Towards a Paradigm for Visual Modeling

Fahd Husain

Director, AI Research

uncharted

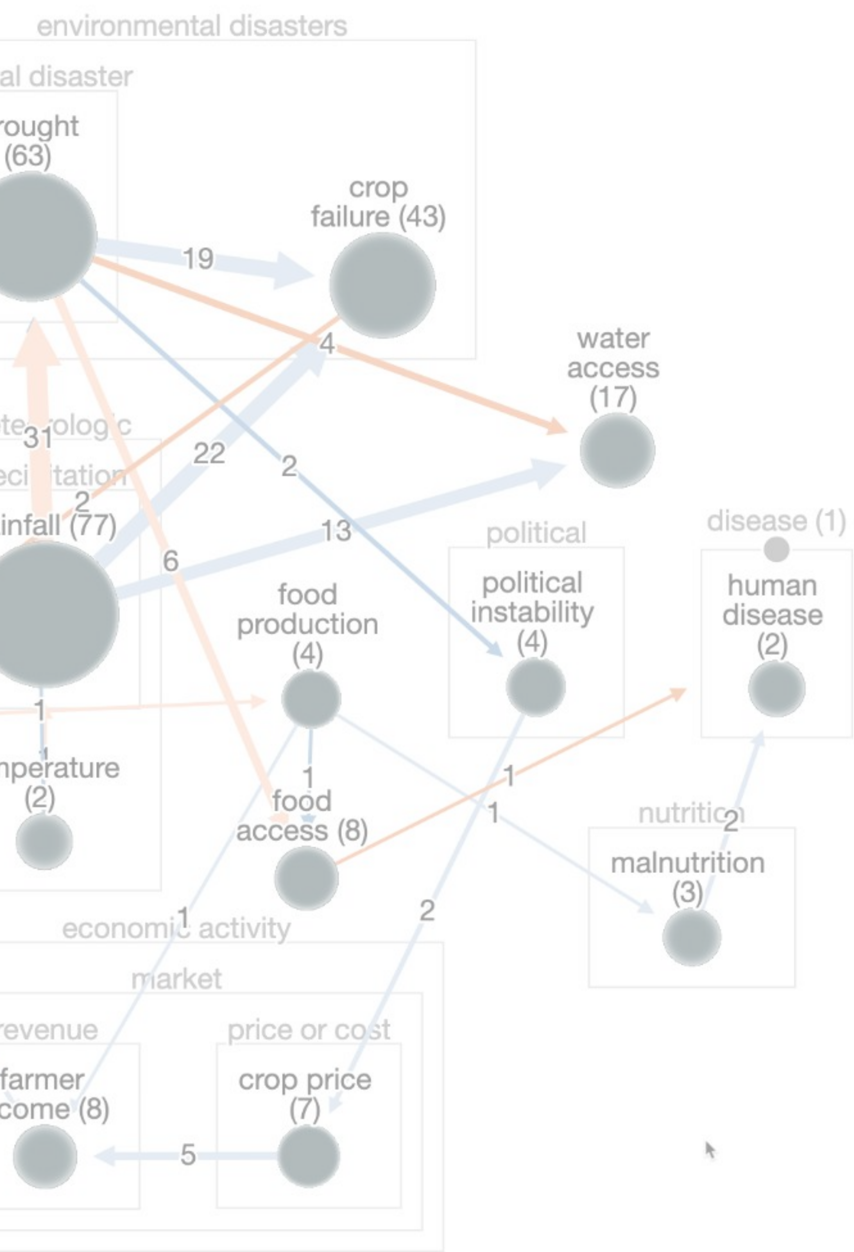

Outline

I: State of the Art & Philosophy

- Visualization & ML
- Augmented Intelligence

II: Research Efforts in DARPA Programs

- World Modelers
- Automated Scientific Knowledge Extraction

III: Towards a Paradigm

- Challenges
- Visual Modeling

I. SOTA & Philosophy

Visualization & ML

Visualization recently gained a foothold in the field of AI research.

Typically, work focuses on visualizing modules or specific dynamics of ML models for:

- debugging
- explainability or interpretability
- pedagogy

"Auto-ML" platforms have also emerged, focused on machine-assisted model workflows so ML solutions can be easily leveraged for common problems.

Explainable AI

A host of recent efforts have been focused on model debugging, interpretability and explainability.

- Tracing the activations across an unfolded RNN or LSTM for a given input.
- Attention mechanisms highlight aspects of training or prediction to help with interpretation.
- Visualizing image activations at various CNN network layers.
- Allowing for real-time experimentation or ‘what-ifs’ by changing input data features and execution parameters.
- Linking classified data samples in terms of prediction.

Focus is typically on a single model, interrogating its training and prediction dynamics. Model is taken to have been assembled externally, and these visualizations are for experimentation or debugging.

[15]

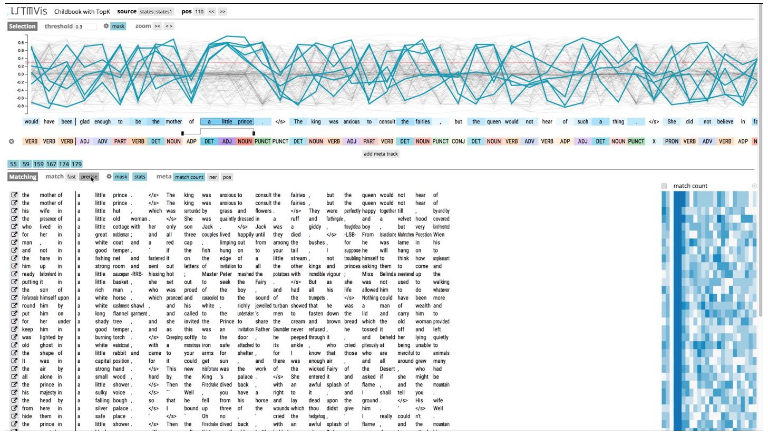

[18]

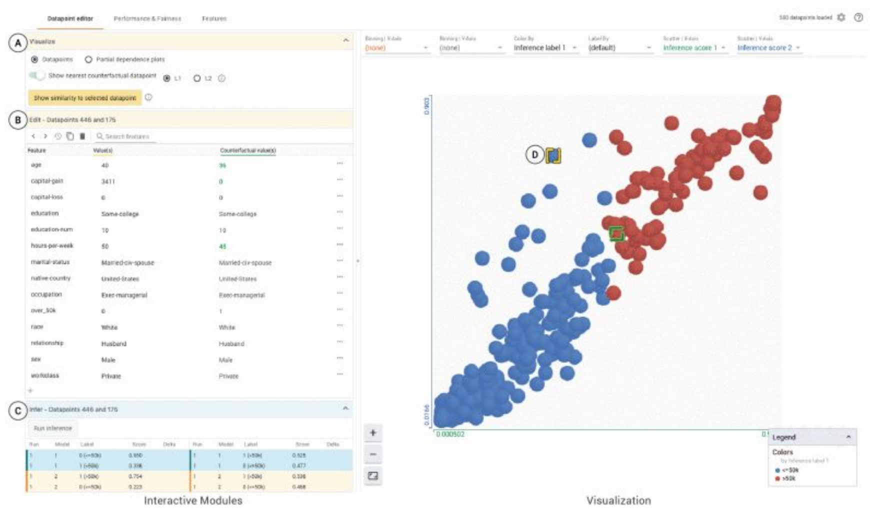

[1]

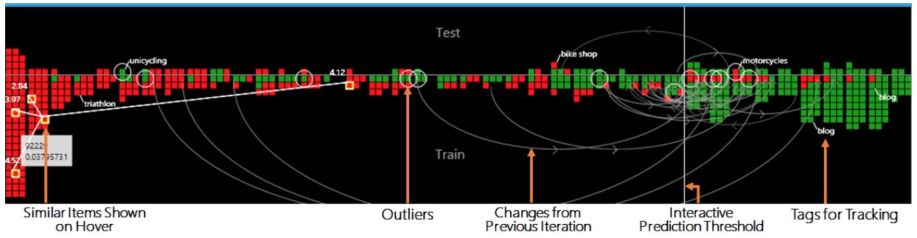

Pedagogy

Similar focus of various interactive tutorials and ML ‘playgrounds’.

Pedagogical resources to understand core ideas and visualize complex data transformations and training dynamics.

Interactive features allow users to explore input and parameter space within some predefined intervals to see outcomes

[11]

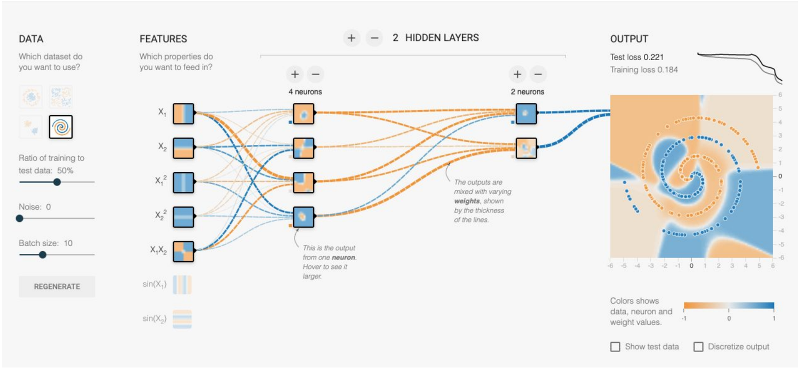

[17]

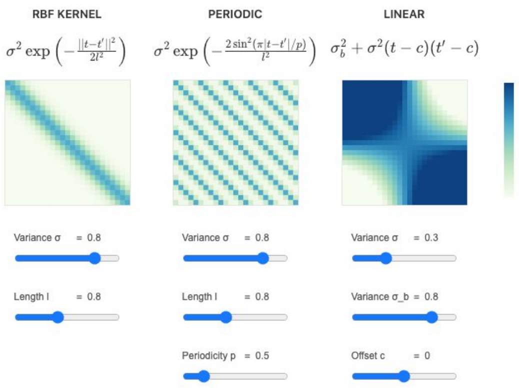

[2]

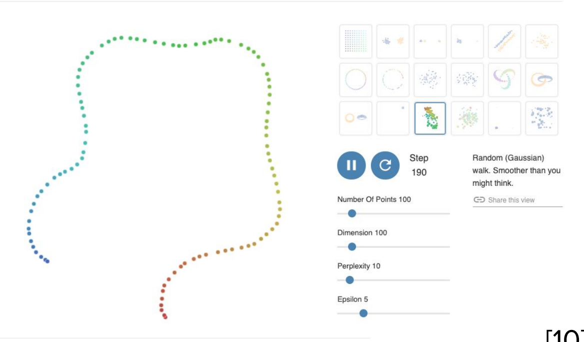

[10]

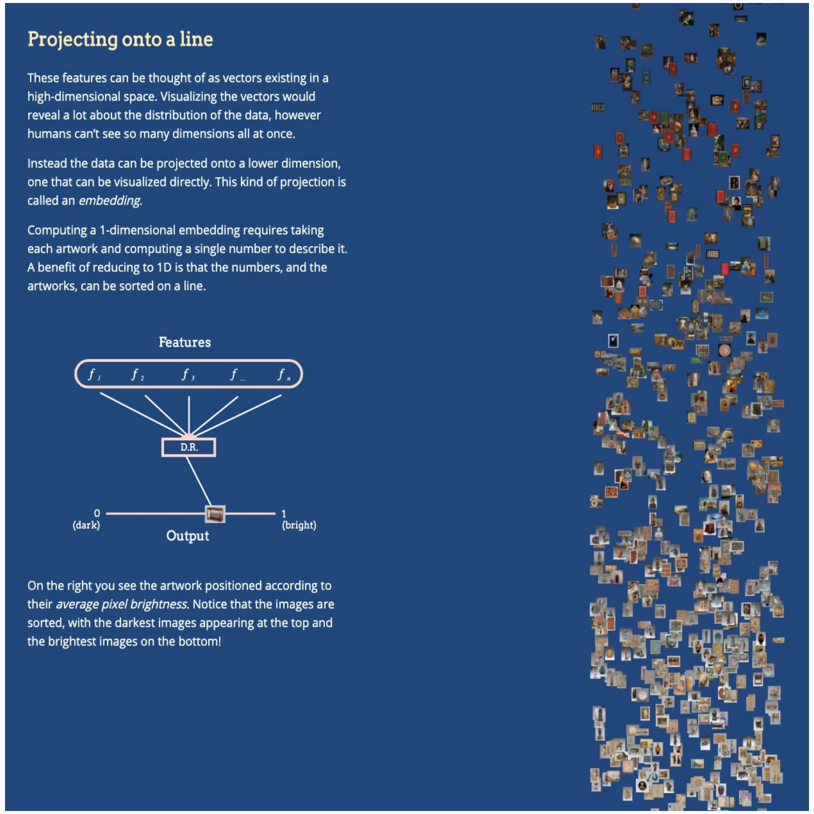

[16]

AutoML interfaces

Larger AutoML systems are emerging that look to streamline the full ML workflow from ingesting input data to model inference.

Geared towards enabling non-experts to use ML tools or accelerating some work by ML practitioners.

Typically focused on common industry problems (classification, clustering, linear models).

Sometimes remain in an awkward place: either users of AutoML systems still need ML knowledge for model selection, training, parameter optimization; or these choices are entirely black-boxed away.

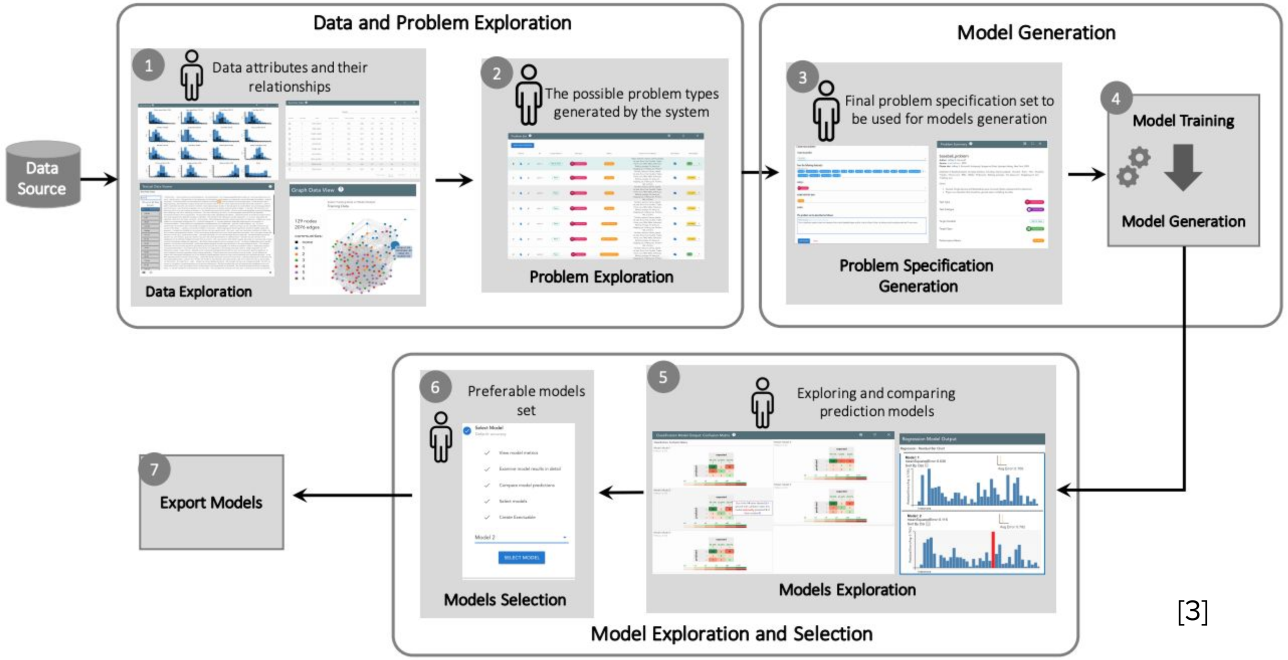

[3]

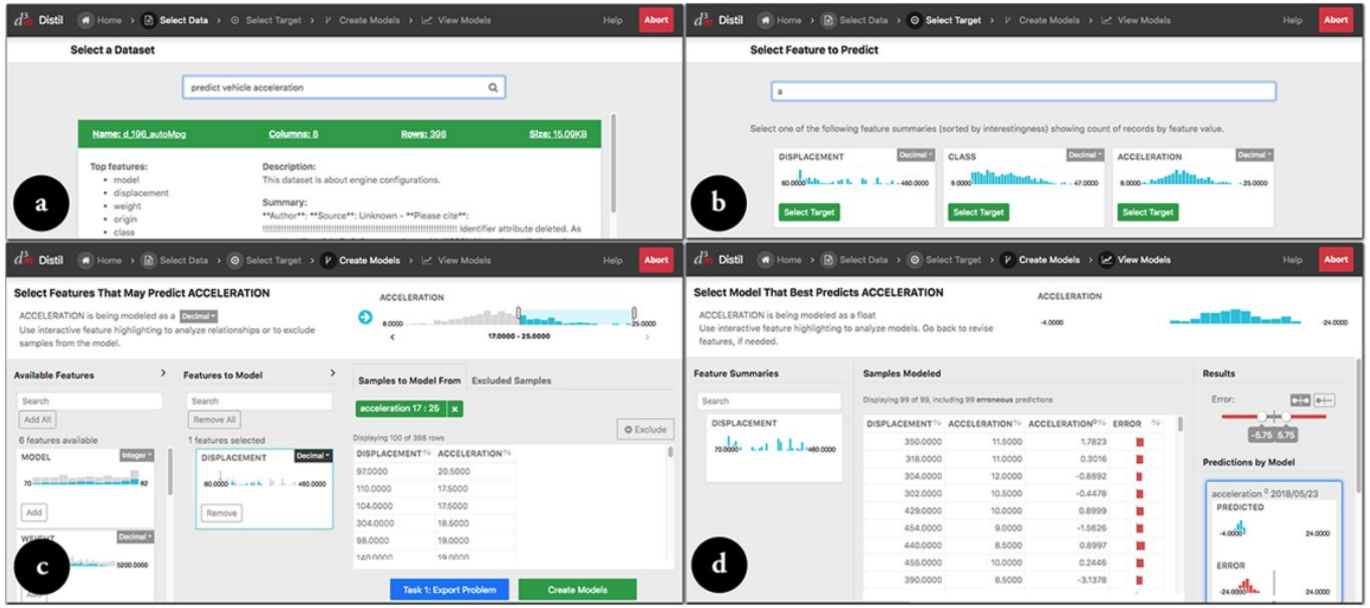

[13]

Visual Modeling & Augmented Intelligence

Artificial intelligence, as often benchmarked against human intelligence, attempts to replicate human skill. It is better understood as a geometric intelligence: it has an entirely different, non-embodied ontology, lacks intuition and creativity, struggles with generalization and synthetical understanding. Has near-perfect memory and recall, calculation capability, pattern detection at inhuman scale.

Augmented intelligence is a historical concept dating back to Doug Engelbert and JCR Licklieter. Core emphasis is on the human-in-the-loop, where human and machine intelligence serve to augment each other in real time, interactively, in a human-machine interface. Not trying to replicate human skill in a machine; but technologically amplify human skill.

Hypothesis is that (while we wait for AGI) augmented intelligence is the paradigm to follow today for large-scale systems attempting to solve compositional problems (not end-to-end learning tasks such as translation, vision, etc).

Visual modeling as explored here is cast as a form of augmented intelligence.

II. Research Efforts in DARPA Programs

World Modelers (WM)

- DARPA program aiming to model socio-natural complex systems such as drought, famine, food security [4].
- Unsolved multi-domain, multi-resolution problem.
- Involving systems of systems, with intricate linkages and causal dynamics.
- Problem has both qualitative and quantitative dimensions and types of information, with large stores of data and knowledge artifacts.
- Current computational models in the space typically siloed in their domains (agricultural yield, political instability, disease spread, rainfall estimates, flood prediction).

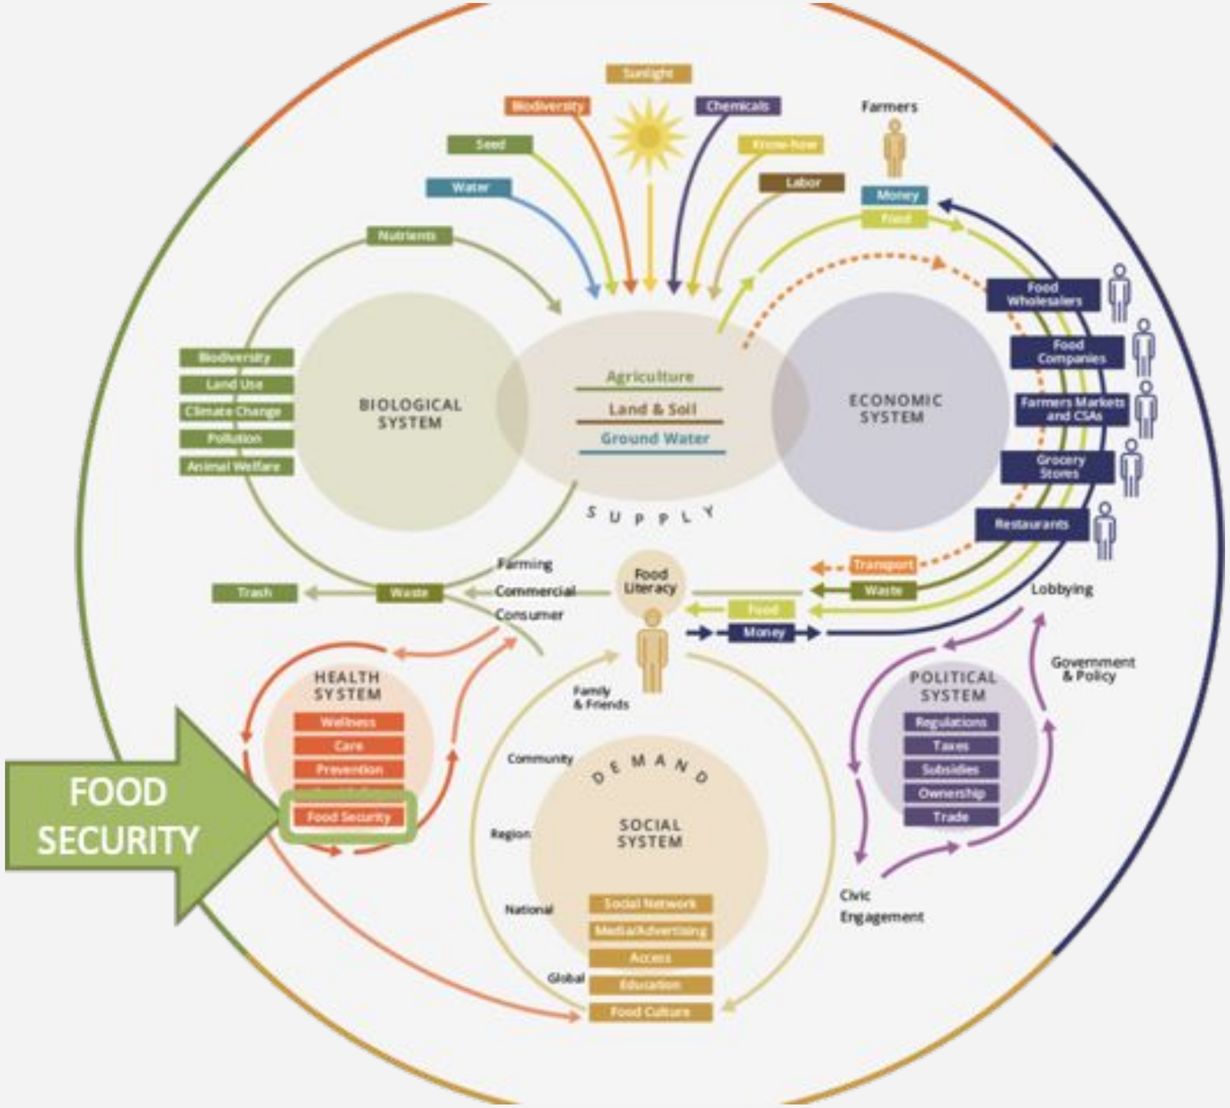

World Modelers

Scenario(s):

- Emergency response: famine event has occurred; aid is required.
- Forecasting dynamics: rain is predicted to be below average - which areas will be hardest hit? What other consequences will result?
- System awareness: decision-maker needs a briefing on main drivers of food insecurity, and linked causal factors and impacts.
- Multi-domain analysis is required, as use-cases shift quite drastically (for e.g. locusts, conflict).

Current status:

- Takes months to assemble data; scour literature for new factors; build and test models; make predictions; capture them in artifacts; present to decision maker; deliberate over policy; execute policy; engage logistics network to deliver aid or provide support.

Need to enable:

- Exploration of networks of factors and knowledge hypotheses in literature at large.
- Access to disparate data from expert models and open-source indicators.
- Quick assembly of models of complex systems, compare / contrast alternatives.
- Intervention analysis to experiment with causal dynamics.
- Adaptation to new domains as necessary.
- Summary artifacts for downstream decision-makers.

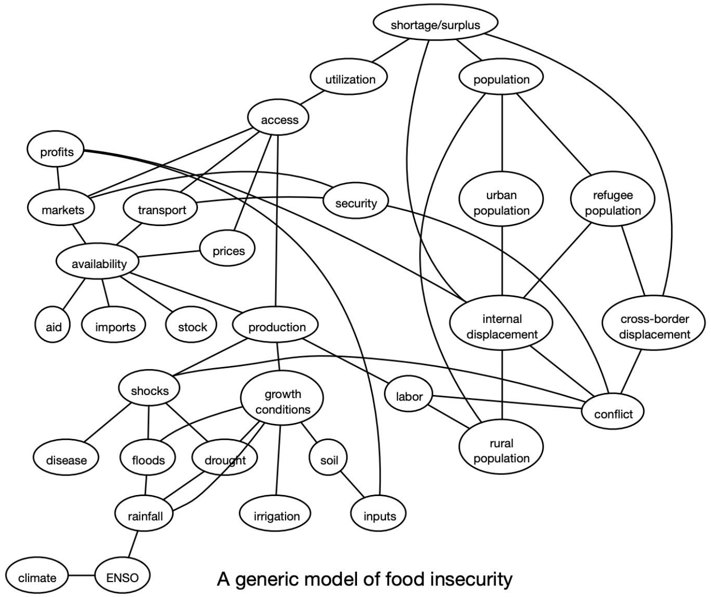

A generic model of food insecurity

[5]

- Step 1: Generate or retrieve a *generic* qualitative, causal food security model;
- Step 2: Modify the model for the specific analyses of Southern Sudan;
- Step 3: Build workflows of expert, quantitative models, where available;
- Step 4: Parameterize quantitative models and the qualitative, causal model;
- Step 5: Configure scenarios and run analyses, producing quantitative results for factors of interest (e.g., food prices, calorie intake);
- Step 6: Produce an "uncertainty report" that documents sources of uncertainty, run uncertainty-reduction procedures and sensitivity analyses;
- Step 7: Identify possible actions to affect factors of interest (e.g., peacekeepers at markets).

Platform: Causemos

- Platform developed for the WM program.
- **Abstraction layer** that synthesizes various data, knowledge and model artifacts.
- Visual, interactive, scalable support for:
  - Exploration and curation of **knowledge** graphs to extract qualitative theories of change.
  - **Data** analysis capabilities across open-source indicators and high-resolution outputs.
  - Iterative assembly of computational/probabilistic graphical **models**.
- Utilizes principles of augmented intelligence in a human-machine interface.

Causemos Conceptual Schema

Data | Knowledge | Models

- ➔ Structured data (numerical data / tables)
- ➔ Spatial outputs from expert high-resolution models
- ➔ Data analysis and transformation capabilities

- ➔ Source knowledge (documents)
- ➔ Structured knowledge (ontology, knowledge graph)
- ➔ Graph exploration and curation capabilities

- ➔ Library of assembled models
- ➔ Single-model parameterization, validation, execution
- ➔ Design of experiments, intervention analysis

- Overall **conceptual structure**, ‘tabs’ in Causemos.
- Represent three different **information streams / artifacts**.
- Sections allow for sorting external capabilities, and helps compartmentalize **integration** along section axes.
- Also helps formulate **workflow** via linkages within and between different sections.

WM Program Architecture

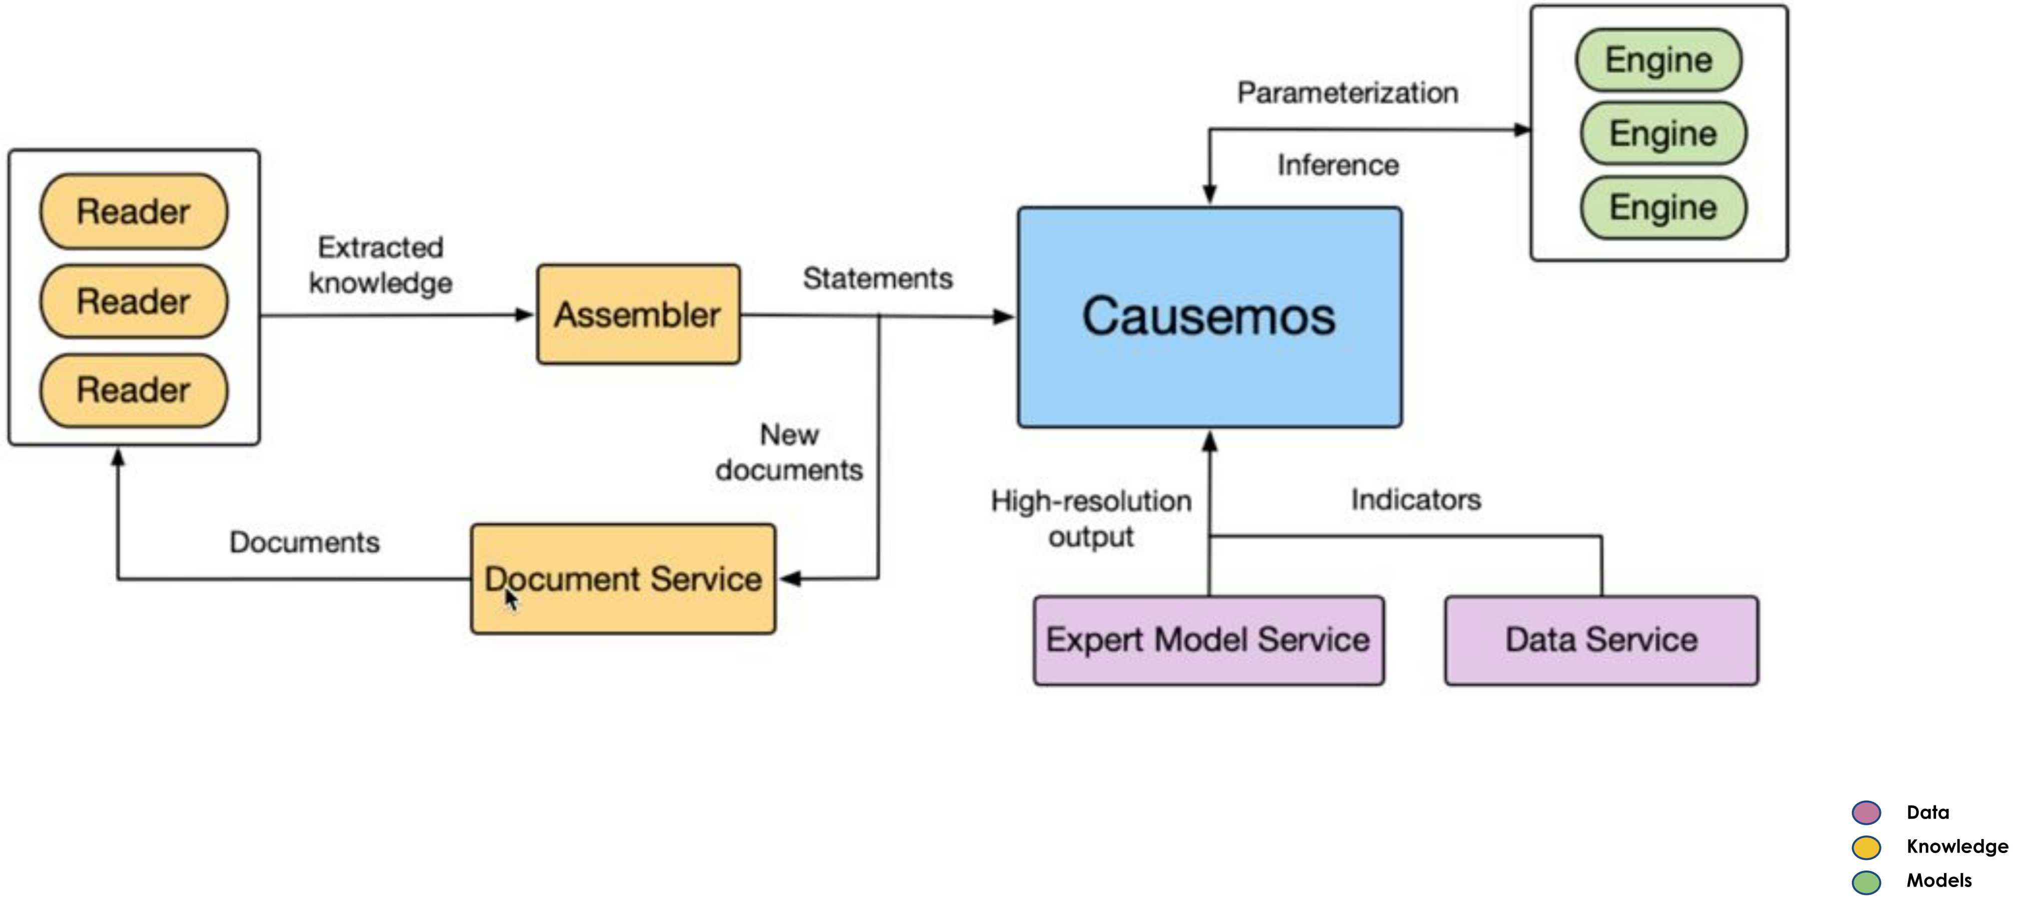

Source Knowledge

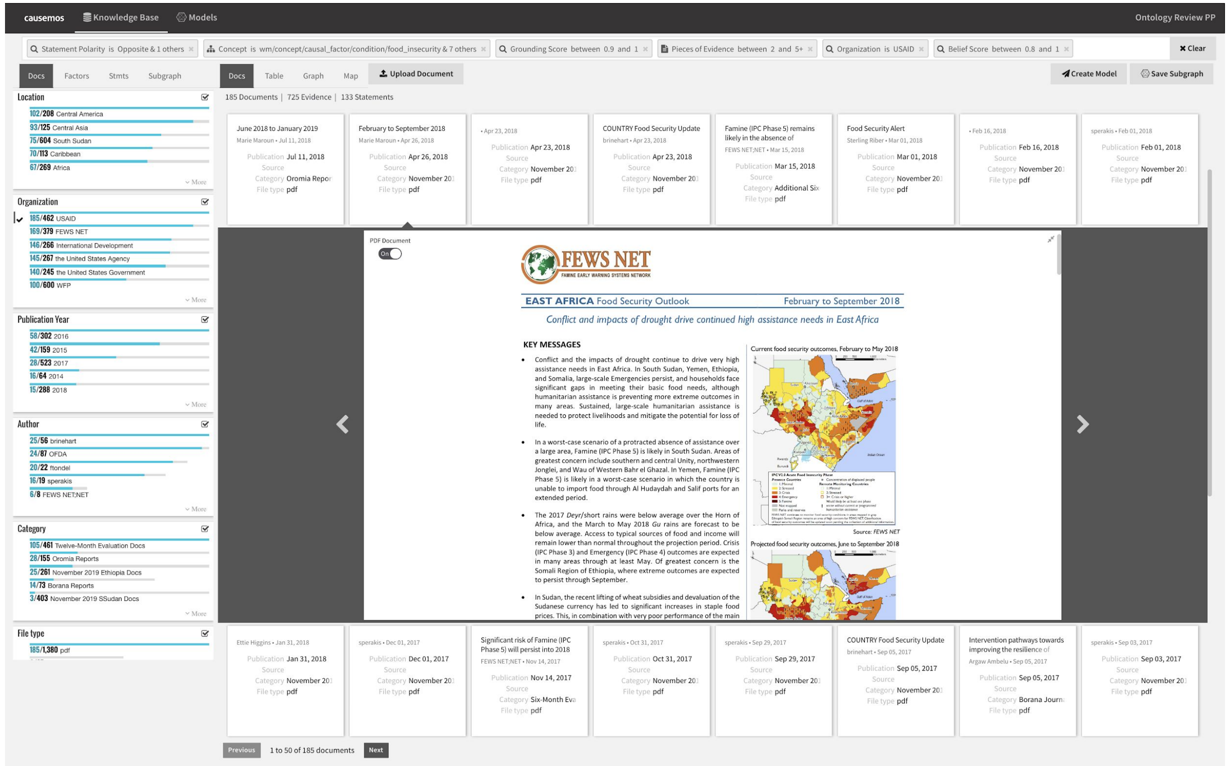

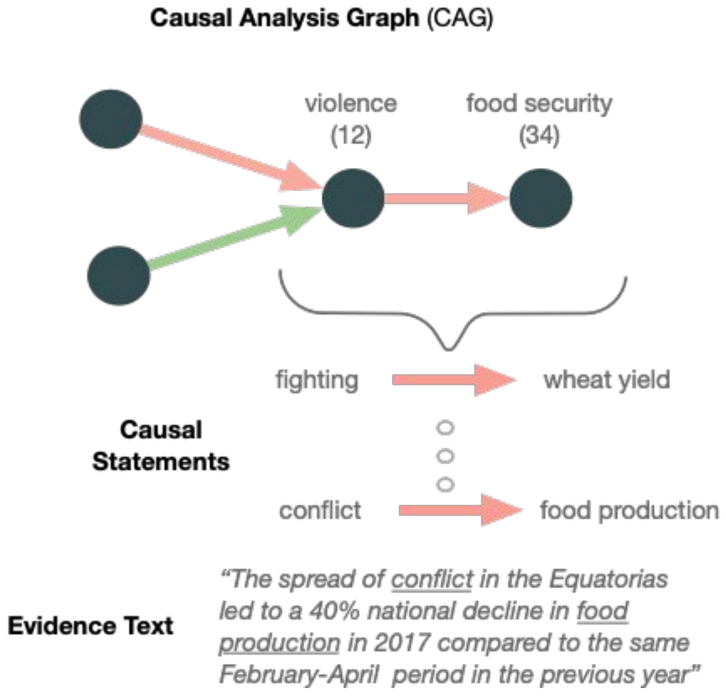

UN reports, news articles and NGO literature are examples of **source documents** (left) from which **causal statements** are extracted, which are then aggregated and cast as a higher-level relationships between **concepts in background ontology** (right).

Causemos Views

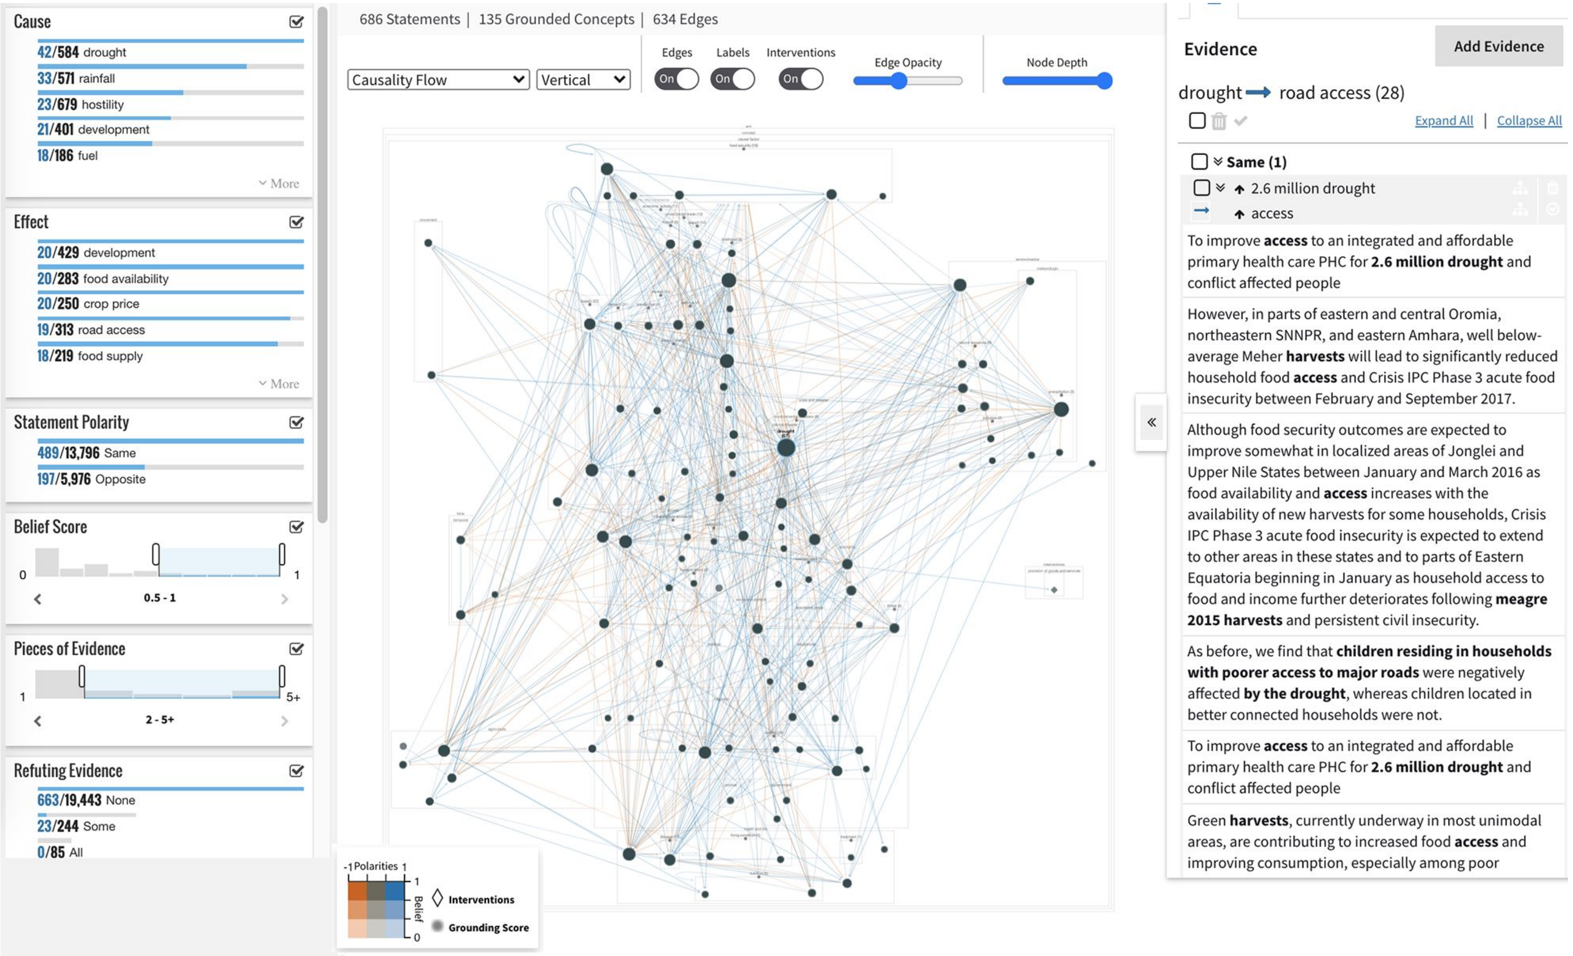

1. A section of the full **knowledge graph** in Causemos, with attribute facets on the left and evidence drill-down on the right.

2. An extracted knowledge **subgraph of interest** represented as a flow graph.

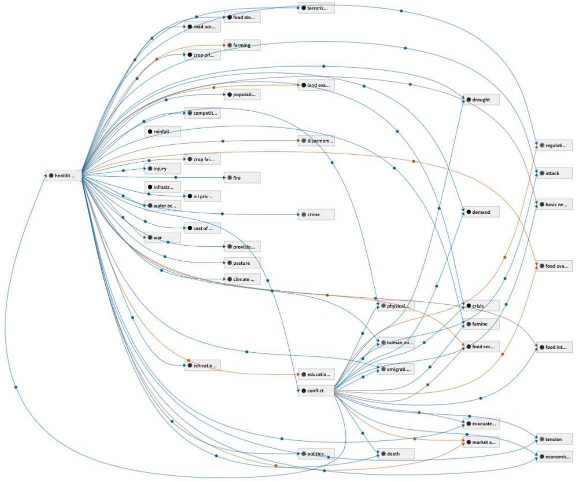

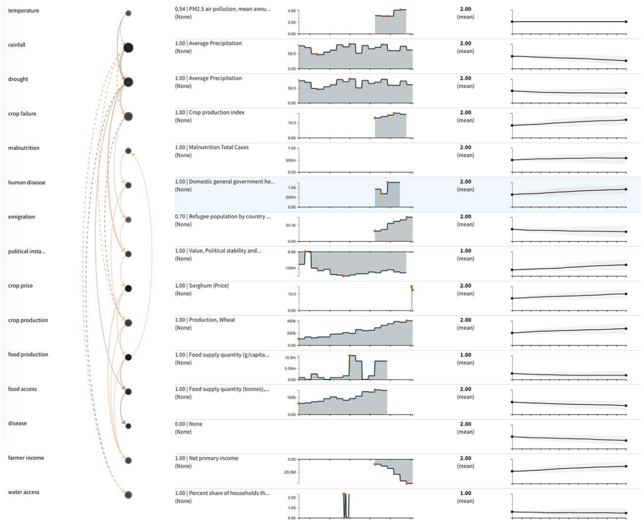

3. A **parameterized graphical model** (arc diagram), with each concept (left) grounded in a time-series (dark grey), with the model being projected forward in time (right)

see talk for video

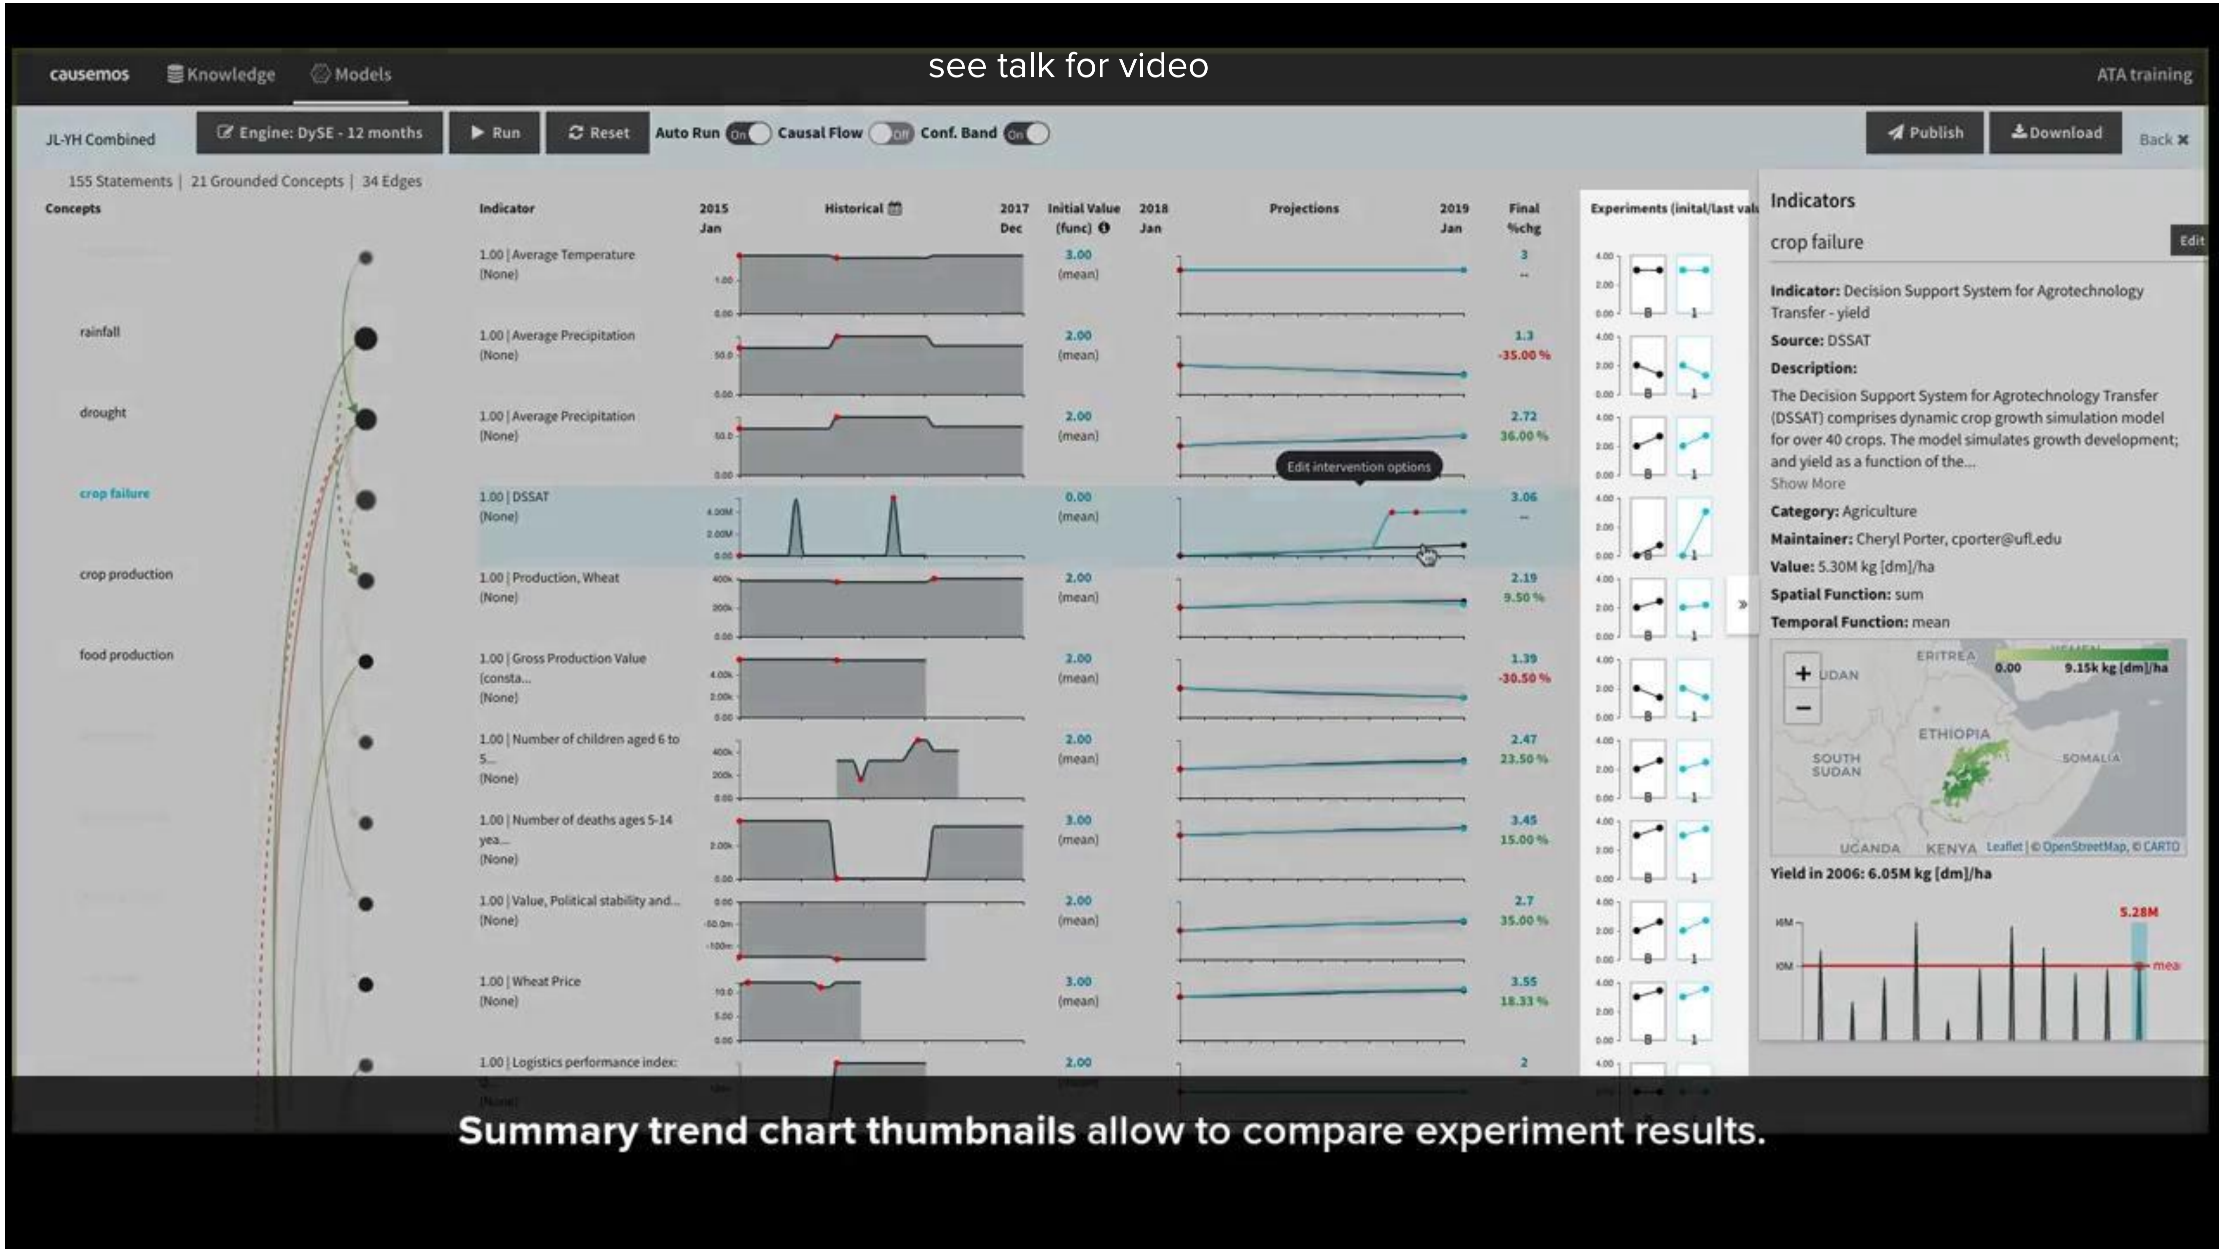

Data View

- Developed to support empirical questions of available structured or high-resolution map data.
- Analysis functionality for algebraic queries across outputs to see where conditions hold across spatial regions.
- Eventually want to use composite data to ground qualitative graphs.

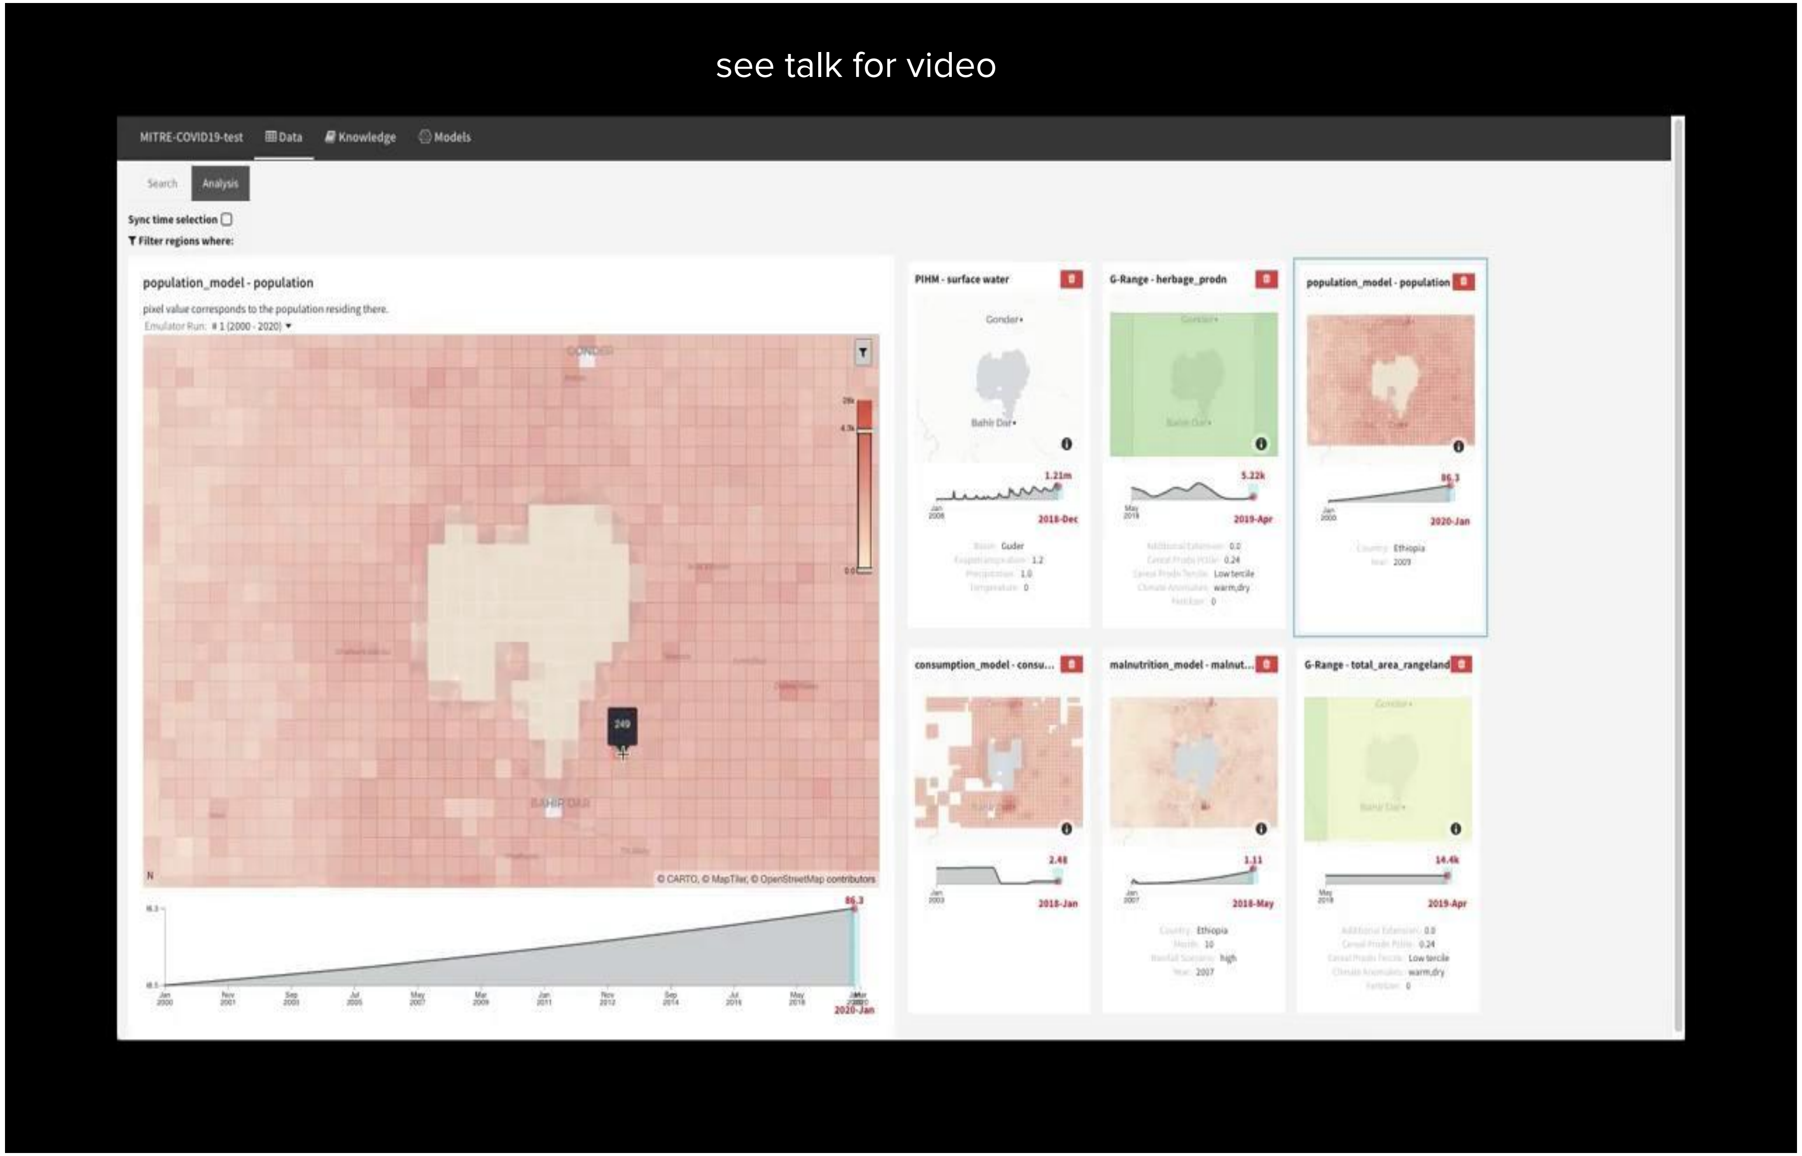

User Engagements

- Current use case around food security and linked issues in **Ethiopia**.

- Ethiopian stakeholders (academics, modelers, agricultural consultants) with support from the **Bill & Melinda Gates Foundation** and **Luma consulting**.

- Multiple successful user-focused experiments and system evaluations.

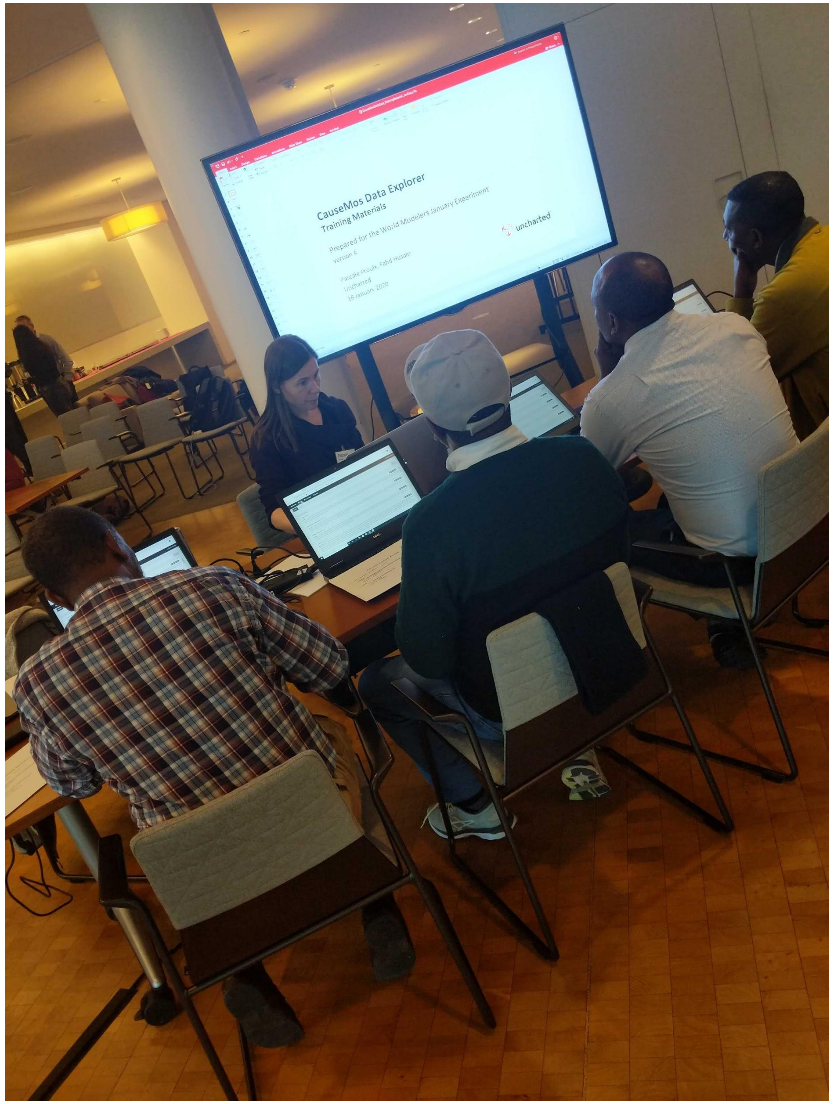

Future Causemos R&D

- Ongoing R&D to extend system to incorporate **user feedback** for next round of engagements, focusing in particular on iteration between qualitative and quantitative modeling workflows.

- Much **viz and ML** to be furthered, particularly in spatial analysis, large-scale graph visualization, model execution / validation, hierarchical clustering, mixed-initiative recommendation systems, and causal inference.

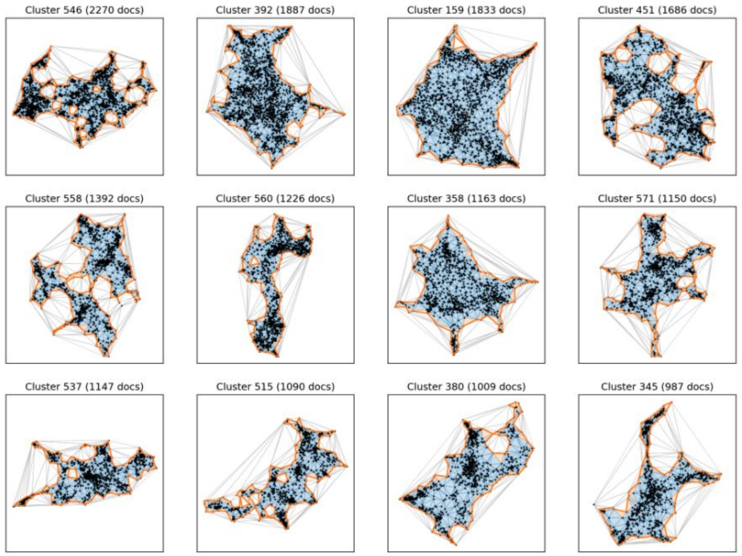

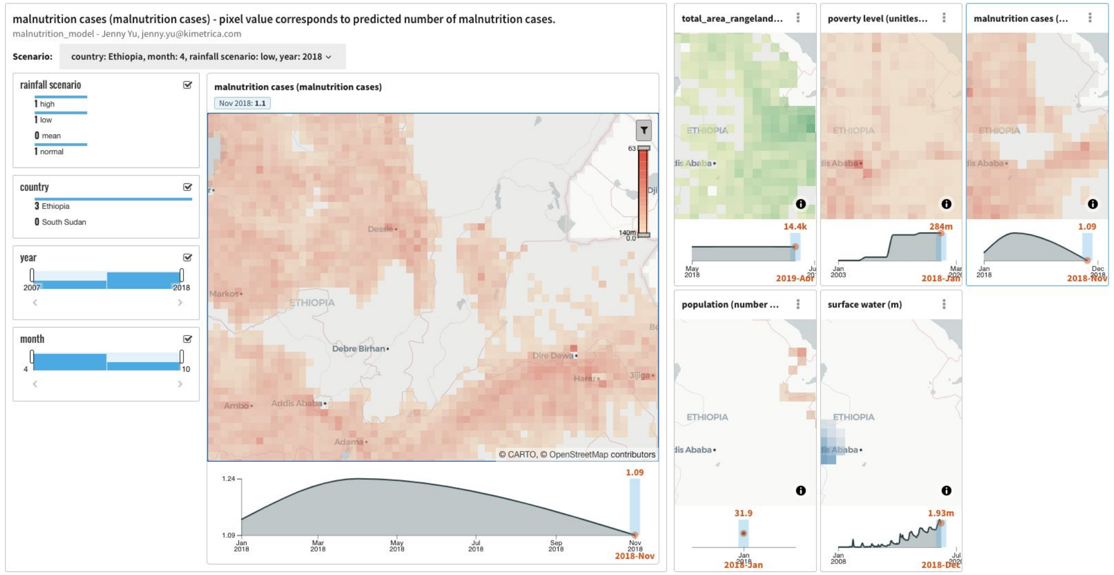

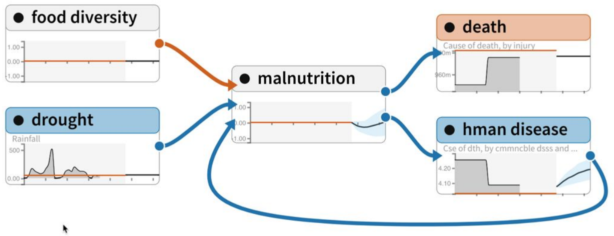

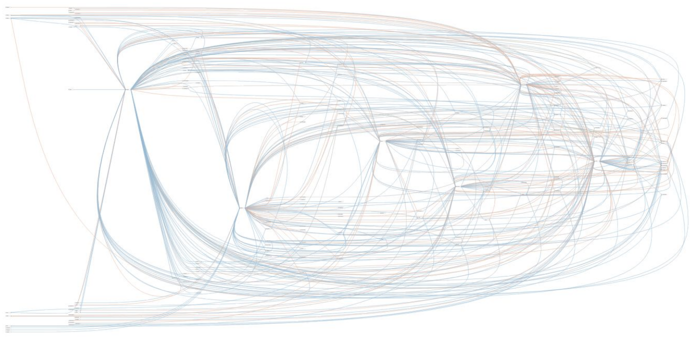

Automated Scientific Knowledge Extraction (ASKE)

Covid-19 Changed How the World Does Science, Together

Never before, scientists say, have so many of the world's researchers focused so urgently on a single topic. Nearly all other research has ground to a halt.

By Matt Apuzzo and David D. Kirkpatrick

Coronavirus exposes the problems and pitfalls of modelling

Models based on assumptions in the absence of data can be over-speculative and 'open to gross over-interpretation'

- Coronavirus - latest updates
- See all our coronavirus coverage

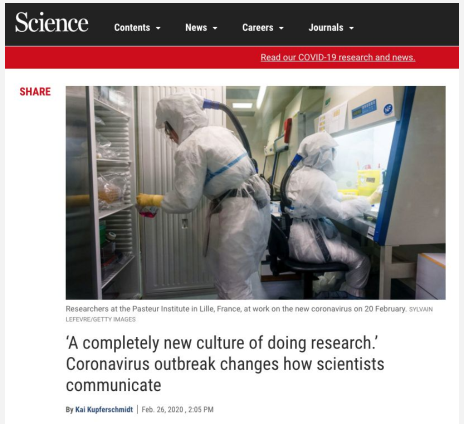

Science Contents News Careers Journals

Read our COVID-19 research and news.

SHARE

Researchers at the Pasteur Institute in Lille, France, at work on the new coronavirus on 20 February.

'A completely new culture of doing research.' Coronavirus outbreak changes how scientists communicate

By Kai Kupferschmidt | Feb. 26, 2020 , 2:05 PM

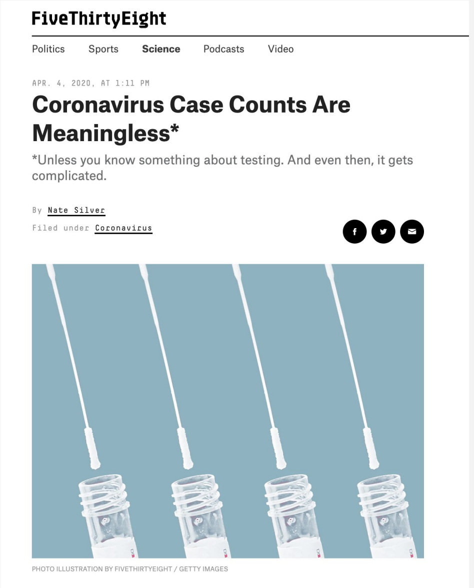

FiveThirtyEight

Politics Sports Science Podcasts Video

APR. 4, 2020, AT 1:11 PM

Coronavirus Case Counts Are Meaningless*

*Unless you know something about testing. And even then, it gets complicated.

By Nate Silver

Filed under Coronavirus

PHOTO ILLUSTRATION BY FIVETHIRTYEIGHT / GETTY IMAGES

- Explosion of models and scientific literature.
- Model representations diverse.
- Models difficult to understand, execute, augment.
- Lack of transparency and interpretability
- Model comparison rare (structural vs parametric).
- Uncertainty compounds across modeling workflow.
- Provenance and metadata critical for understanding.
- A tense stand-off between 'model fatigue' and 'immediate action'.
- All of the above amplified at scale.

[6]

Research Approach

- **Integrated visual modeling platform** for the exploration, augmentation and analysis of extracted scientific models.
- Promote **usability, transparency, interpretability** for diverse model artifacts in terms of representation, execution, augmentation and extension.
- Go from single-model to **multi-model capabilities** for reasoning *across* different model artifacts, enabling model comparison and ensembling.
- Platform to link model representations back to **source knowledge,** enabling large-scale exploration of background knowledge for contextualization and discoverability.
- Eventually moving towards visualization of **multi-resolution or hierarchical** model structures.
- Future directions: Development of **spatial data visualizations** for hyperlocal analysis.

Conceptual Schema

Data | Knowledge | **Models** | Experiments

- Initial focus on **models** and **knowledge**
- Articulate core **model representation(s)**
- **Contextualizing models** with background knowledge, both in terms of a structured knowledge representation (ontology / knowledge graph) and in terms of source knowledge (documents, code)
- **Model comparison**: juxtaposition, structural, parametric.
- Model execution, mutation and coupling longer-term.

Models as Graphical Models

Types

- Probabilistic graphical models
- Computational graphs
- Circuit diagrams
- Biological signal networks

Challenges

- Systems of systems of systems
- System features and metrics
- Hierarchies and complex nesting
- Various node, link and group attributes
- Heterogeneous semantics
- Nodes and links explode as system complexity grows

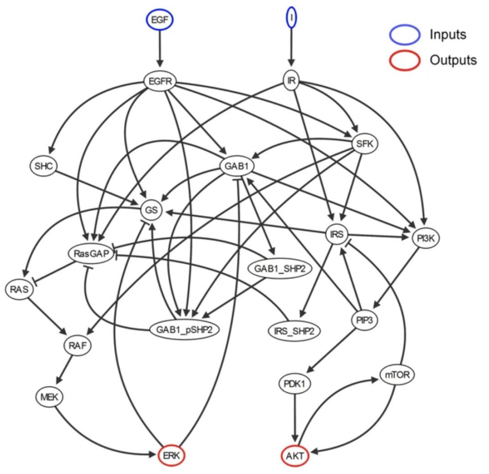

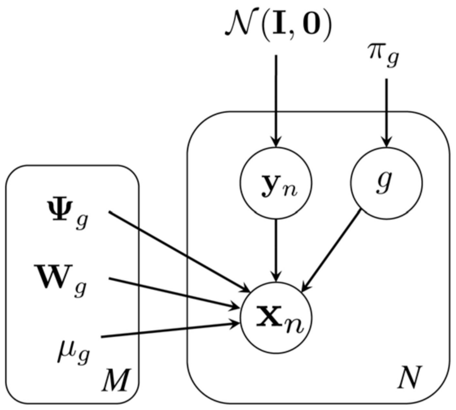

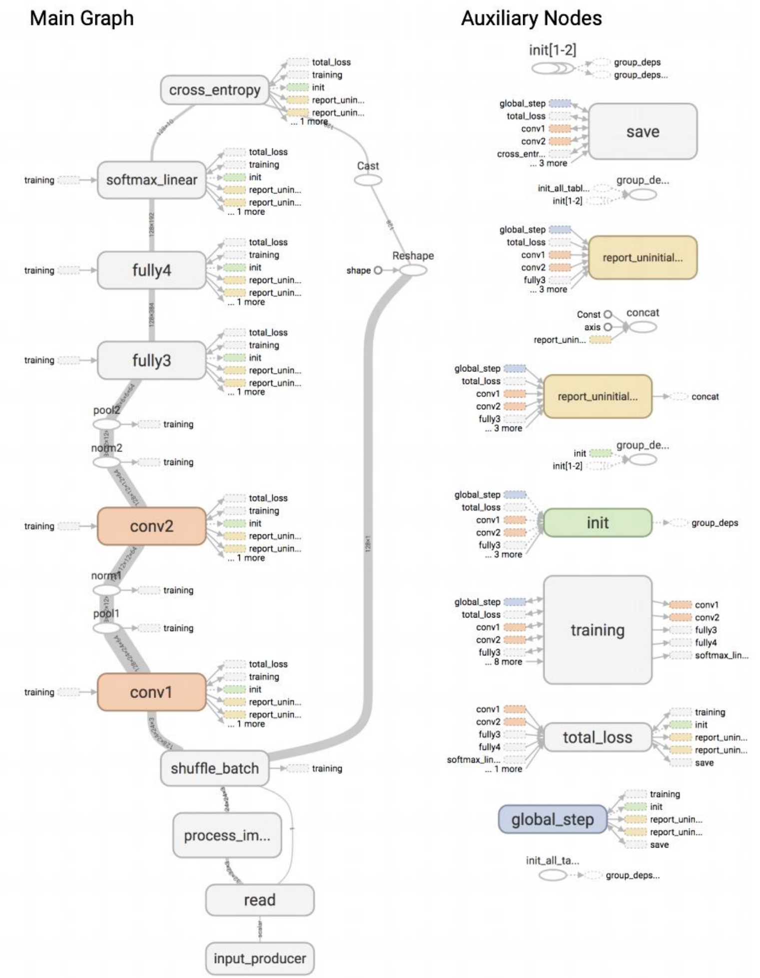

[19]

Enriched Model Representations

Start with **simple node-link representations**

- Simple attributes such as node values, link polarity, functions, global parameters (for execution).

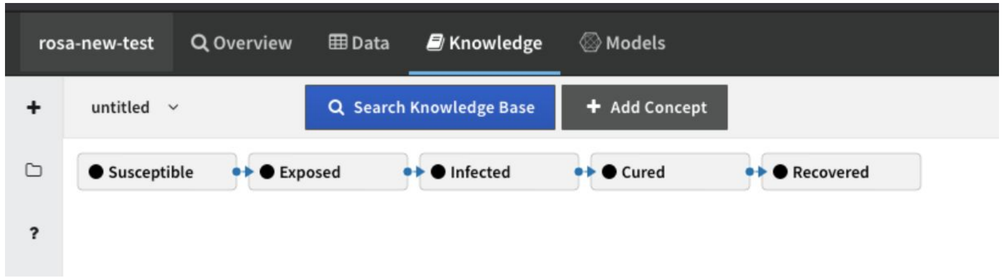

Move to **enriched graphical representations**

- Variety of node and link attributes
- Group structures: ontological categories, decomposable routines.
- Visualization of parameter surfaces, and corresponding outcomes and rewards.
- Linked visualizations connecting different model structures and components.
- Articulation of time? Evolving graph structures?

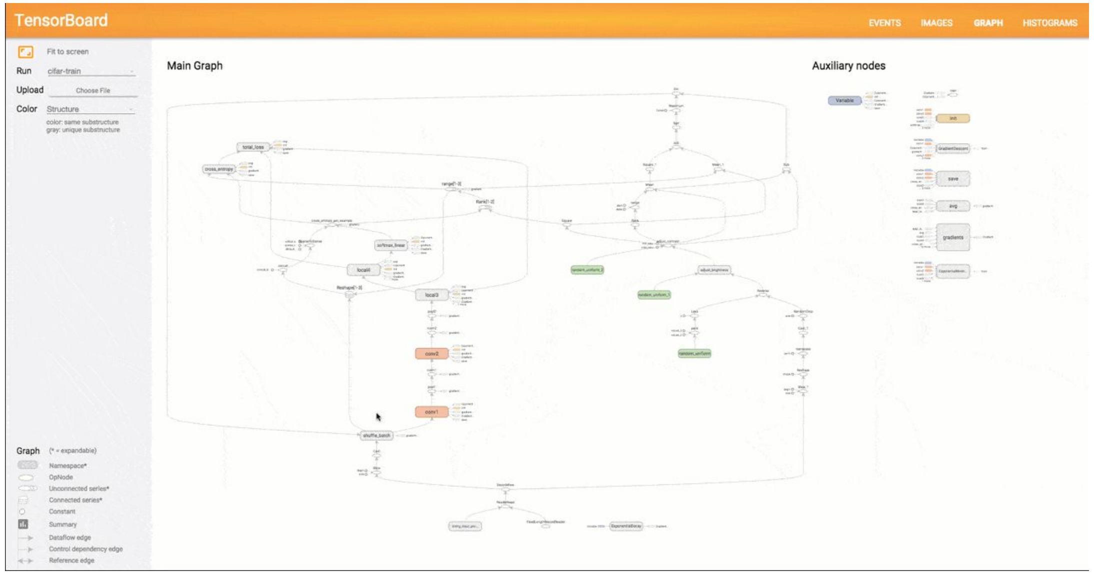

[19]

Single Model -> Multi-Model

Trajectory is to move from (extracted) single model representations to a (comparative) multi-model space.

Multiple versions of a single model:

- Juxtaposition of structural alternatives, parametric variations, multiple outcomes.
- Comparison of model augmentations and extensions.

Joint analysis of multiple models:

- Identifying converging and diverging predictions across models from different families.
- Exploring joint optimization strategies across various models.
- Working with model ensembles

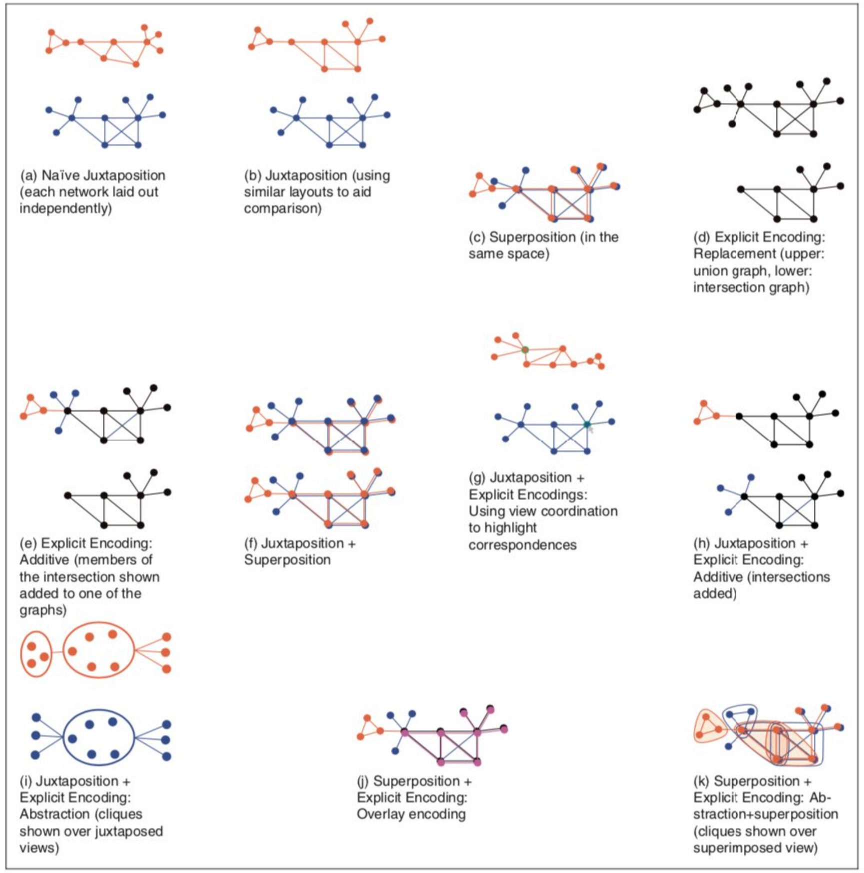

[9]

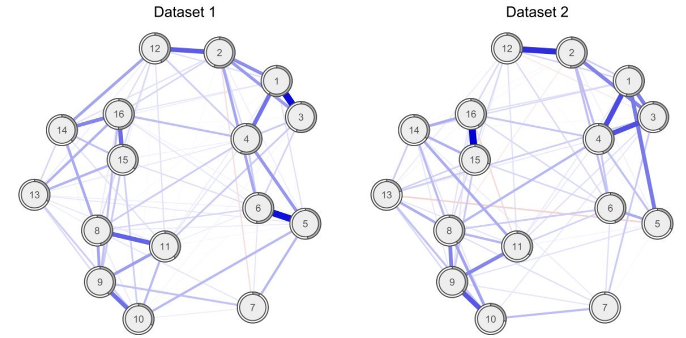

[8]

Extracted Model < > Background Knowledge

Model representations also need background knowledge representations.

These can take the form of model assumptions, model and parameter metadata, input data requirements, links back to source literature or code, juxtaposition of extracted model in larger knowledge graph…

Background knowledge exploration will be important for model contextualization, augmentation and extension of existing models, as well as discoverability of related models.

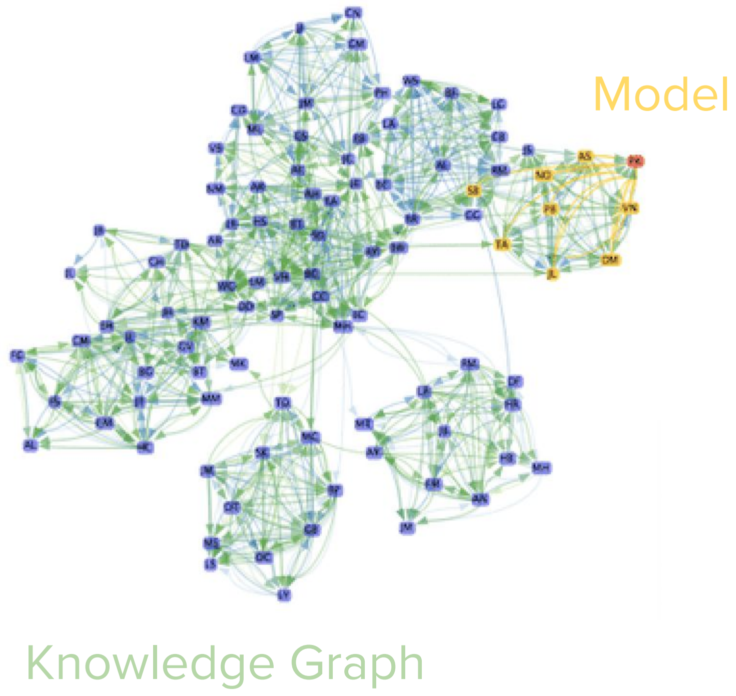

[7]

Hierarchical Models?

- Eventual vision to move towards visualizing interactive hierarchical model structures.
- **Single-model** case: model decomposable as nested sub-systems, where nodes can be unpacked to lower-level functional subgraphs or subsystems.
- **Multi-model, multi-resolution** case: stacked hierarchical layers of models from 'macro to micro' (from population-level models to biological models, for example).
- Various challenges:
  - Need full articulation of a single model as nested sub-systems, with interactions at lower levels propagating up in transparent and traceable ways.
  - Multi-model, multi-resolution case very complex, requiring a formal expression of hierarchy levels, correct placement of models at levels, establishing links between coarser and finer graph structures, articulating functional mappings (if any) between levels…

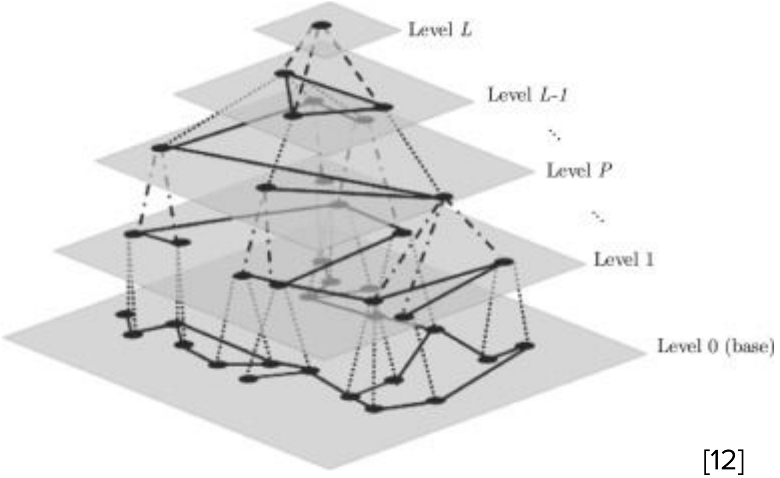

[12]

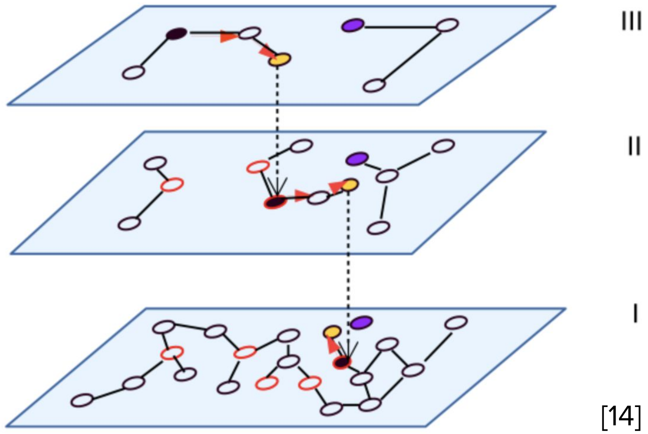

[14]

Initial Prototypes + Design

- Platform must cater to models at different levels: population-level epidemiological models and molecular-level biological models.
- Initial prototyping with different representations to understand semantics.
- Iterative exercises to build on design structures in light of potential upcoming research.
- Conceptual work to establish representation hierarchies, workflows tying together capabilities, etc.

see talk for video

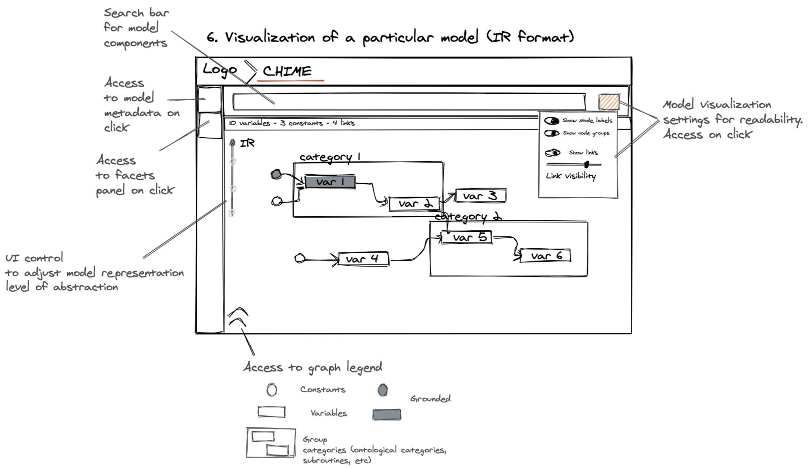

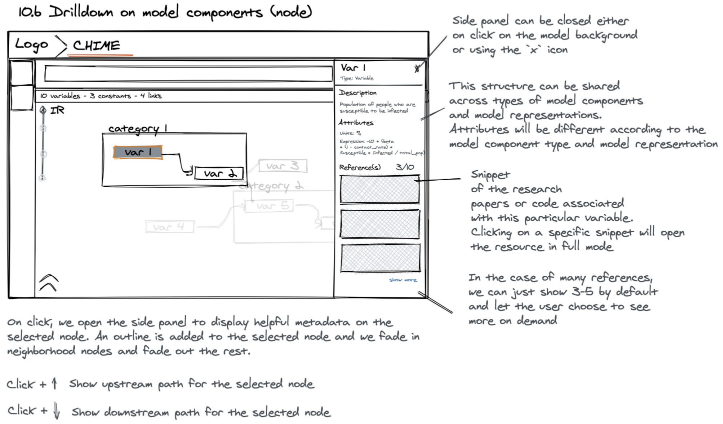

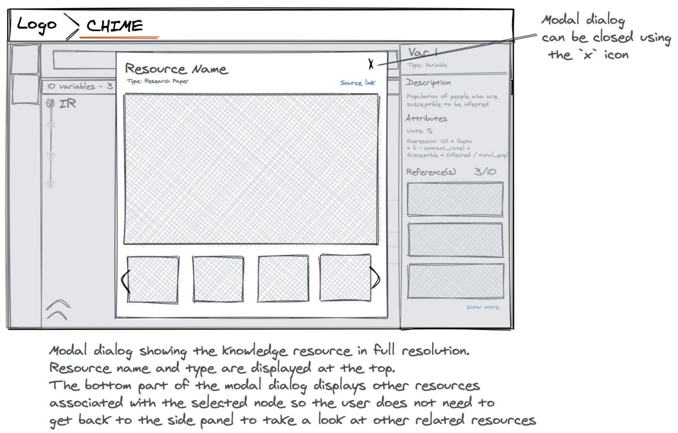

Scalable Graph Visualization Prototypes

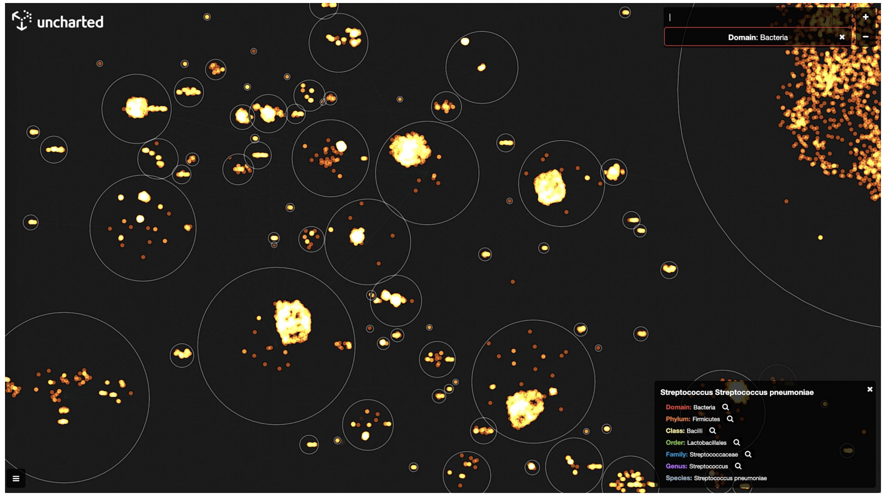

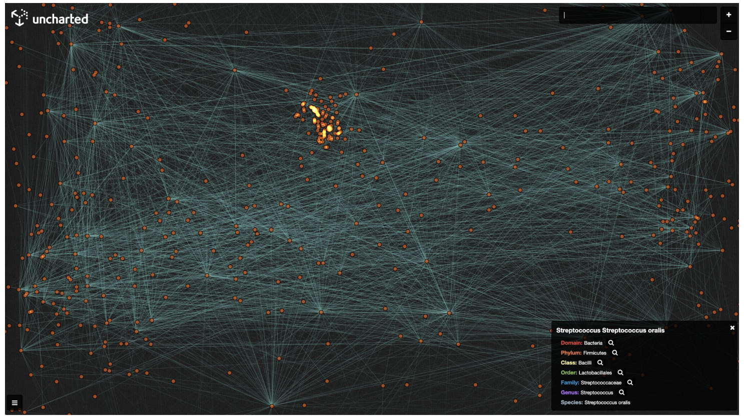

Genome data with 60k nodes / 60 mil edges. Hierarchical Louvain clustering used to allow interactive zoom to decompose coarse clusters into finer graphs.

see talk for video

Layout preserving 2D to 3D conversion and high performance rendering experiment for graph of thousands of nodes with coloured groupings.

III. Towards a Paradigm

Challenges (1/2)

Various challenges keep reappearing in these visual modeling research efforts:

- User resistance to change. Familiarity with existing workflows leads to bias towards the contours (and limitations) of existing technologies.
- Preconceived ideas and prior knowledge carried in by users, so there must be support for the insertion of user mental models into existing knowledge. This sort of flexibility is a modeling and engineering challenge.
- Different user roles such as model builders, model tweakers, model consumers. Different levels of expertise and background knowledge, which again requires design flexibility.
- Encoding metadata (descriptions, parameter interpretations, sensible output characteristics) of models or model outputs remains mostly a manual process. Real obstacle for users looking to understand knowledge.

Challenges (2/2)

- Even with models as graphs, common forms are hard to articulate. At the API level, node-link objects get complicated quickly with various exceptions. At the visual representation level, semantics often do not align, and interactions become inconsistent across models.
- Uncertainty compounds at each step of the workflow. Mechanisms are needed to track and trace the sources of uncertainty inherent to subjective workflows.
- Conceptual hierarchies are needed to provide aggregation for visual simplification and organization - but aggregations are often lossy, ill-fitting, lack proper coverage or granularity.
- Engineering a platform that spans complex interconnected services, with each evolving with a particular research agenda, is not trivial.

Visual Modeling: Representation

- Core representation is all models cast as graphical models of a few core types: node-link directed graphs, computational graphs, more complex flow structures like circuit / wiring diagrams.

- Core model components are nodes, edges, groups, or combinations thereof, with different types and semantics possible.

- Ongoing adjustment to maintain the abstraction sweet spot for visual representation as model types, components and functionality grows: need the right amount of generalization to span various models, but with a flexible structure that captures the critical particulars for each model.

Visual Modeling: Users & Workflows

- Visual modeling aims to enable a multi-domain, multi-model space. Modelers with expertise in a domain immediately become general modelers when immersed.

- Iterative, non-linear workflows are required, consistently moving between knowledge, models and data. Enable the user to jump between different spaces and stitch together different artifacts.

- As capabilities and ease of iteration in workflows increases, the space explored expands very quickly. Provenance of data, knowledge and modeling decisions - along with bookmarking, annotations and user management - all are needed very quickly.

Visual Modeling: ML

- User interactions full of implicit and explicit contextual clues. Machine-assisted guidance and machine-intelligence suggestion / recommendation services are essential, as are mechanisms for eliciting, capturing and incorporating user feedback.

- Interactive ML services for such a platform must be interpretable, lightweight, user-tweakable, and developed with active-learning in mind. Offline or back-end ML services can be deployed in more typical ways.

- Interpretation, transparency and debugging of models and output is critical for establishing user trust - work from current AI viz research can be leveraged here.

Visual Modeling: Linking & Abstraction

- Given the scale of the knowledge and data often involved, conceptual organization (an ontology, for e.g.) is critical for browsing knowledge and data artifacts at scale.

- Graph aggregations, hierarchical clustering, data summaries are all heavily used visual components. Principles of progressive disclosure at play.

- Frequent oscillation between background and foreground, between overview and drill-down, between high-level and granular. Linked views and linked artifacts are consistently utilized.

- Consistent, ever-present contextualization of model artifacts against knowledge, or knowledge artifacts against data, or data artifacts against models.

Visual Modeling: A Paradigm

- A platform aligned with a **paradigm of visual modeling** enables general modelers to move through various stages of a modeling workflow interactively, iteratively, visually.

- Capabilities include engaging or assembling diverse model representations from different domains; persistently have access to knowledge and data spaces for both context and separate analysis; allow for modeling and meta-modeling perspectives - *and* enable all this through the medium of visualization itself.

- **Core components** of this approach involve graphical modeling, knowledge contextualization, hierarchies/abstraction, user-in-the-loop, linked artifacts and non-linear workflows.

- Some **benefits** are speed of iteration, synthesis of information, general visual forms for broader engagement, able to engage different aspects of compositional problems.

- Thoroughly **interdisciplinary** paradigm: equal emphasis on conceptual architecture, design, user engagement, interaction, visualization, workflows, engineering *and* modeling.

- Lots of **research and synthesis** still to come to fully articulate the paradigm of visual modeling.

Thank you!

fhusain@uncharted.software

Uncharted

WM team:

Pascale Proulx (co-PI), Daniel Chang, Rosa Romero-Gómez, Adamo Carolli, Tom Choi, Anthony Kalsatos, Jaehwan Ryu, Azad Memon, Ben Codrington, Nelson Liu, Kevin Branigan.

ASKE team:

Rosa Romero-Gómez, Dario Segura, Luis Antunes, Adamo Carolli, Nelson Liu

Special thanks to Rosa Romero-Gómez for her thoughtful contributions to this seminar talk.

External Partners

Atlas AI

Galois

Harvard Medical School

Jataware

Kimetrica

University of Arizona

University of Florida

University of Pittsburgh

... and more.

Referenced Works

1. Amershi, Saleema, et al. "Modeltracker: Redesigning performance analysis tools for machine learning." Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems. 2015.
2. A Visual Exploration of Gaussian Processes, https://distill.pub/2019/visual-exploration-gaussian-processes/
3. Cashman, Dylan, et al. "A user‐based visual analytics workflow for exploratory model analysis." *Computer Graphics Forum*. Vol. 38. No. 3. 2019.
4. DARPA World Modelers, https://www.darpa.mil/program/world-modelers
5. DARPA World Modelers Reports, https://www.darpa.mil/attachments/WorldModelers%20v2.pdf
6. DARPA ASKE, https://www.darpa.mil/program/automating-scientific-knowledge-extraction
7. Federico, Paolo, et al. "A visual analytics approach to dynamic social networks." *Proceedings of the 11th International Conference on Knowledge Management and Knowledge Technologies*. 2011.
8. Fried, Eiko I., et al. "Replicability and generalizability of posttraumatic stress disorder (PTSD) networks: A cross-cultural multisite study of PTSD symptoms in four trauma patient samples." *Clinical Psychological Science* 6.3 (2018): 335-351.
9. Gleicher, Michael, et al. "Visual comparison for information visualization." *Information Visualization* 10.4 (2011): 289-309.
10. How to Use T-sne effectively, https://distill.pub/2016/misread-tsne/
11. Kahng, Minsuk, et al. "Gan lab: Understanding complex deep generative models using interactive visual experimentation." *IEEE transactions on visualization and computer graphics* 25.1 (2018): 1-11.
12. Mousavi, Seyedeh Fatemeh, et al. "Hierarchical graph embedding in vector space by graph pyramid." *Pattern Recognition* 61 (2017): 245-254.
13. Langevin, Scott, et al. "Distil: A Mixed-Initiative Model Discovery System for Subject Matter Experts." International Conference on Machine Learning AutoML Workshop. 2018.
14. Lin, Peng-Cheng, and Wan-Lei Zhao. "A comparative study on hierarchical navigable small world graphs." *arXiv preprint arXiv:1904.02077* (2019).
15. Strobelt, Hendrik, et al. "Lstmvis: A tool for visual analysis of hidden state dynamics in recurrent neural networks." *IEEE transactions on visualization and computer graphics* 24.1 (2017): 667-676. http://lstm.seas.harvard.edu/
16. The Beginner's Guide to Dimensionality Reduction, https://idyll.pub/post/visxai-dimensionality-reduction-1dbad0a67a092b007c526a45/
17. TensorFlow PlayGround, https://playground.tensorflow.org/#activation=tanh&batchSize=10&dataset=circle&regDataset=reg-plane&learningRate=0.03&regularizationRate=0&noise=0&networkShape=4,2&seed=0.21639&showTestData=false&discretize=false&percTrainData=50&x=true&y=true&xTimesY=false&xSquared=false&ySquared=false&cosX=false&sinX=false&cosY=false&sinY=false&collectStats=false&problem=classification&initZero=false&hideText=false
18. What-If tool, https://pair-code.github.io/what-if-tool/
19. Wongsuphasawat, Kanit, et al. "Visualizing dataflow graphs of deep learning models in tensorflow." *IEEE transactions on visualization and computer graphics* 24.1 (2017): 1-12.