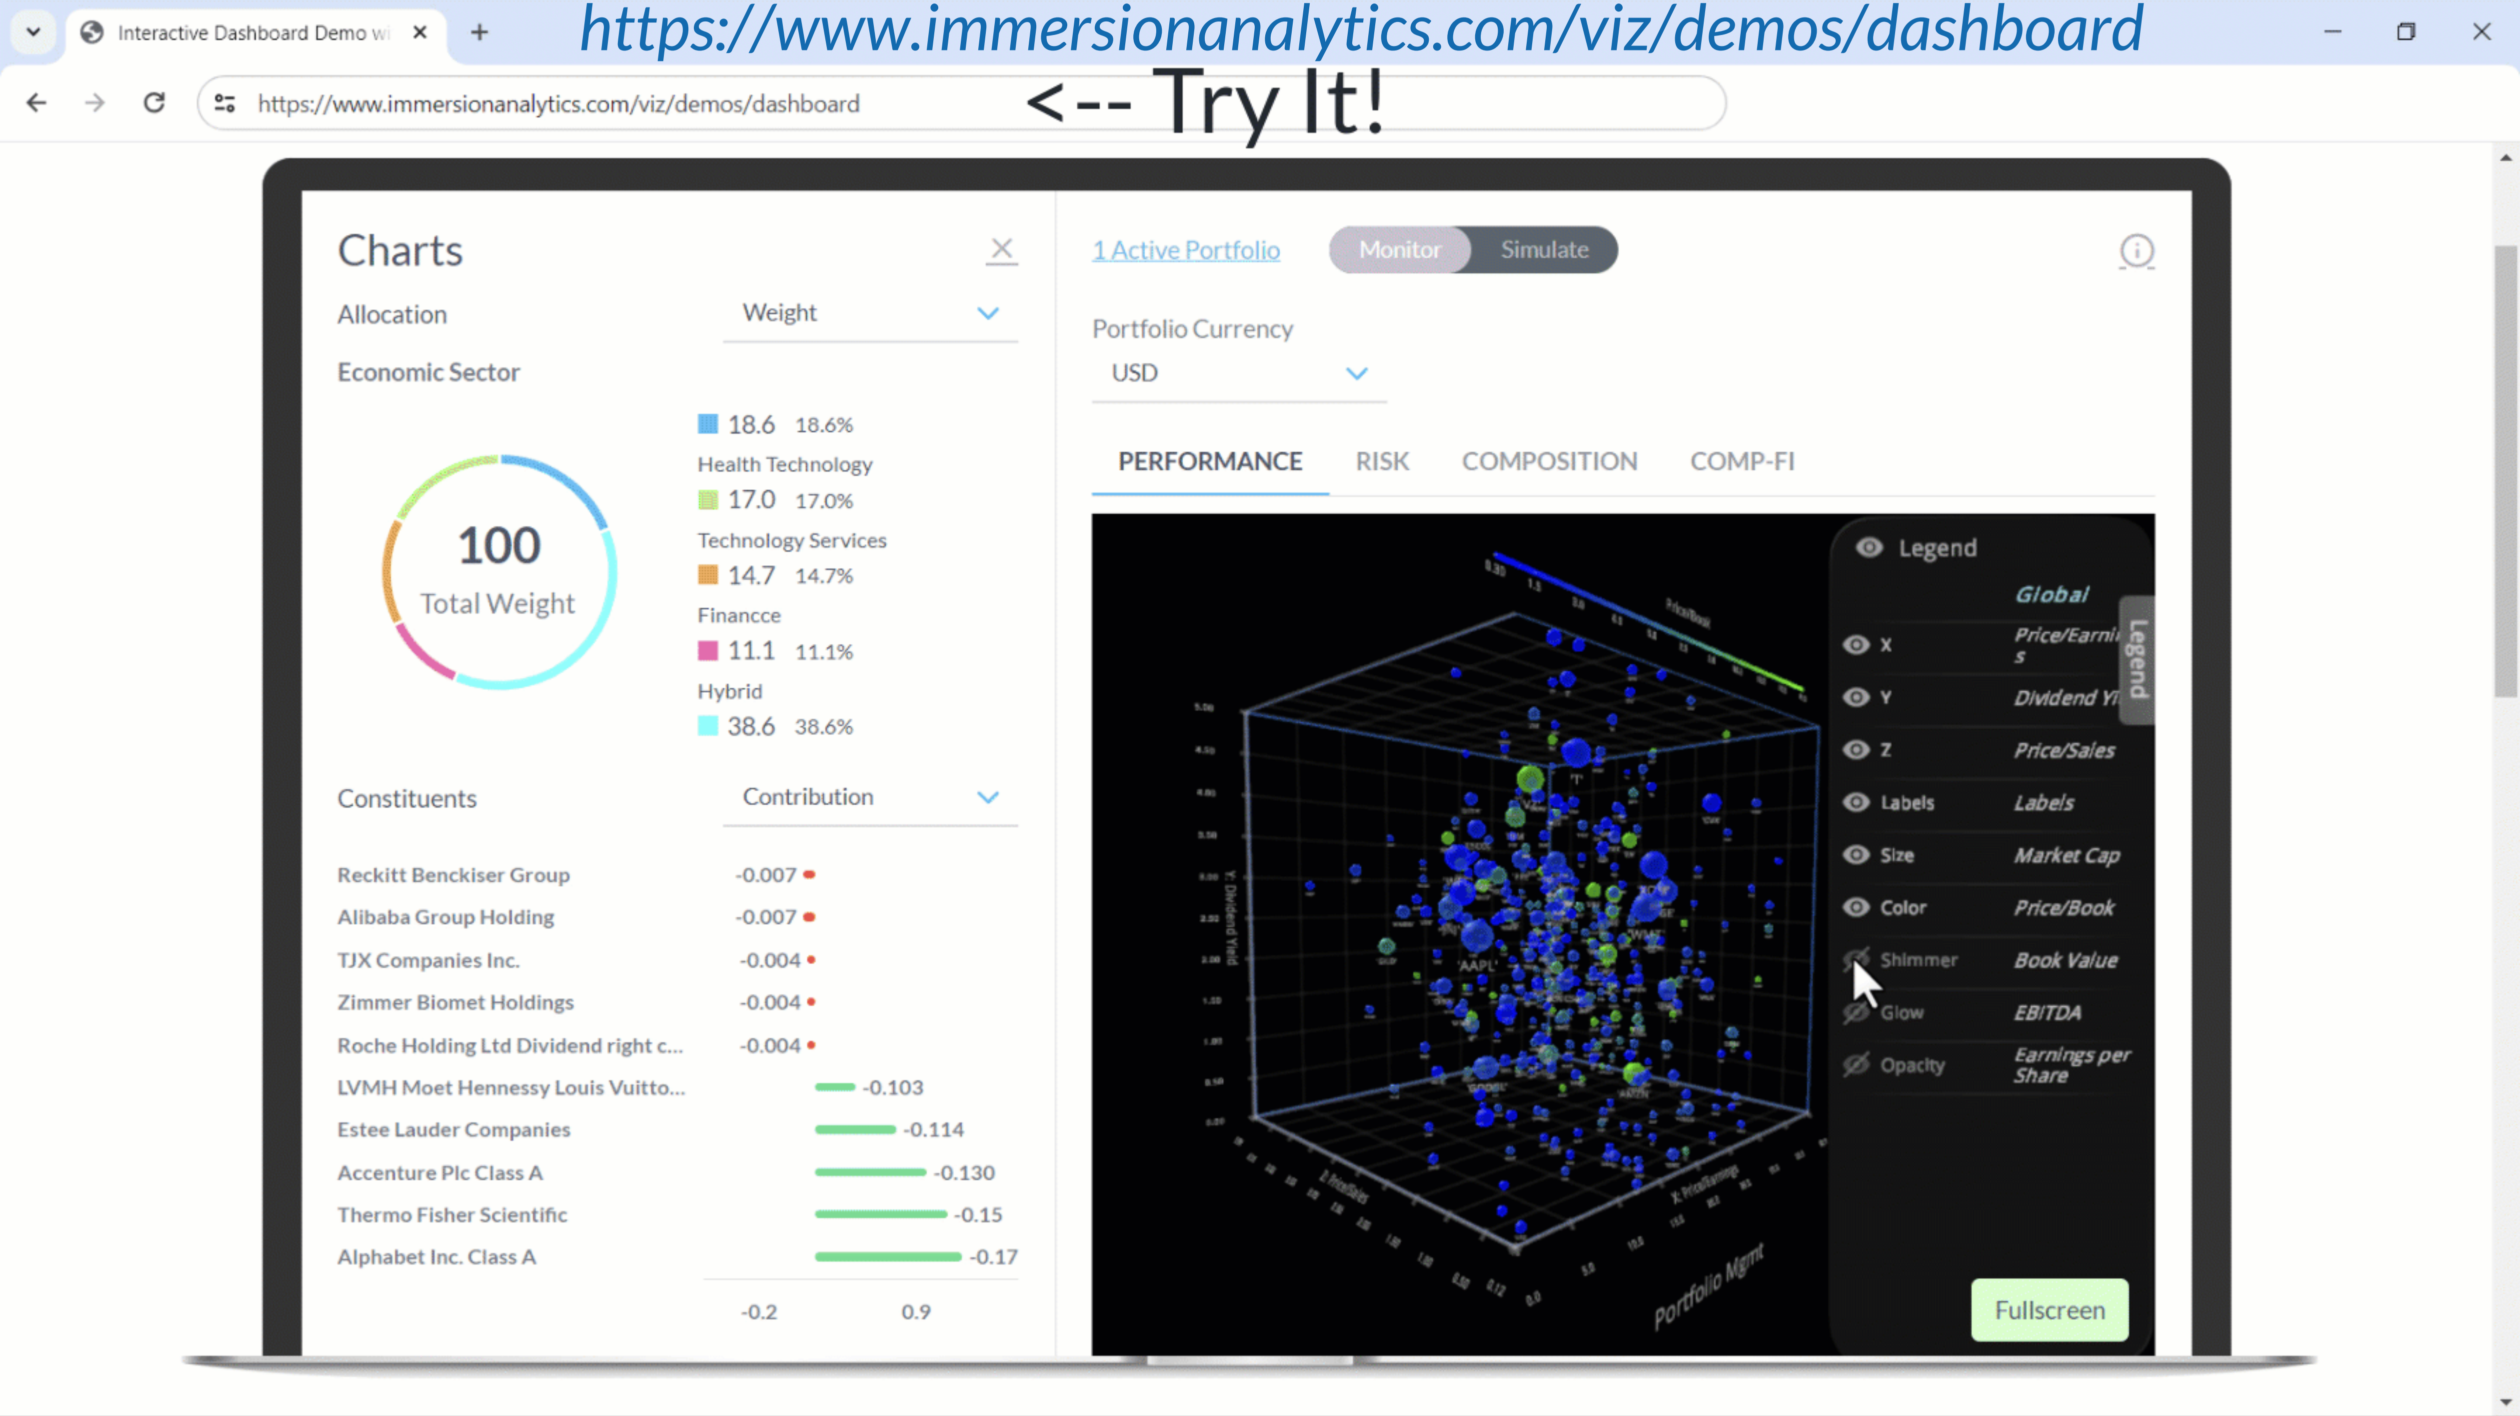

Charts

Allocation: Weight

Economic Sector

100
Total Weight

Sector	Weight	%
🟦	18.6	18.6%
Health Technology 🟩	17.0	17.0%
Technology Services 🟧	14.7	14.7%
Financce 🟪	11.1	11.1%
Hybrid 🟦	38.6	38.6%

Constituents: Contribution

Constituent	Contribution
Reckitt Benckiser Group	-0.007
Alibaba Group Holding	-0.007
TJX Companies Inc.	-0.004
Zimmer Biomet Holdings	-0.004
Roche Holding Ltd Dividend right c...	-0.004
LVMH Moet Hennessy Louis Vuitto...	-0.103
Estee Lauder Companies	-0.114
Accenture Plc Class A	-0.130
Thermo Fisher Scientific	-0.15
Alphabet Inc. Class A	-0.17

Portfolio Currency: USD

Legend

Global

X	Price/Earnings
Y	Dividend Yi
Z	Price/Sales
Labels	Labels
Size	Market Cap
Color	Price/Book
Shimmer	Book Value
Glow	EBITDA
Opacity	Earnings per Share

Interaction Amplifies Perception

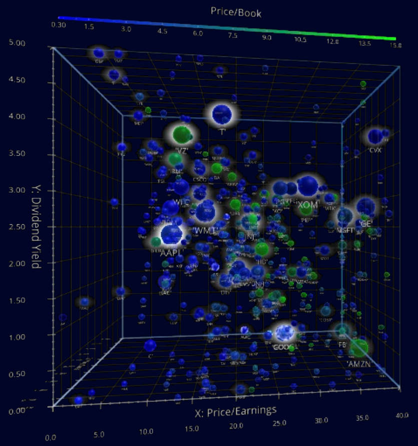

Partial

Full

Static Plot on Paper

Static Plot on 2D display

Animated Plot on 2D display

Animated Plot on 2D touchscreen

VR / AR / 3D Laptop

Patent Issued

Immersion Analytics – 2 Main Elements to our Solution

Immersive Visualizations

IA Visualizer
(Web Embeddable and/or Client Apps)

Runtime Connectivity Platform

IA Runtime Server
- Docker Image
- P2P

IA Runtime Library (SDK)

IA Runtime Client = IA Visualizer

Patent Issued

Immersion Analytics – 2 Main Elements to our Solution

Immersive Visualizations

- Visualization Engines & Real-time Support
  - Scatterplots, Surface Chart, Time Series, Network Graph, Bar Charts & DataTubes
- PC, Mac, iOS, Android, VR & AR devices & Web Embeddable
- Rendering & patented Dimensional Engine™

Runtime Connectivity Platform

- Data Support (proven w/ Tableau, Qlik, MATLAB & Python)
- Bi-Directional Interactivity
- Integration via Familiar Programming Interfaces
- Complements Existing ISV Applications

Patent Issued

Immersive Computing Runtime ~ Deployment

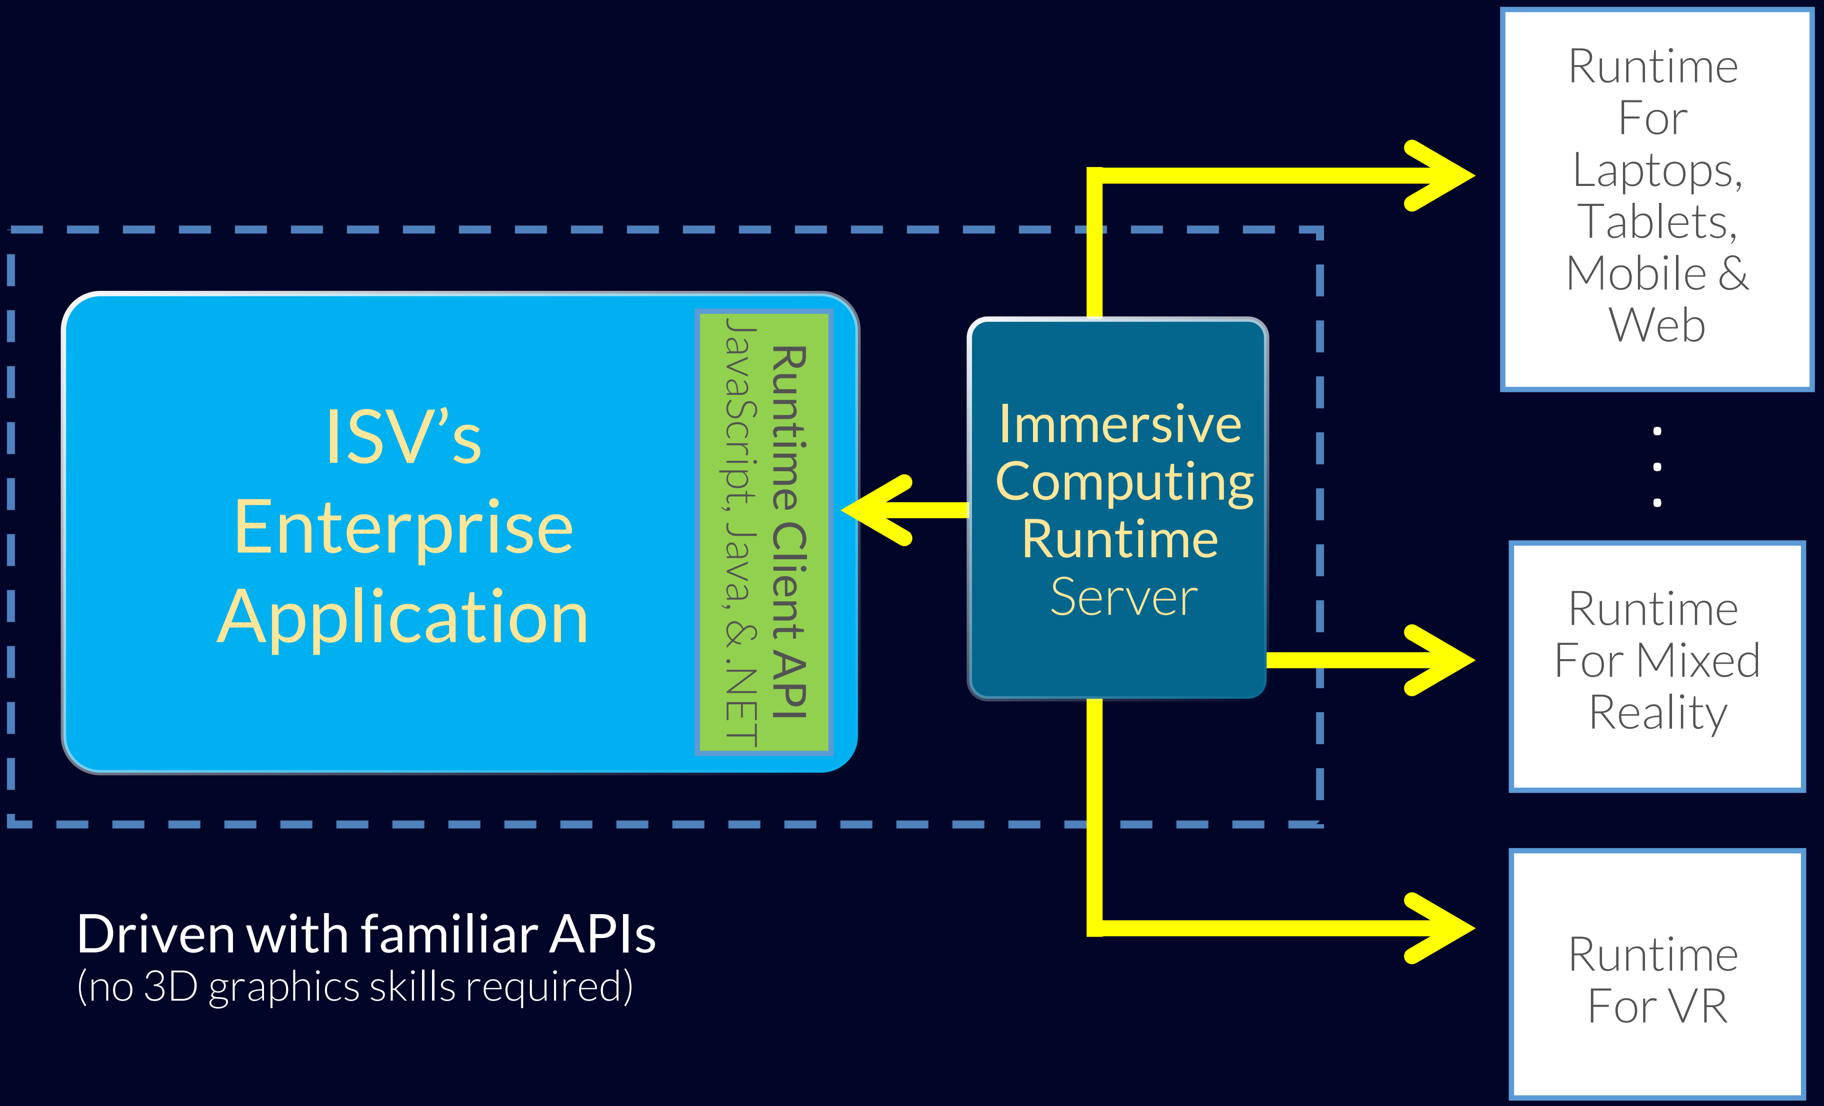

ISV's
Enterprise
Application

Runtime Client API
JavaScript, Java, & .NET

Immersive
Computing
Runtime
Server

Runtime
For
Laptops,
Tablets,
Mobile &
Web

⋮

Runtime
For Mixed
Reality

Runtime
For VR

Driven with familiar APIs
(no 3D graphics skills required)

Patent Issued

Immersive Computing Runtime ~ Protocol Flows

Create Scene Object

```
client = ia.RuntimeClient()
scene = client.Scene
```

Create Dataset Object & Add Data

```
my_dataset = Dataset('My Portfolio Data')
apply_my_data_to_ia(my_data, dataset)
scene.Database.Datasets.Add(my_dataset)
```

Create Visualization Object

```
my_viz = ScatterViz('My Portfolio')
scene.Visualizations.Add(my_viz)
my_viz.Axes.Map(
     'X', 'Volatility',
     'Y', 'Price',
     'Z', 'Market Cap',
     'Size', 'Volume',
     'Color', 'Sales',
     'Glow', 'Margin’ )
```

Create Color Map Object

```
my_viz.ColormapName = 'Red/Yellow/Green’
```

Assign Viz Object to Data Object

```
my_viz.Axes.SourceDataset = 'My Portfolio Data'
```

Patent Issued

Immersive Computing Runtime:
Real-time Streaming & Bi-directional Controls

Real-Time Streaming

- Polling at regular intervals
- Event driven data updates

```
while True:
    my_data = read_data_from_server(...)
    apply_my_data_to_ia(my_data, my_dataset)
    thread.sleep(.5)
```

```
def update_dataset(new_data):
   apply_my_data_to_ia(new_data, my_dataset)
datasource.on_data_received += update_dataset
```

Bi-Directional Interactive Controls

- Events triggered by the user

```
def handle_selection_event(selection):
    ... # do something
viz.SelectedPointProperty.OnChanged += handle_selection_event
```

Patent Issued