|galois|

ORCA: Orchestrating Symbiotic Intelligence for Agile and Adaptable, Crisis Response Decision Making

Sam Cowger, Sourya Dey, Matt LeBeau, Ethan Lew, Panchapakesan Shyamshankar, Ted Hille, *Eric Davis**

ewd@galois.com

This work funded in part by:

DARPA:	DARPA-PA-18-02-AIE-FP-039	AFRL:	FA8650-22-C-6420
	HR001119S0083-SDCPS-FP-009		
	HR001118S0057-Office-Wide-BAA-FP-057		

Crisis Response and AI/ML: The Ideal

- Domain experts capable of agile, performable, and high confidence modeling in response to crisis.
- Leverage AI/ML and data science to make decisions on actions, policy, with continuous improvement and feedback to improve outcomes.

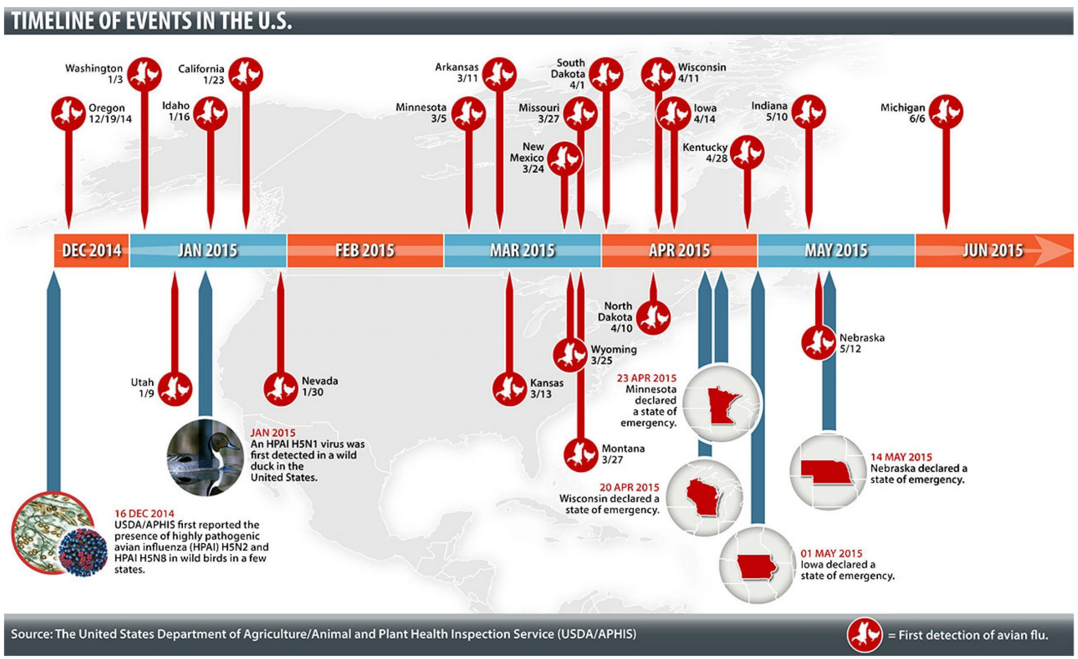

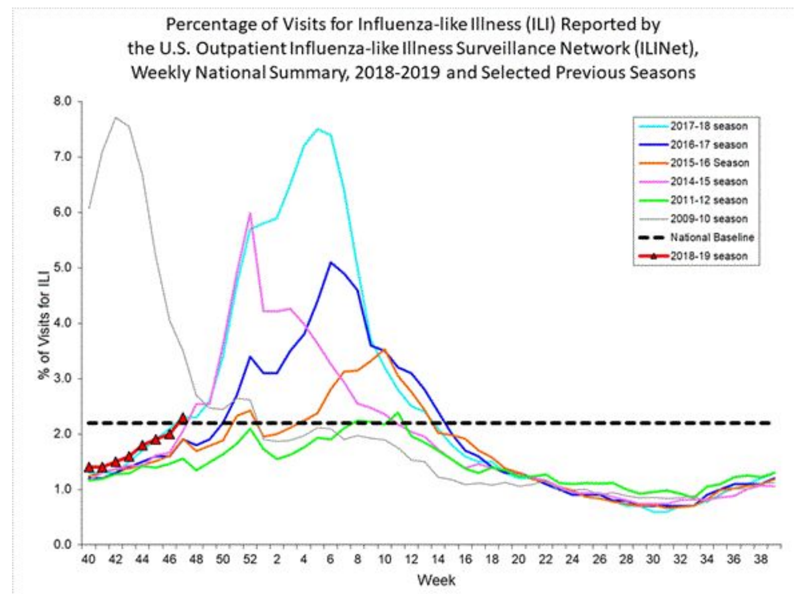

Crisis Response and AI/ML: The Reality

- Analysis of complex crises:

Requires large teams of modelers, mathematicians, and software engineers.

Not reusable.

Not maintainable.

Not easily developed.

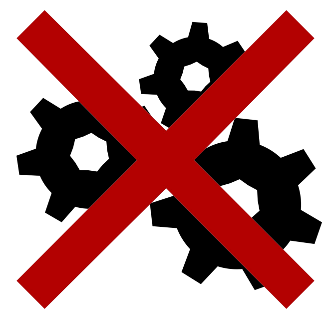

Not performable.

Initial Models

First Model-Based Workflow

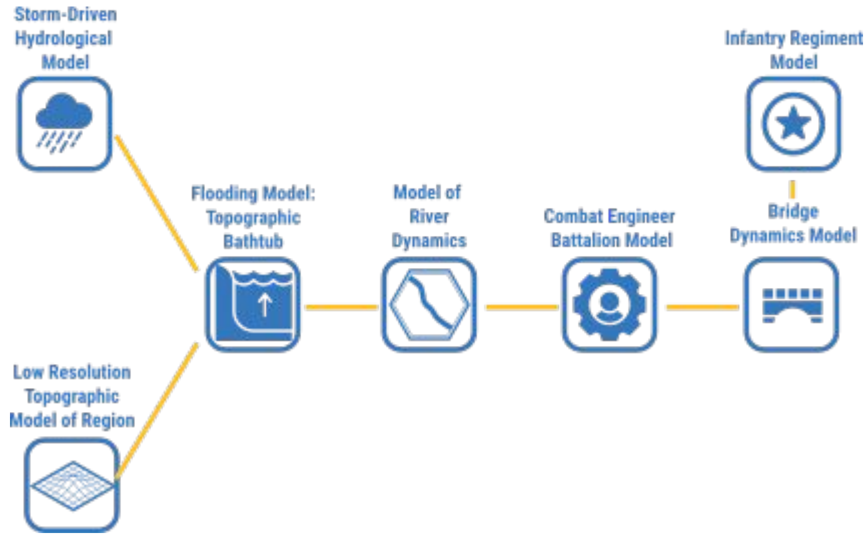

Adjustments Based on Other Stakeholders

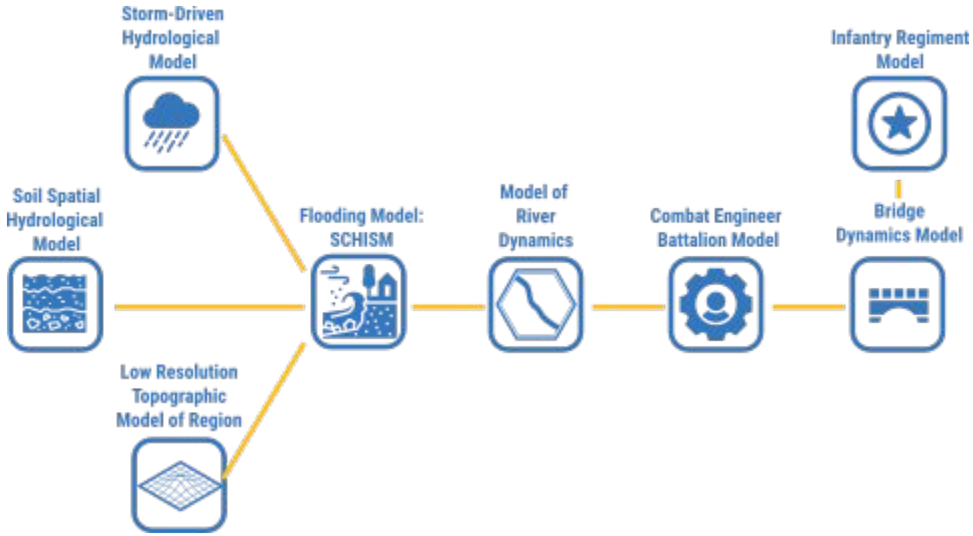

Improvements to Accuracy and Coverage

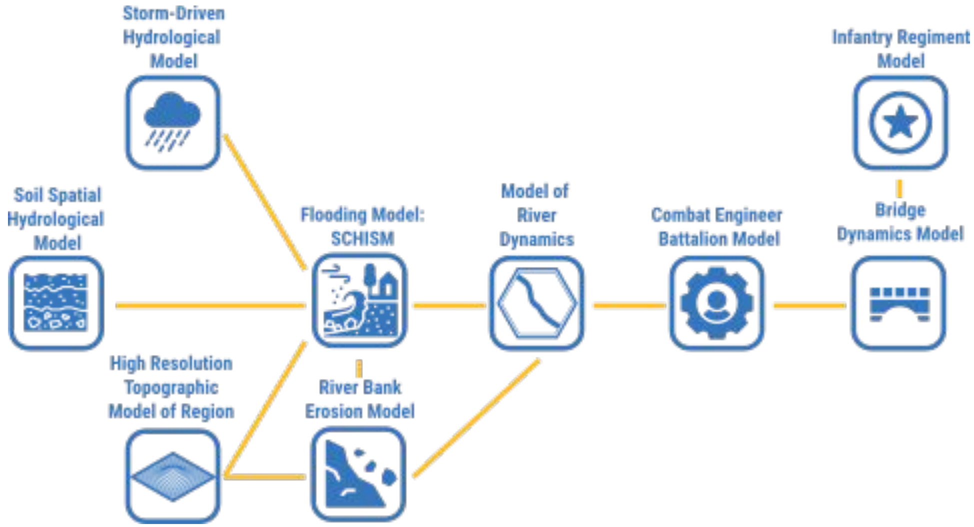

Extension to Other Domains

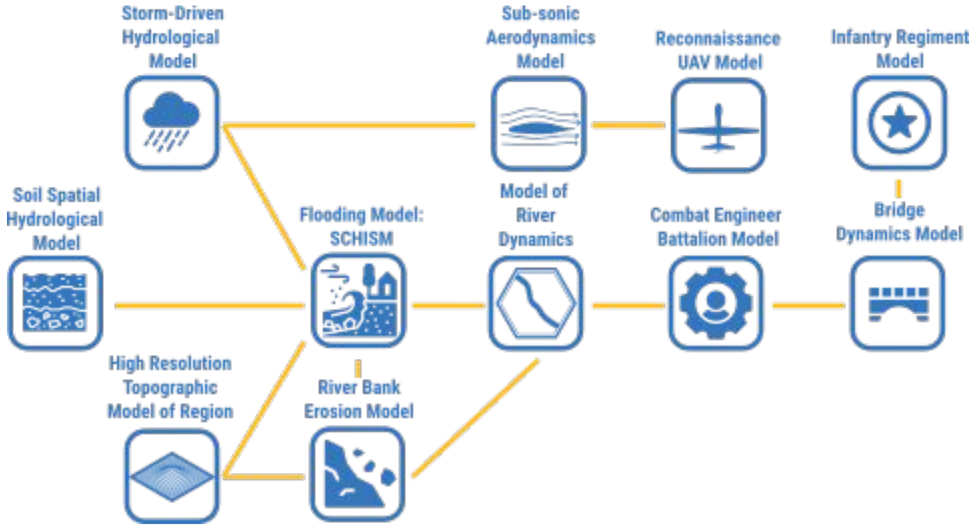

SME Computation Needs in Practice

Models of Varied Provenance

- Models come from varied producers.
- Varying quality and assumptions.

Initial Exploration

- Models evolve over time
- The need for **fidelity, believability, accuracy**, and **alignment to mission goals** increases over time.

Model Resolution and “Cost”

- Expensive/complex models may be necessary in the final step.
- Initial models during feasibility studies just need to prove out the workflow.
- Often room for improvement in a given domain before final analysis.

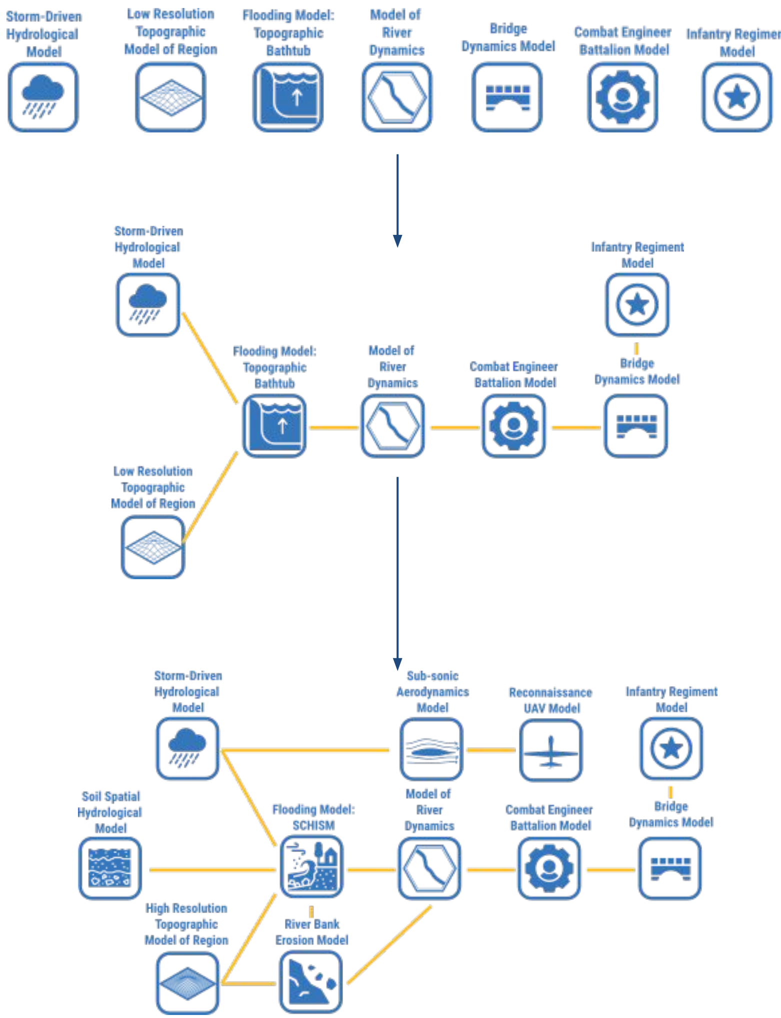

SME Centric Workflows

Knowledge Representations

Mission planning, analysis, and performance utilizes a variety of knowledge usually encoded in incompatible forms.

Human Artifacts: Papers, documentation, studies, reports.

Mathematical Artifacts: Formal specifications, models, statistical data.

Executable Artifacts: Code, executables, software.

UI/UX Artifacts: Dashboards, visualizations.

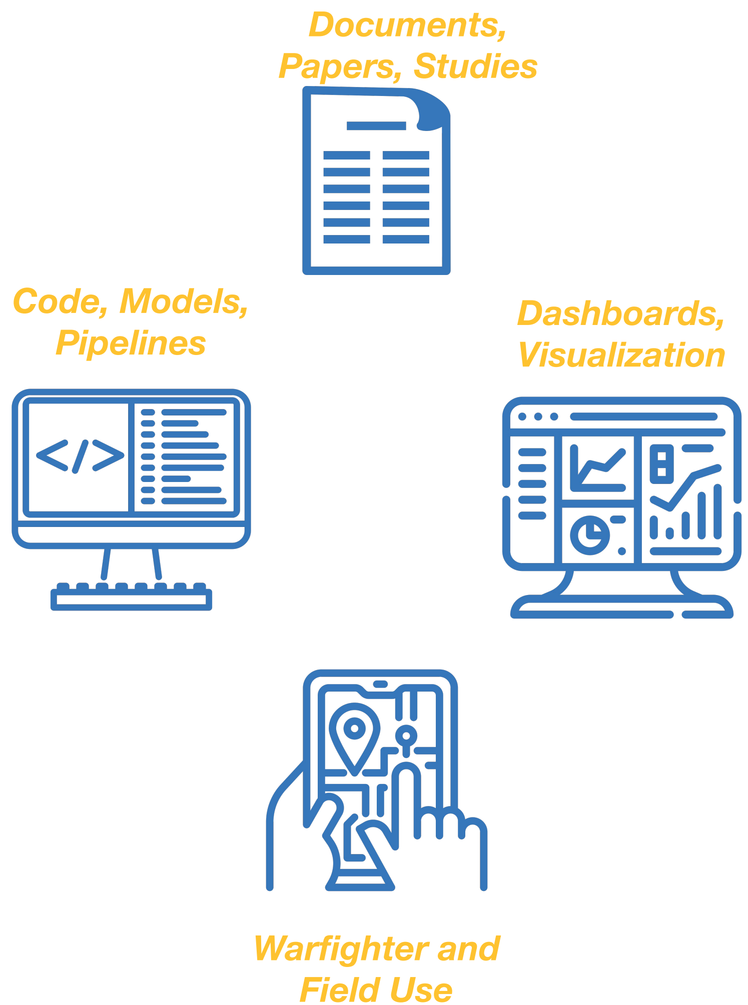

SME Centric Workflows

Knowledge Layers and Representations

Improve the ability of decision makers to **directly** interact with artifacts in forms they find **intuitive, usable,** and **explainable**.

Improve response times by having AI do the hard work of translation, and allowing SMEs to make changes, explore differences, and validate models and representations of individual processes.

Automated (Re)implementation

Avoid issues of being locked into legacy solutions, libraries, or code.

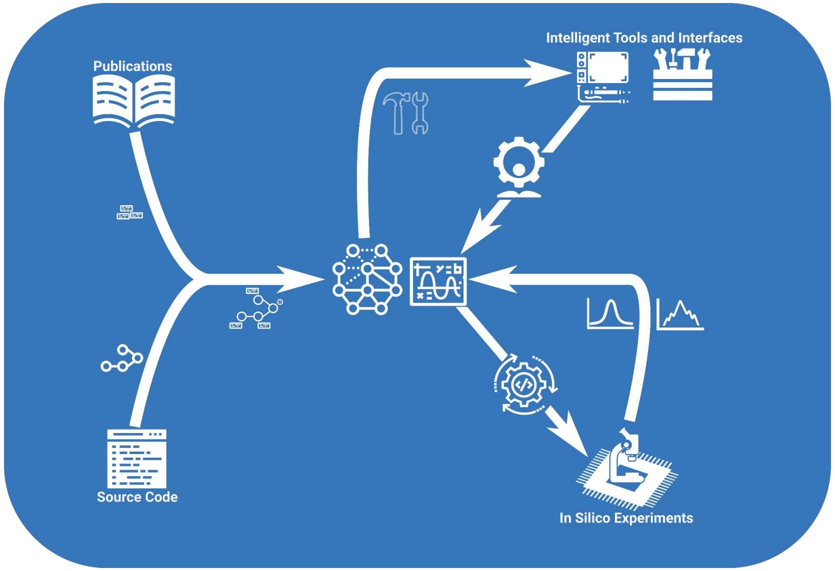

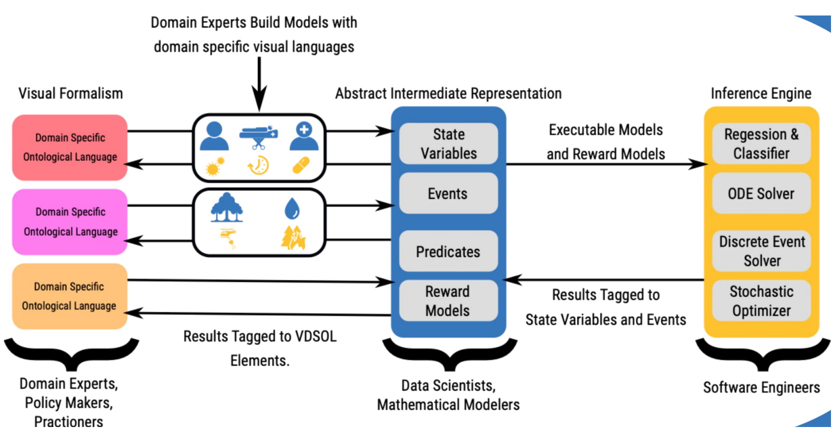

CDC models all forecast different futures

The CDC's model credibility problem is so large, it's a hard problem in and of itself.

Building an environment for continuous model iteration, testing, development, and improvement.

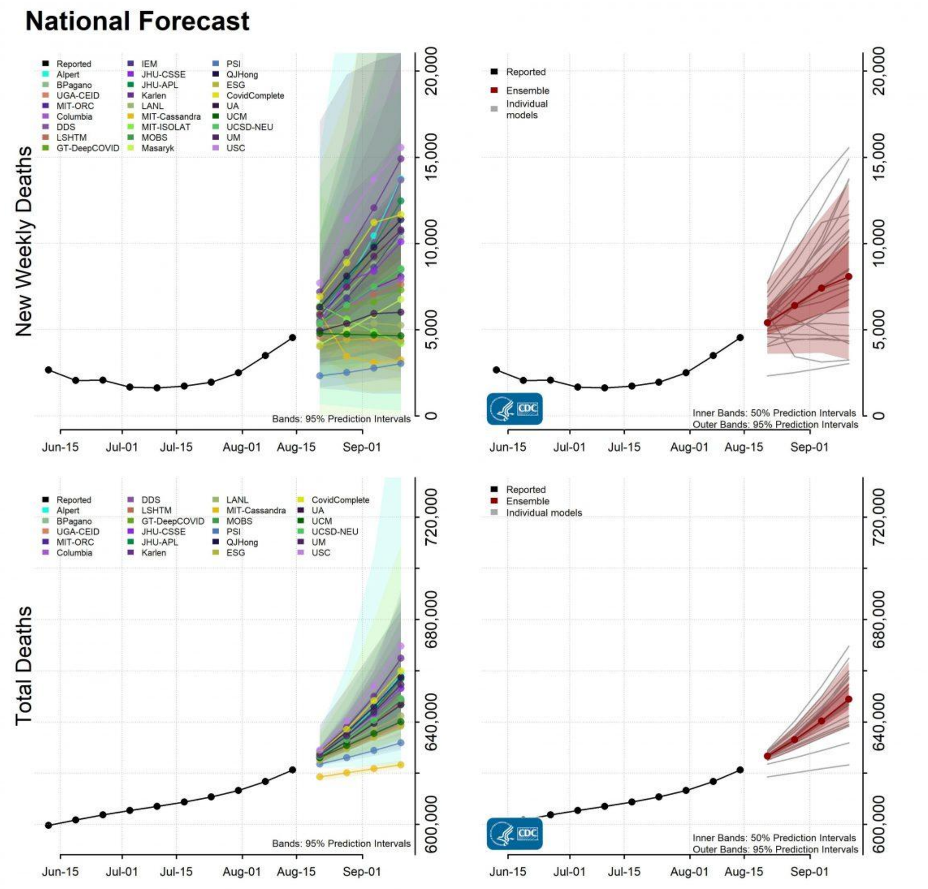

Case Study: COVID-19

Epidemic Calculator

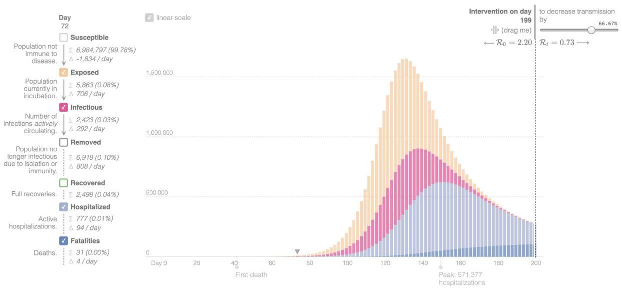

$$\frac{dS}{dt} = -\frac{\mathcal{R}_t}{T_{\text{inf}}} \cdot IS, \quad \frac{dE}{dt} = \frac{\mathcal{R}_t}{T_{\text{inf}}} \cdot IS - T_{\text{inc}}^{-1}E, \quad \frac{dI}{dt} = T_{\text{inc}}^{-1}E - T_{\text{inf}}^{-1}I, \quad \frac{dR}{dt} = T_{\text{inf}}^{-1}I$$

Gabriel Goh's SEIR

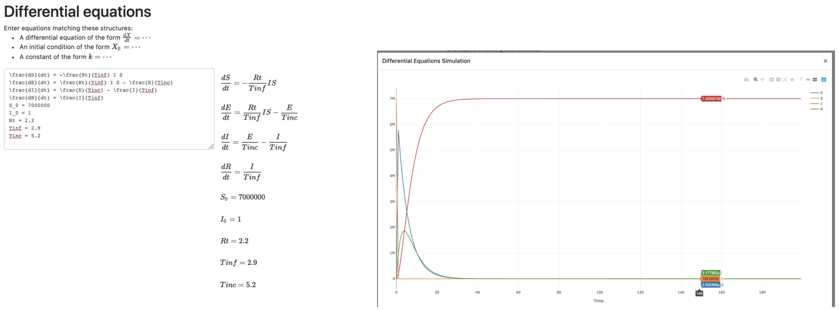

AMIDOL Synthesized Model doesn't match the reported results.

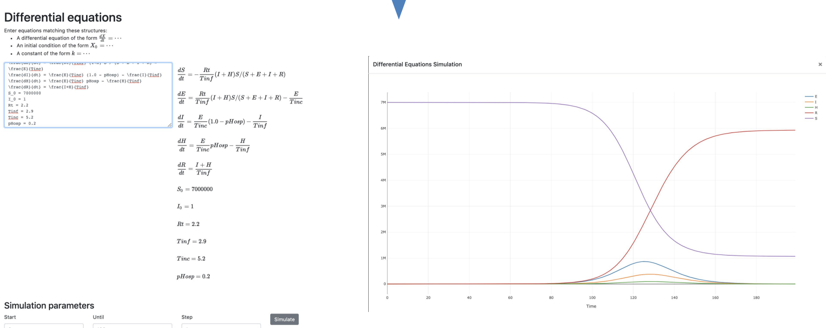

AMIDOL identifies missing term, synthesizes correct model which matches reported results.

Case Study:

Model of COVID-19 spread created by Gabriel Goh and popularized through the NYT and other media. As documented, the model has a flaw in the volumetric terms, leading to results that do not generalize as the total population of Goh's simulation was incorrectly composed with the generalized rate parameters.

As a result, attempts to replicate Goh's results failed. AMIDOL was able to automatically detect and correct the model., synthesizing the missing terms and producing a model which matched reported results.

$$\frac{dS}{dt} = -\frac{R_t}{T_{\text{inf}}} \cdot IS, \quad \frac{dE}{dt} = \frac{R_t}{T_{\text{inf}}} \cdot IS - T_{\text{inc}}^{-1}E \quad \frac{dI}{dt} = T_{\text{inc}}^{-1}E - T_{\text{inf}}^{-1}I, \quad \frac{dR}{dt} = T_{\text{inf}}^{-1}I$$

$$\frac{dS}{dt} = -\frac{R_t}{T_{\text{inf}}} \cdot \frac{IS}{\cdot S+E+I+R}, \quad \frac{dE}{dt} = \frac{R_t}{T_{\text{inf}}} \cdot \frac{IS}{\cdot S+E+I+R} - T_{\text{inc}}^{-1}E \quad \frac{dI}{dt} = T_{\text{inc}}^{-1}E - T_{\text{inf}}^{-1}I, \quad \frac{dR}{dt} = T_{\text{inf}}^{-1}I$$

Case Study: COVID-19

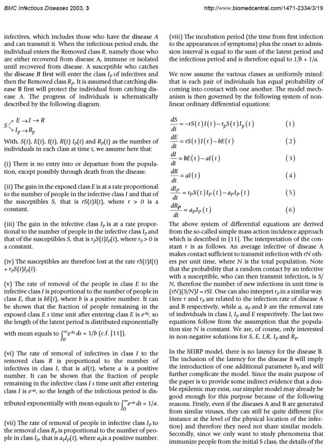
BMC Infectious Diseases 2003, 3

http://www.biomedcentral.com/1471-2334/3/19

infectives, which includes those who have the disease A and can transmit it. When the infectious period ends, the individual enters the Removed class R, namely those who are either recovered from disease A, immune or isolated until recovered from disease. A susceptible who catches the disease B first will enter the class I_P of infectives and then the Removed class R_P. It is assumed that catching disease B first will protect the individual from catching disease A. The progress of individuals is schematically described by the following diagram.

$$S \begin{matrix} \nearrow E \to I \to R \\ \searrow I_p \to R_p \end{matrix}$$

With, $S(t)$, $E(t)$, $I(t)$, $R(t)$ $I_P(t)$ and $R_P(t)$ as the number of individuals in each class at time t, we assume here that:

(i) There is no entry into or departure from the population, except possibly through death from the disease.

(ii) The gain in the exposed class E is at a rate proportional to the number of people in the infective class I and that of the susceptibles S, that is $rS(t)I(t)$, where $r > 0$ is a constant.

(iii) The gain in the infective class I_P is at a rate proportional to the number of people in the infective class I_P and that of the susceptibles S, that is $r_PS(t)I_P(t)$, where $r_P > 0$ is a constant.

(iv) The susceptibles are therefore lost at the rate $rS(t)I(t) + r_PS(t)I_P(t)$.

(v) The rate of removal of the people in class E to the infective class I is proportional to the number of people in class E, that is $bE(t)$, where b is a positive number. It can be shown that the fraction of people remaining in the exposed class E s time unit after entering class E is e^{-bs}, so the length of the latent period is distributed exponentially with mean equals to $\int_0^\infty e^{-bs}\,ds = 1/b$ (c.f. [11]).

(vi) The rate of removal of infectives in class I to the removed class R is proportional to the number of infectives in class I, that is $aI(t)$, where a is a positive number. It can be shown that the fraction of people remaining in the infective class I s time unit after entering class I is e^{-as}, so the length of the infectious period is distributed exponentially with mean equals to $\int_0^\infty e^{-as}\,ds = 1/a$.

(vii) The rate of removal of people in infective class I_P to the removal class R_P is proportional to the number of people in class I_P, that is $a_PI_P(t)$, where a_P is a positive number.

(viii) The incubation period (the time from first infection to the appearances of symptoms) plus the onset to admission interval is equal to the sum of the latent period and the infectious period and is therefore equal to $1/b + 1/a$.

We now assume the various classes as uniformly mixed: that is each pair of individuals has equal probability of coming into contact with one another. The model mechanism is then governed by the following system of nonlinear ordinary differential equations:

$$\frac{dS}{dt} = -rS(t)I(t) - r_pS(t)I_p(t) \quad (1)$$
$$\frac{dE}{dt} = rS(t)I(t) - bE(t) \quad (2)$$
$$\frac{dI}{dt} = bE(t) - aI(t) \quad (3)$$
$$\frac{dR}{dt} = aI(t) \quad (4)$$
$$\frac{dI_P}{dt} = r_PS(t)I_P(t) - a_PI_P(t) \quad (5)$$
$$\frac{dR_P}{dt} = a_PI_P(t) \quad (6)$$

The above system of differential equations are derived from the so-called simple mass action incidence approach which is described in [11]. The interpretation of the constant r is as follows. An average infective of disease A makes contact sufficient to transmit infection with rN others per unit time, where N is the total population. Note that the probability that a random contact by an infective with a susceptible, who can then transmit infection, is S/N, therefore the number of new infections in unit time is $(rN)(S/N)I = rSI$. One can also interpret r_P in a similar way. Here r and r_P are related to the infection rate of disease A and B respectively, while a, a_P and b are the removal rate of individuals in class I, I_P and E respectively. The last two equations follow from the assumption that the population size N is constant. We are, of course, only interested in non-negative solutions for S, E, I,R, I_P and R_P.

In the SEIRP model, there is no latency for the disease B. The inclusion of the latency for the disease B will imply the introduction of one additional parameter b_P and will further complicate the model. Since the main purpose of the paper is to provide some indirect evidence that a double epidemic may exist, our simpler model may already be good enough for this purpose because of the following reasons. Firstly, even if the diseases A and B are generated from similar viruses, they can still be quite different (for instance at the level of the physical location of the infection) and therefore they need not share similar models. Secondly, since we only want to study phenomena that immunize people from the initial S class, the details of the

$$\frac{dS}{dt} = -r(SI)/N - rp(SIp)/N$$
$$\frac{dE}{dt} = r(SI)/N - \beta E$$
$$\frac{dI}{dt} = \beta E - \alpha I$$
$$\frac{dR}{dt} = \alpha I$$
$$\frac{dIp}{dt} = rp(SIp)/N - \sigma Ip$$
$$\frac{dRp}{dt} = \sigma Ip$$

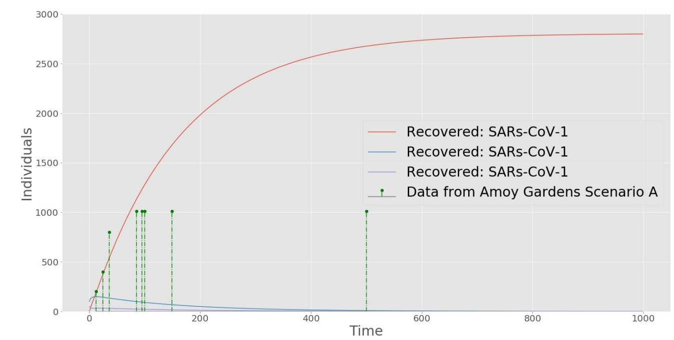

$$Ip_0 = 500000$$
$$S_0 = 6300000$$
$$E_0 = 100$$
$$I_0 = 50$$
$$N = 6800150$$
$$r = 0.6929$$
$$rp = 0.4808$$
$$\alpha = 0.47$$
$$\sigma = 0.461$$
$$\beta = 0.103$$

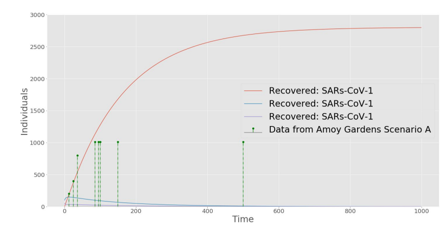

$$r = 0.6929$$
$$rp = 0.4808$$
$$\alpha = 0.3359$$
$$\sigma = 0.2402$$
$$\beta = 0.3015$$

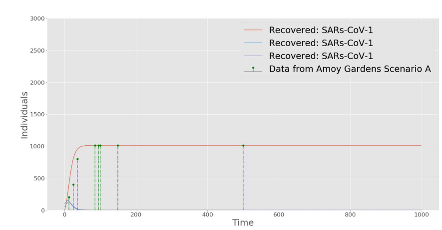

Case Study:

Double epidemic SEIRP model of SARS-CoV-1, which was influential in the development of policy and the initial study of COVID-19. First model presented in the paper for the Hong Kong Amoy Gardens Scenario A contained issues with the published model.

AMIDOL was able to determine the problem resided in three reported parameters: alpha, sigma, and beta.

AMIDOL detected and corrected the issue, yielding a corrected implementation, and specification.

Building Standardized Model “Containers”

Computational workflows can be made more interchangeable.

- Allow for model comparison
- Allow for model refinement
- Allow for alternative paths to a given computation.

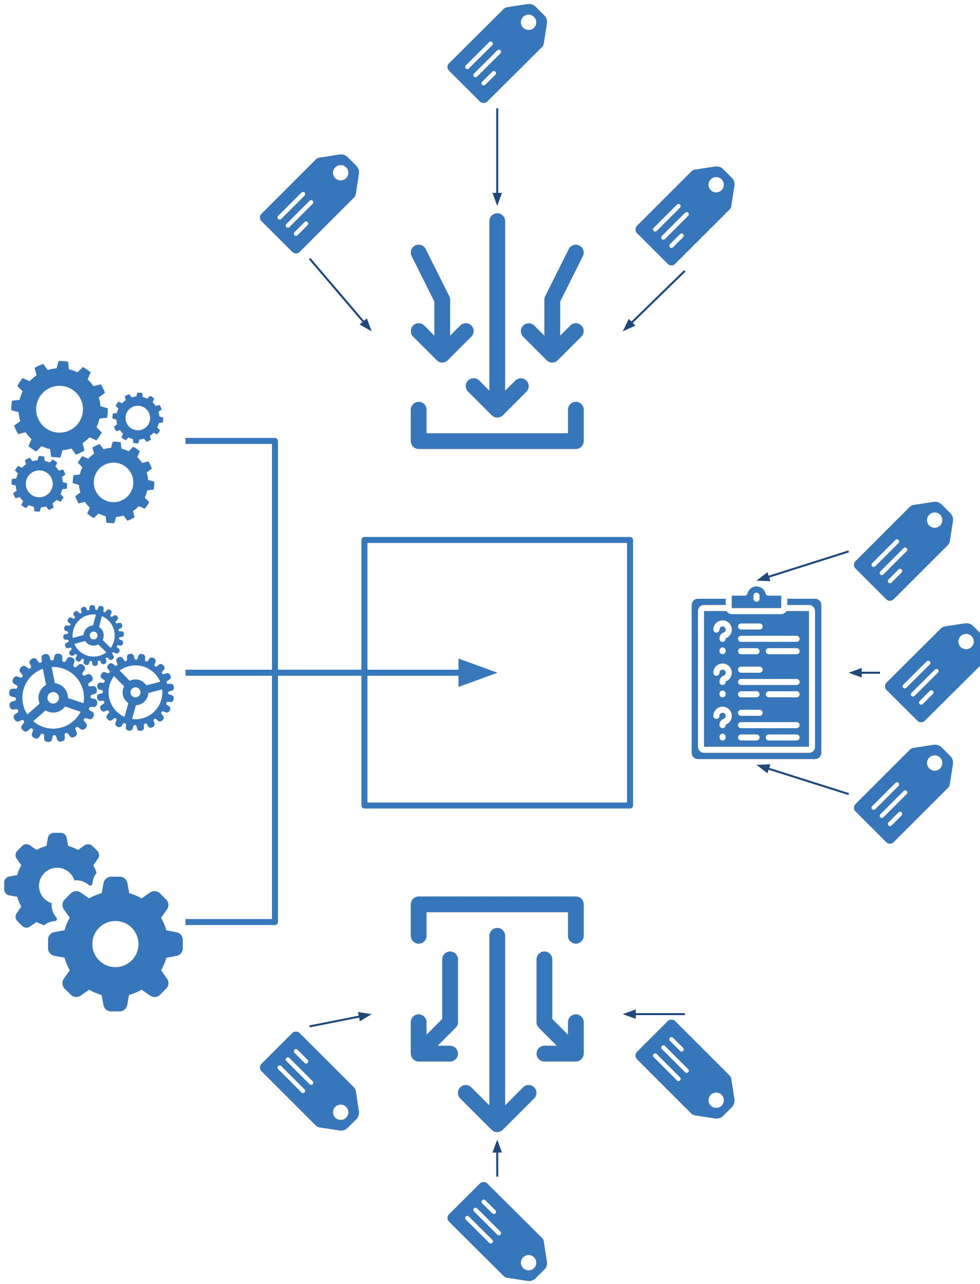

Building Standardized Model "Containers"

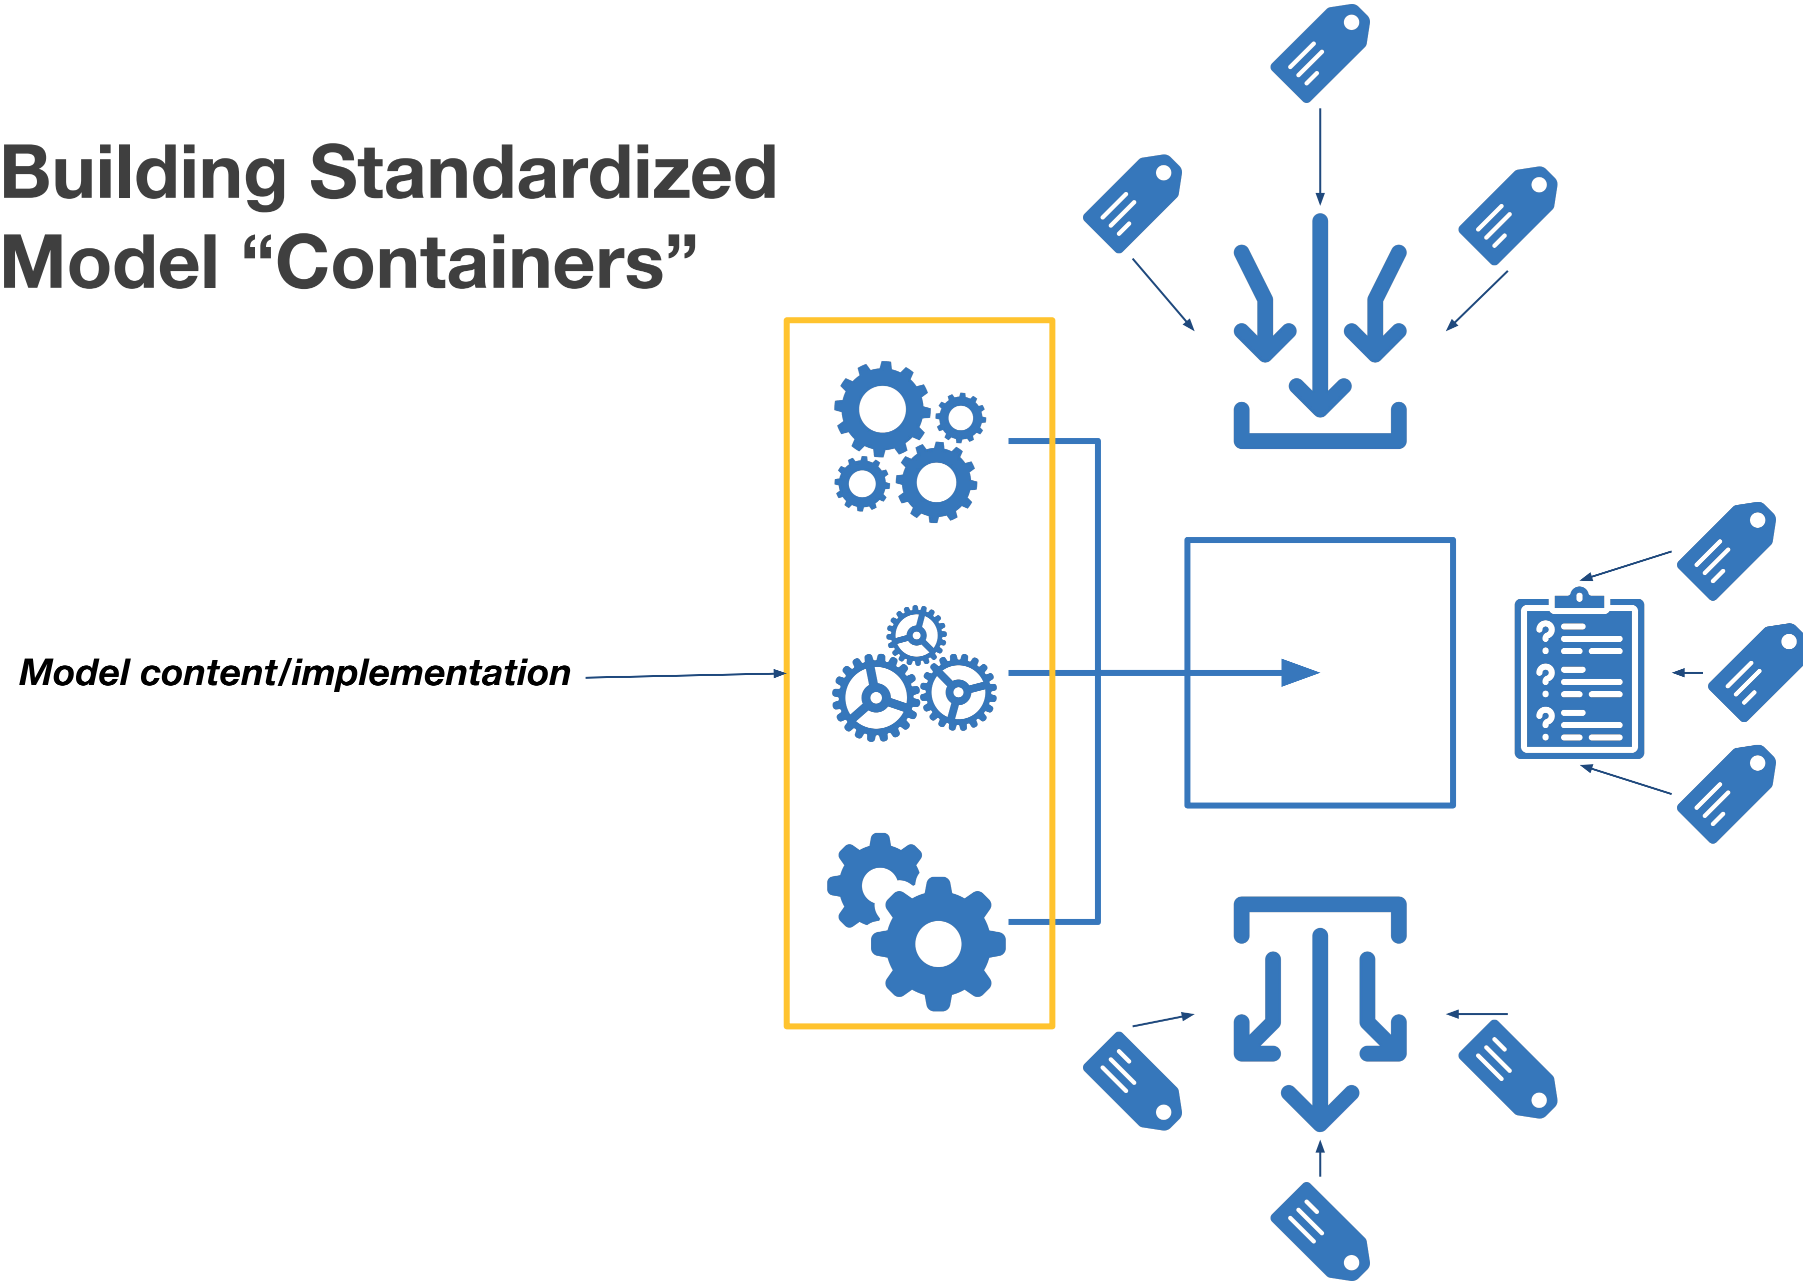

Building Standardized Model "Containers"

Input and Output signatures, types, and semantic tags

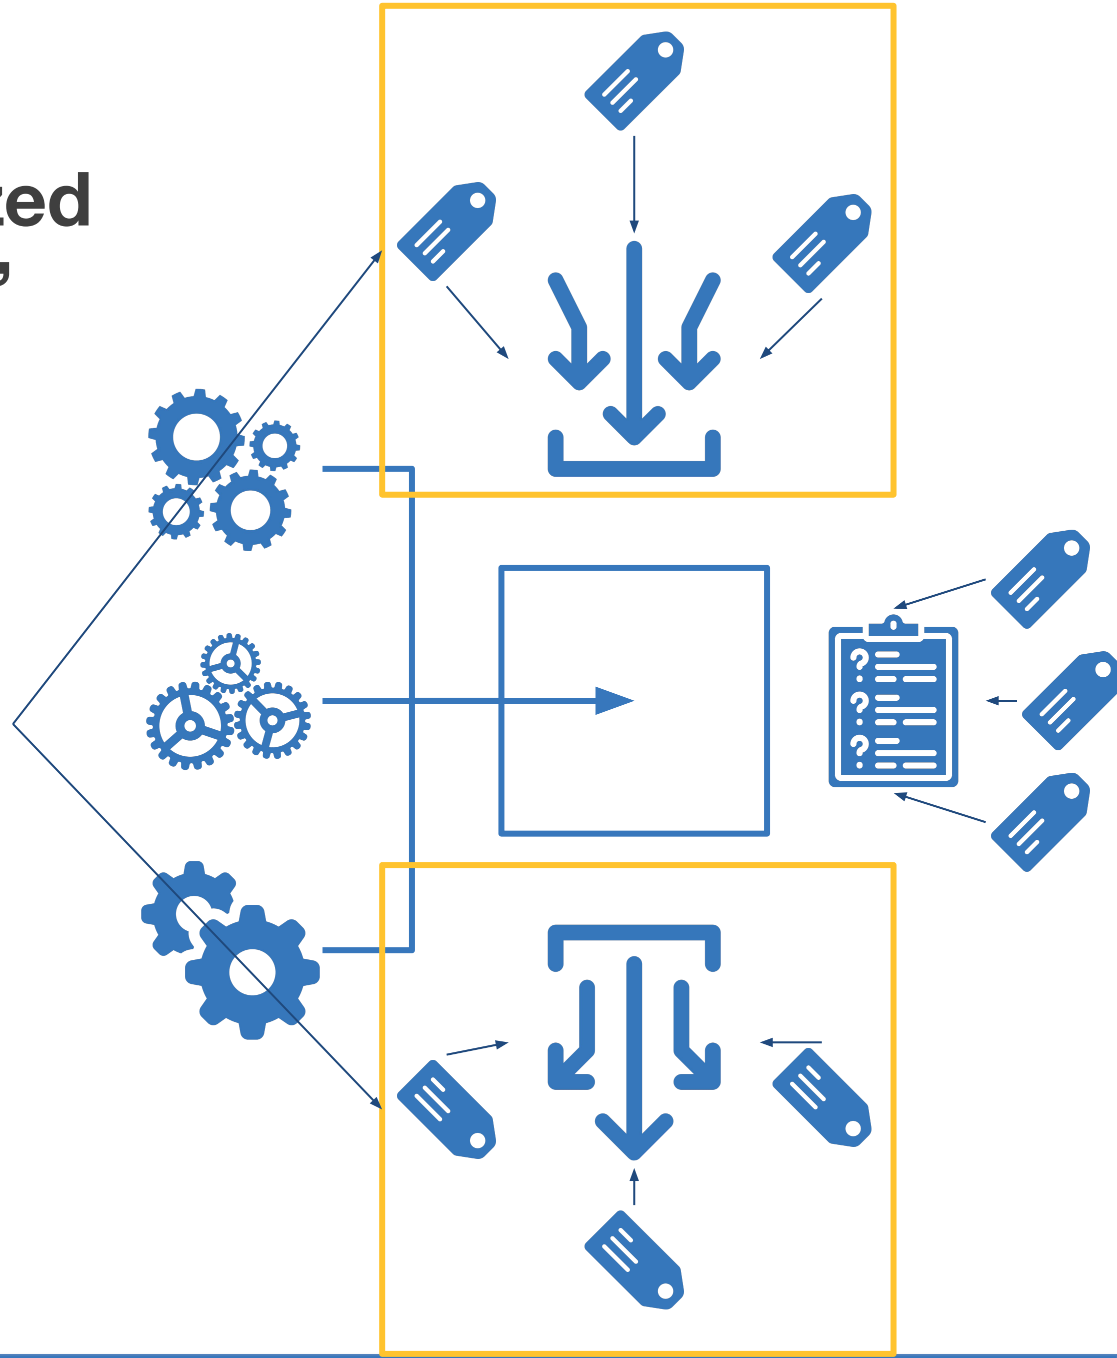

Building Standardized Model "Containers"

Model requirements, constraints, and semantic tags

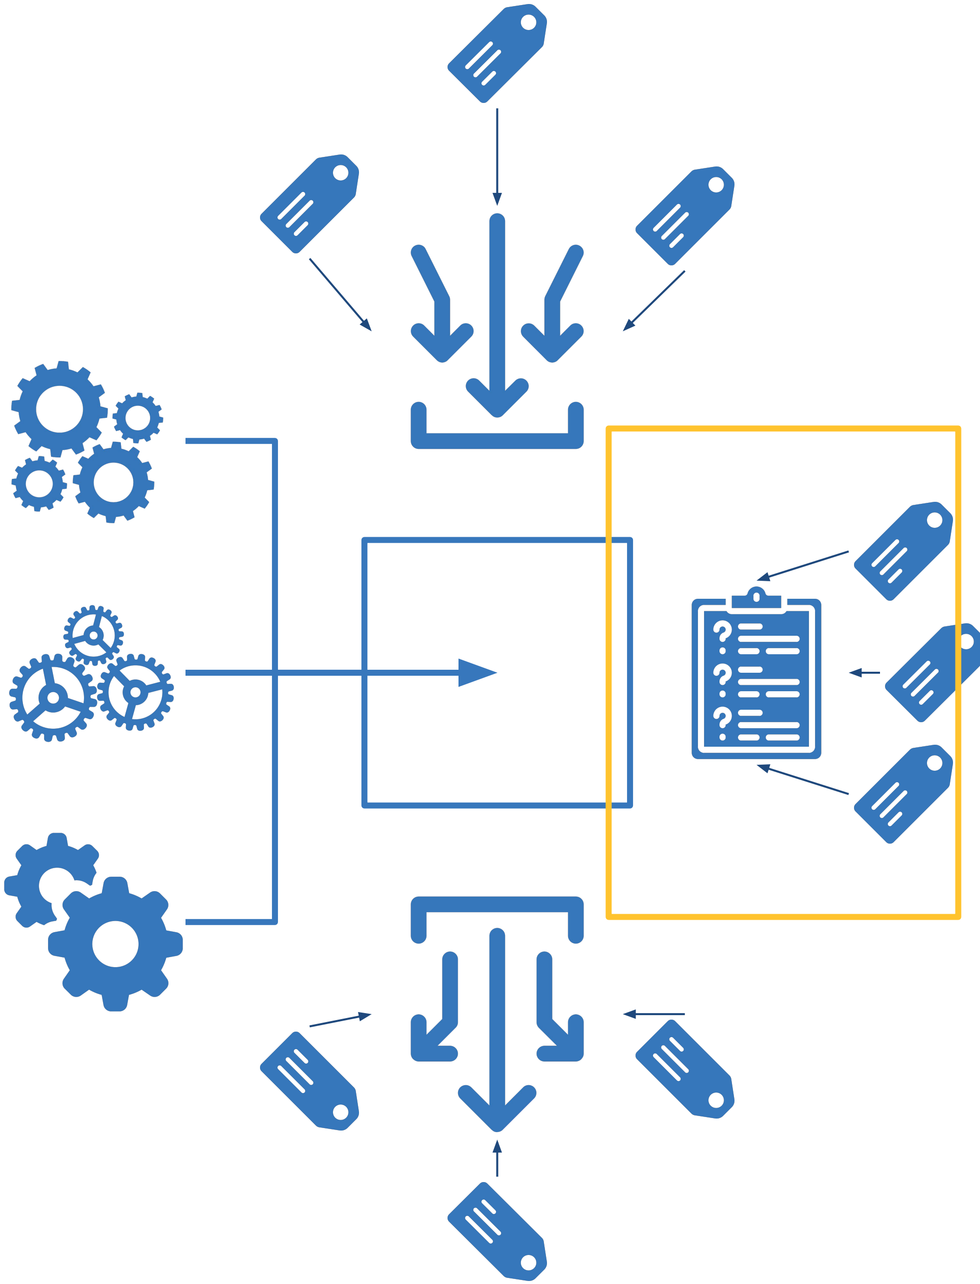

Current Pipelines

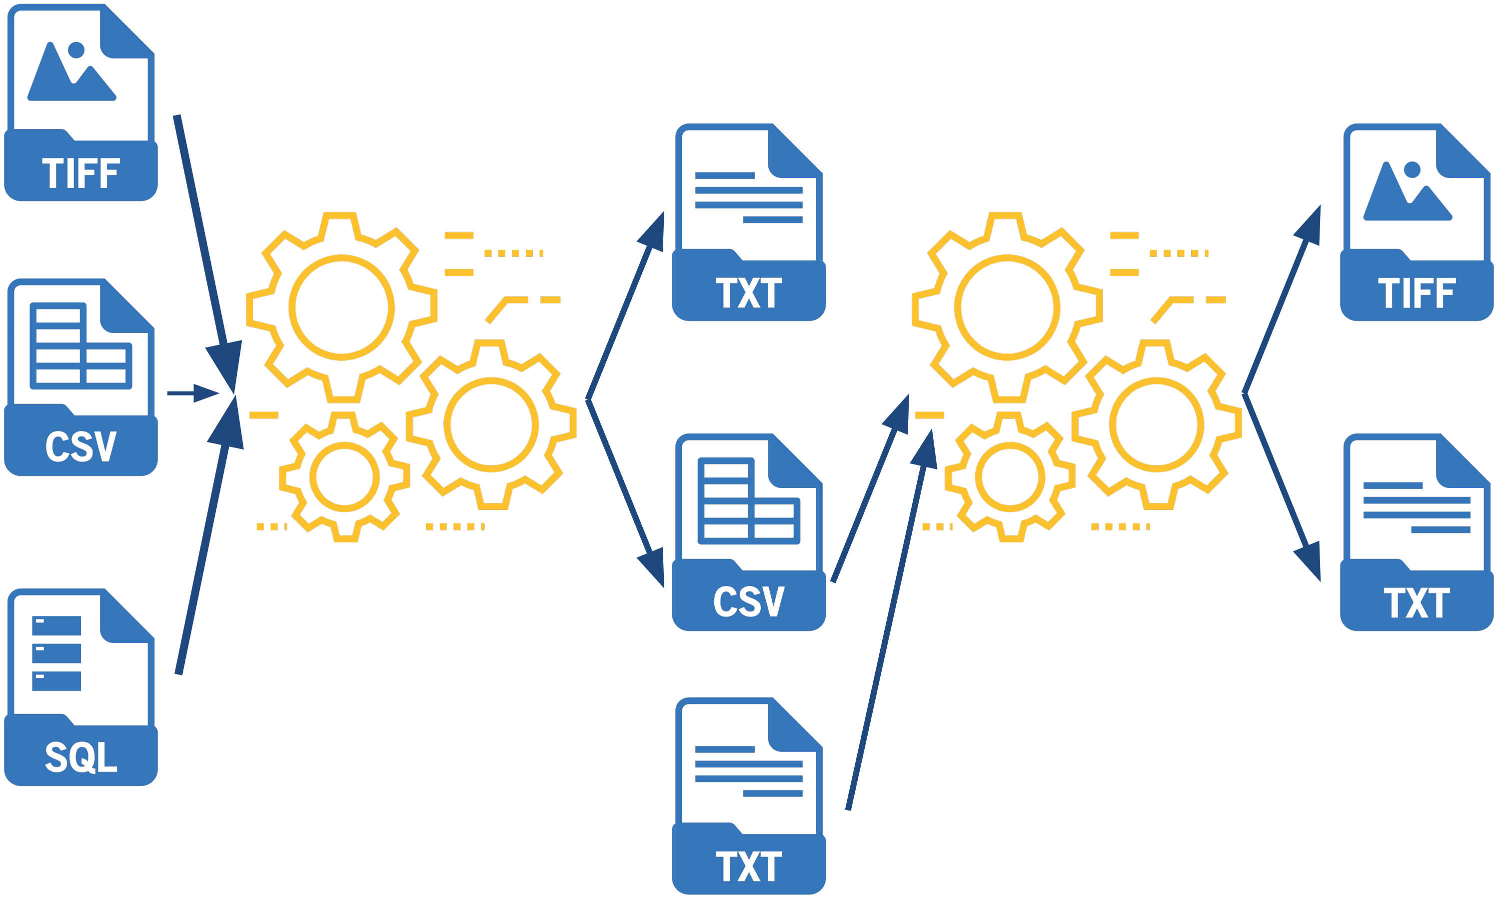

Model Abstraction

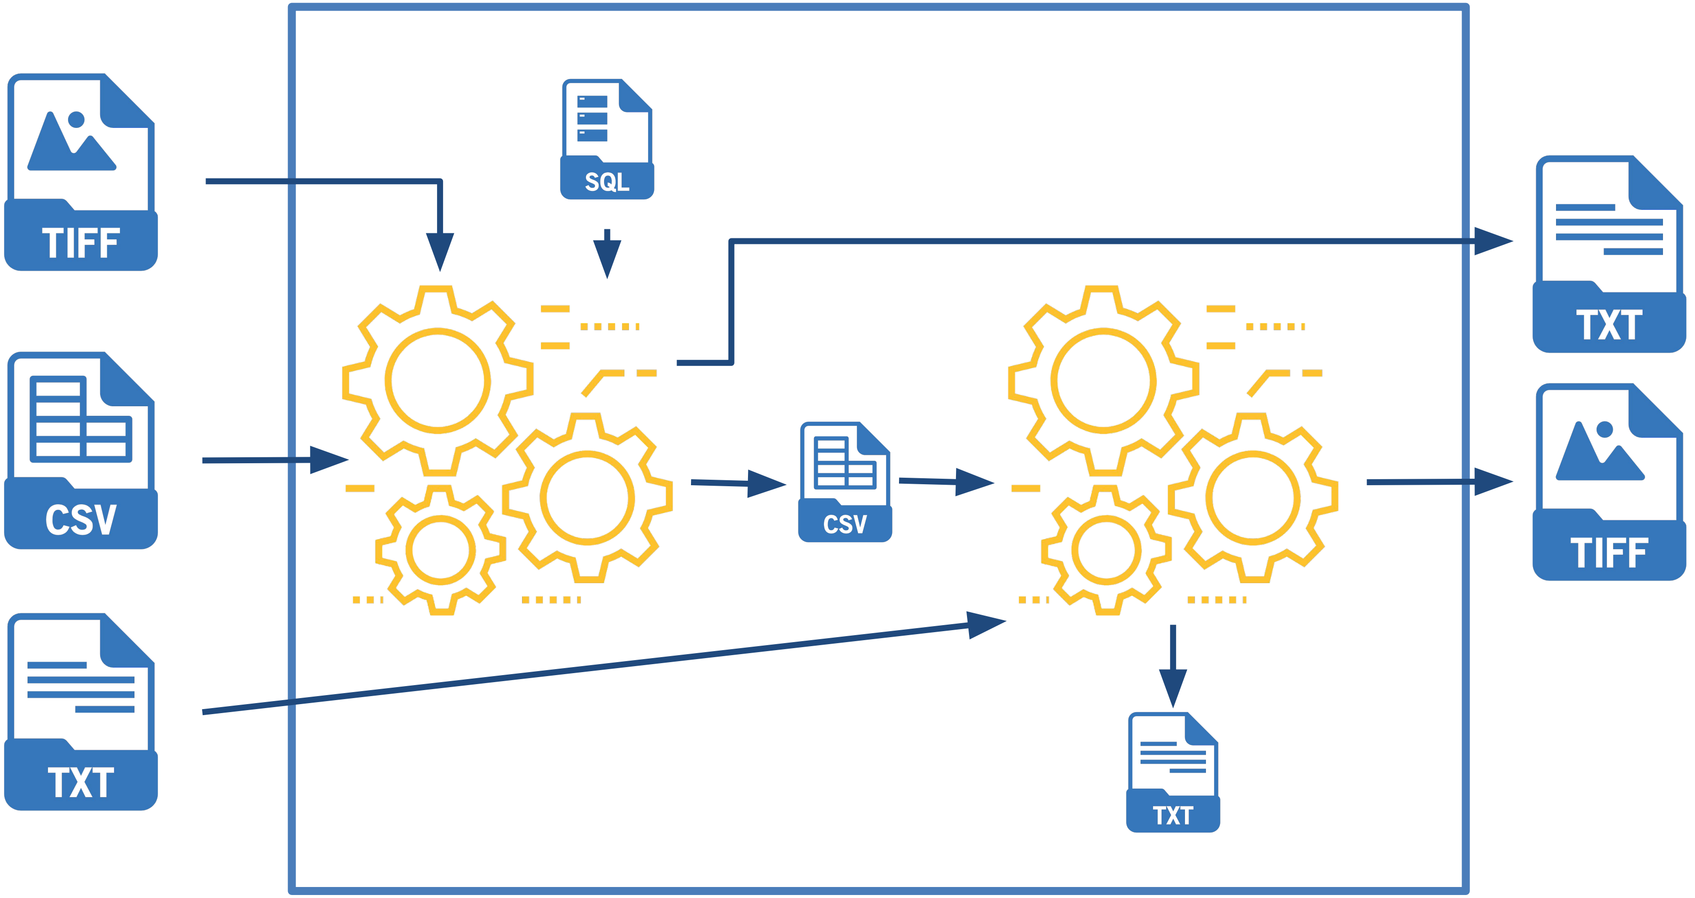

Input Specialization

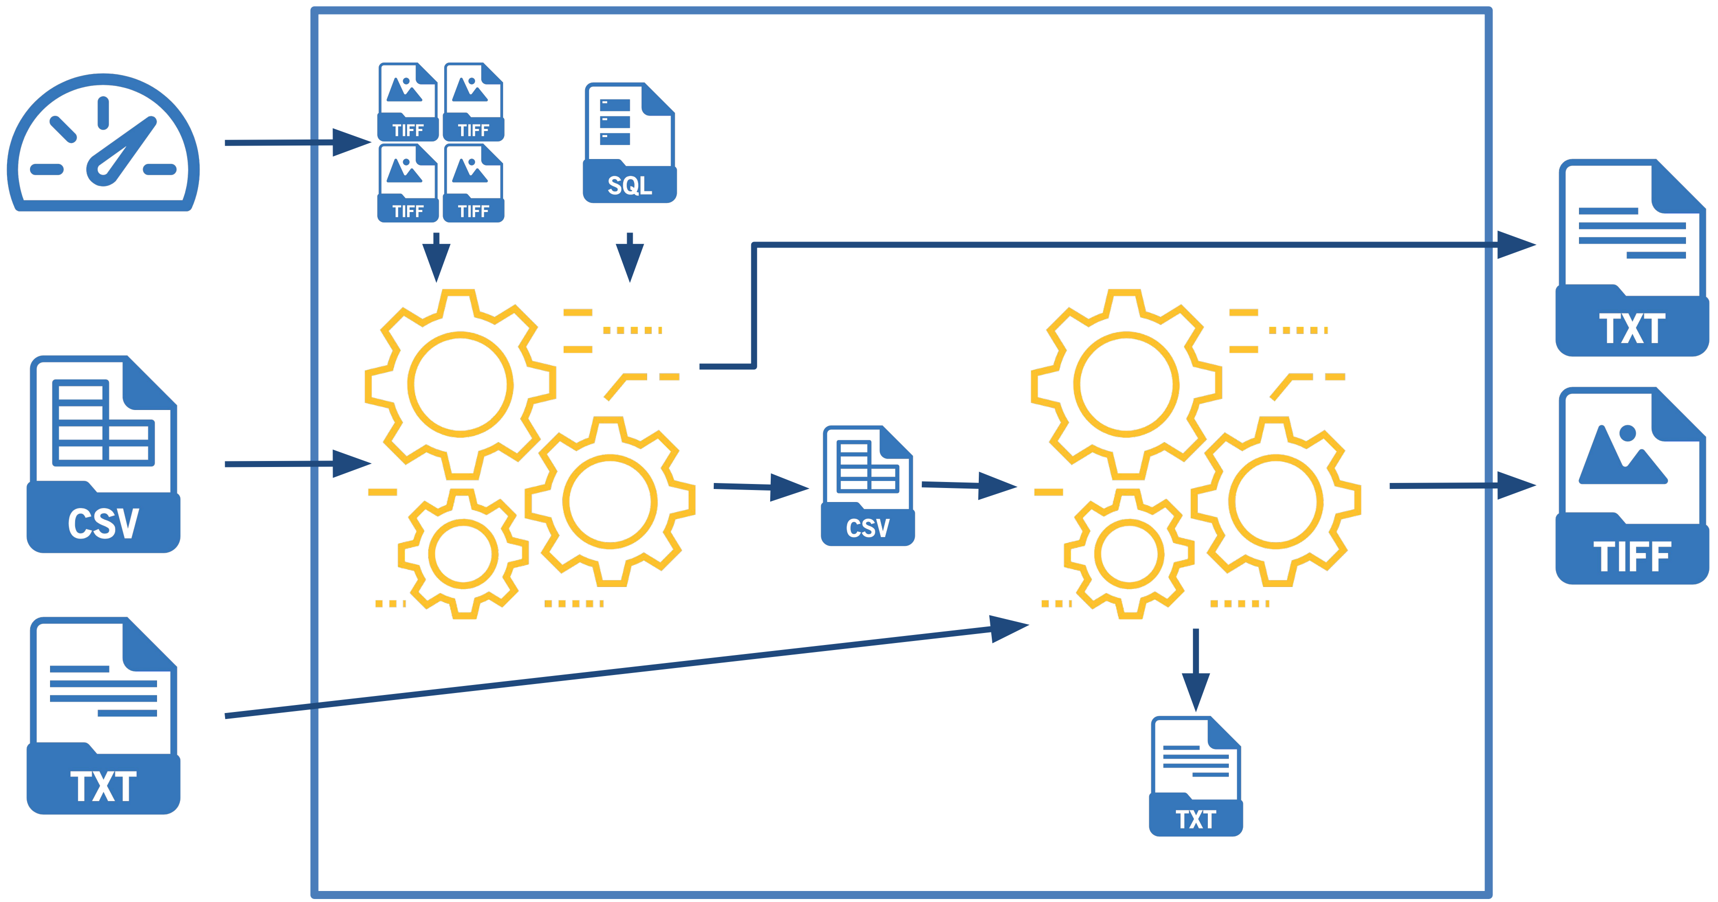

Core Models are the same.
Procedure differs in the Data Selection

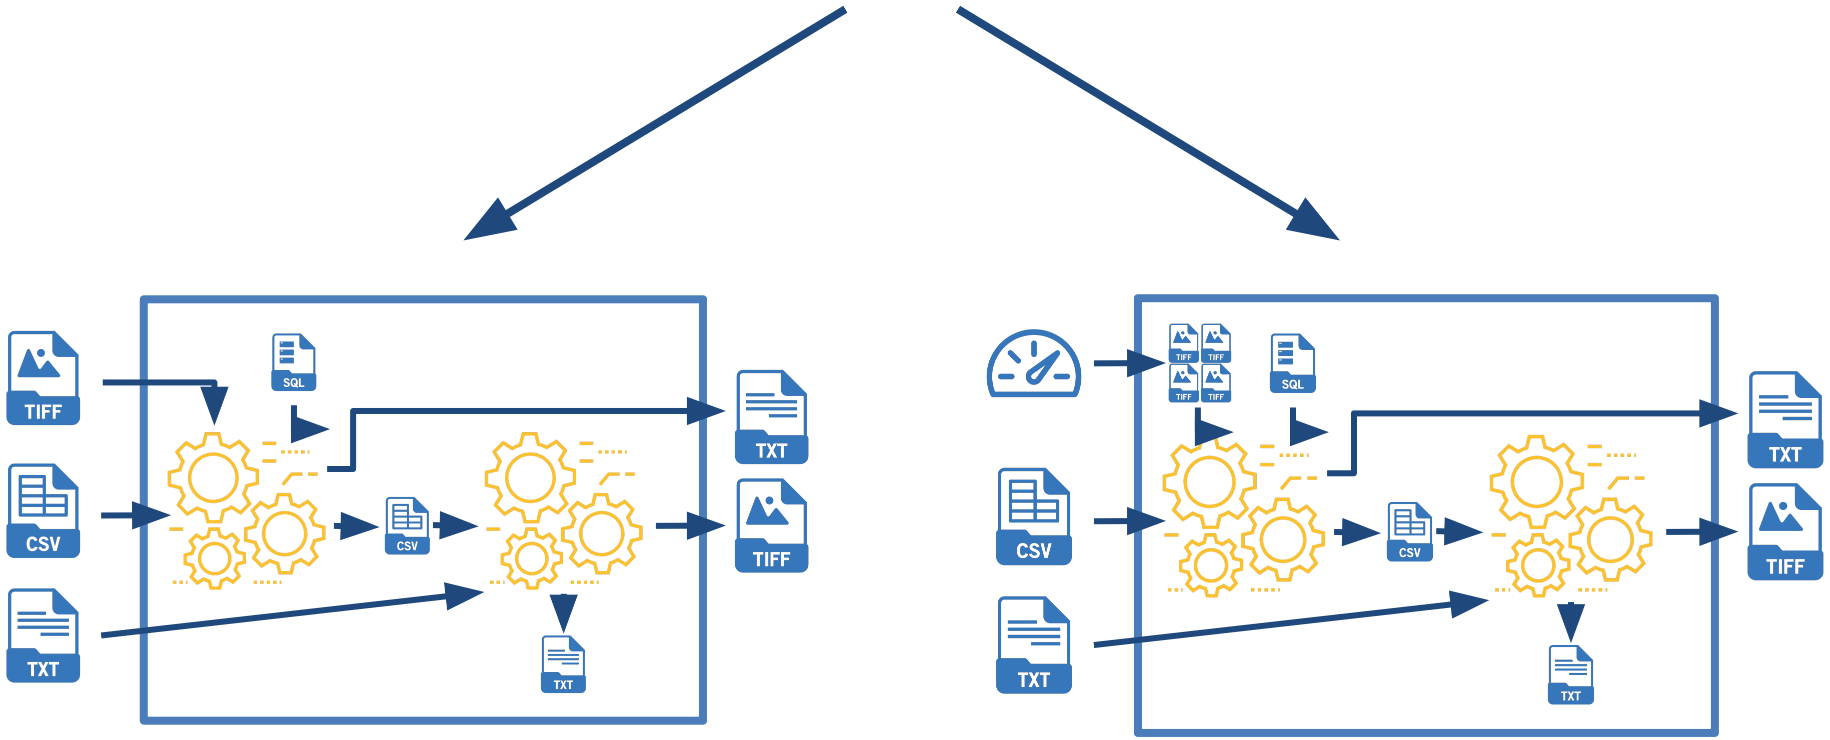

Core Models are the same.
Procedure differs on output, and exposed inputs.

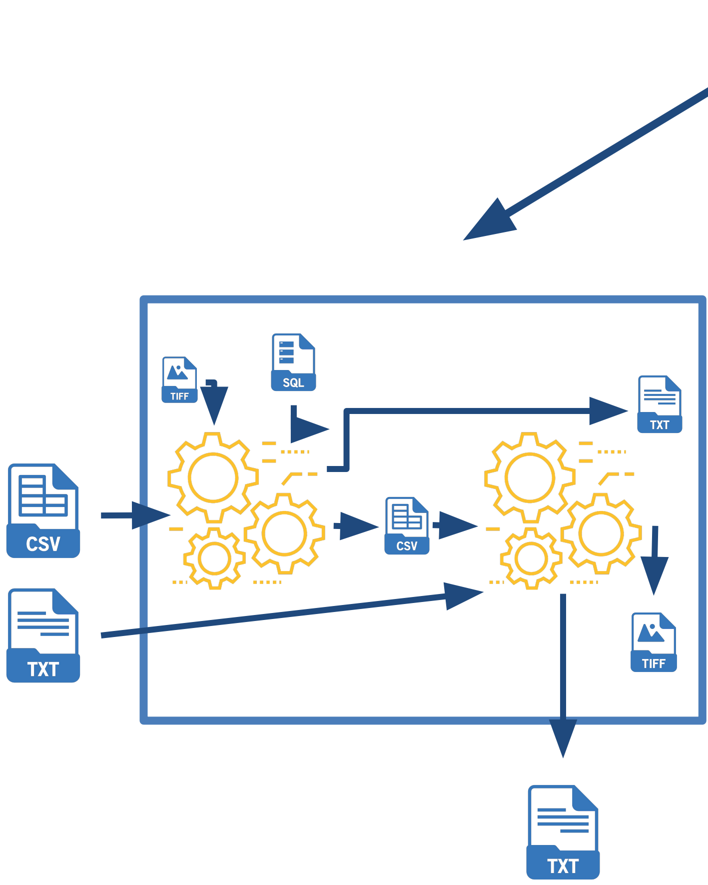

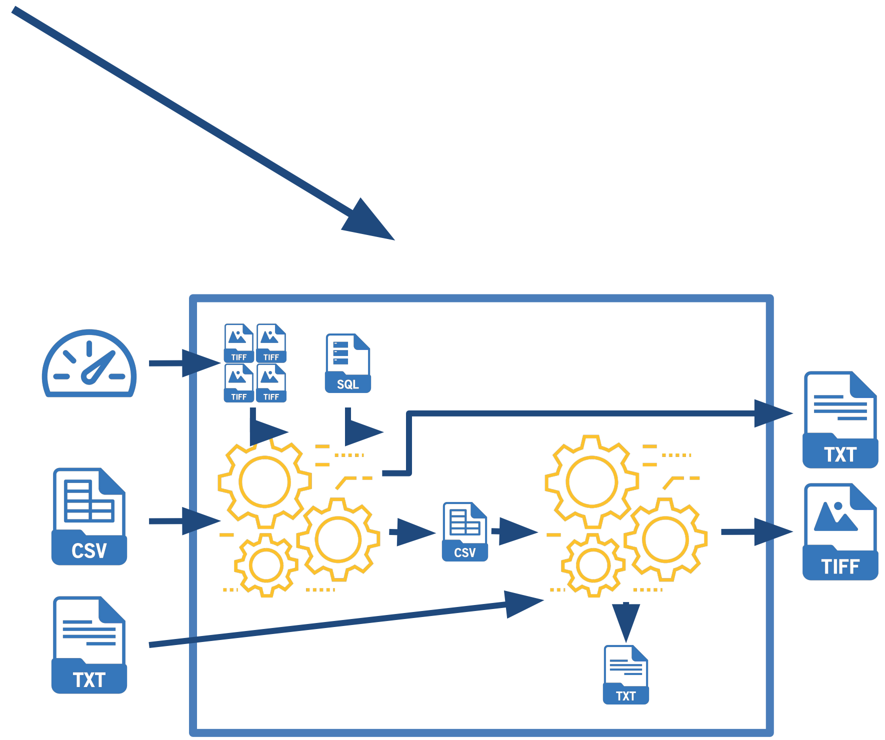

Compositional Modeling

Cooperative, Compositional, Modeling

Enables domain experts, government analysts, and other specialists to work together with versions of models that are abstracted via model registration to provide decision support in composable frameworks.

Abstracts complex model artifacts to more **explainable** and usable forms. Provides **traceability** not only within an given model but in composed pipelines. Pipelines are capable of **introspection** on inputs, outputs, requirements, and specifications to provide **reliable** modeling capabilities during crisis. Artifacts are directly **governable**, and can be controlled with specialized “guard-rails”.

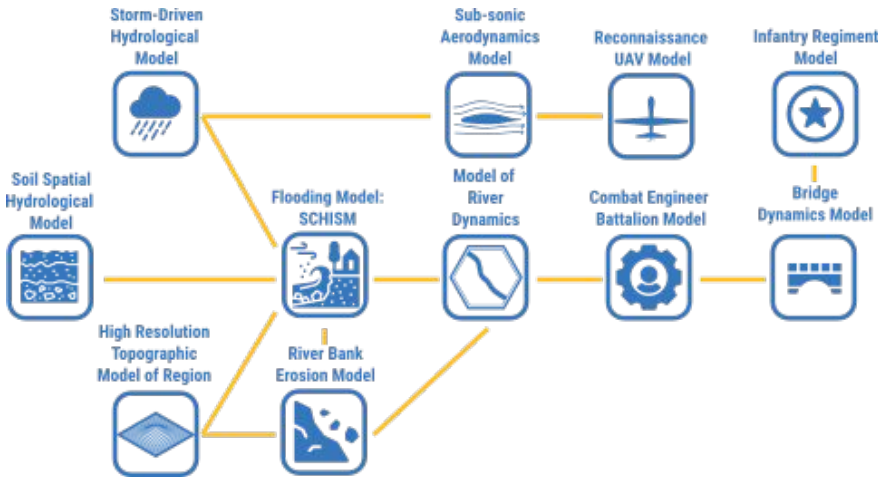

Compositional Modeling

Fusing Models across Domains

Example analysis conducted as part of an exercise modeling ocean currents, wind patterns, and mesoscale coupling. Composes climate models, wind models, and turbulence float models to project the positions of sensors, person-overboard, etc.

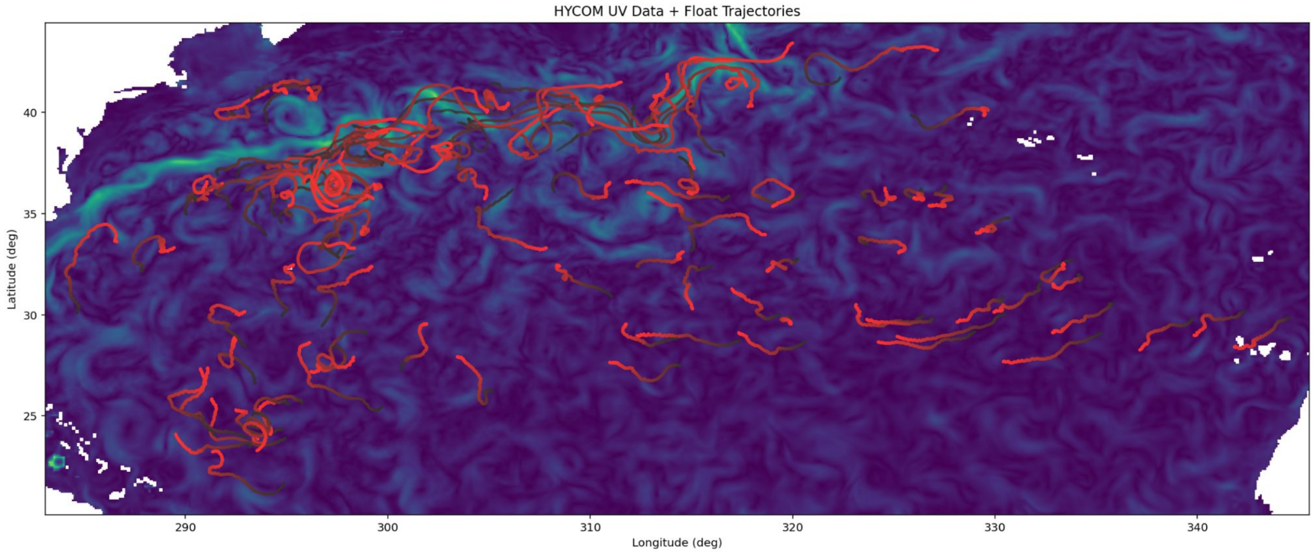

Ethiopia: Case Study

Food insecurity in Ethiopia's Oromia region was studied in a number of exercises with SuperMaaS and models drawn from academic, industrial partners, and open source repositories.

Pipelines to support government analysts were created by inexpert users with little to no experience in the models or their original domains. These pipelines consisted of abstracted hydrology, agronomy, geo-political, human movement, meteorological, and both crisis and intervention models.

Pipelines were used to evaluate the cost and impact of various interventions, and validated against real world scenarios and historical data.

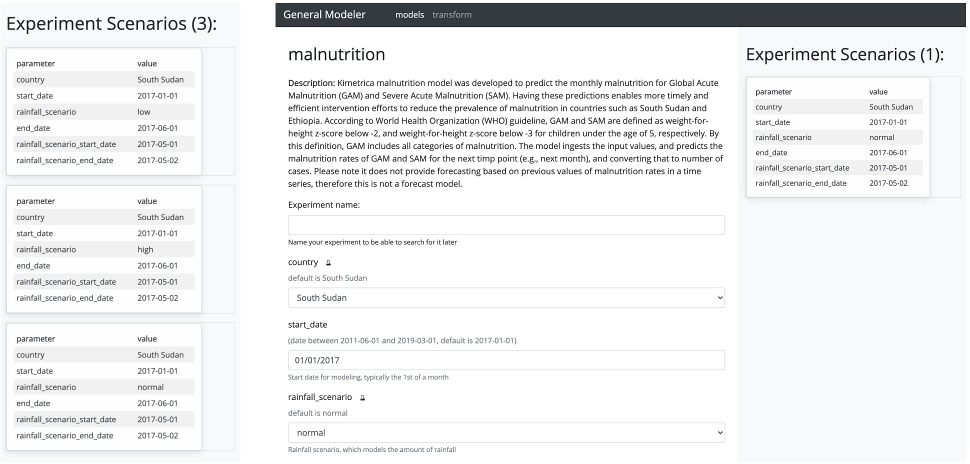

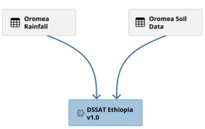

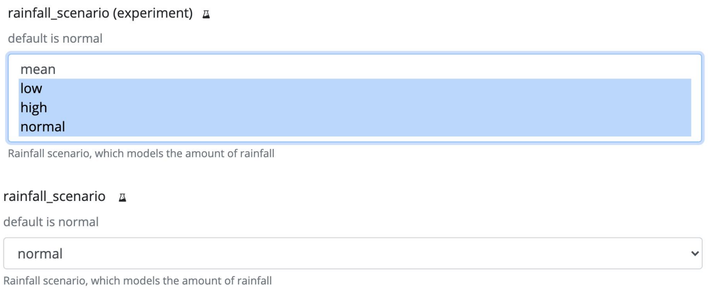

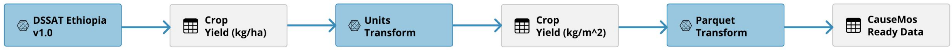

Common Ground as a Meta-Reasoning Tool

A surprising outcome of our use of ORCA has been the automatic construction of model workflows.

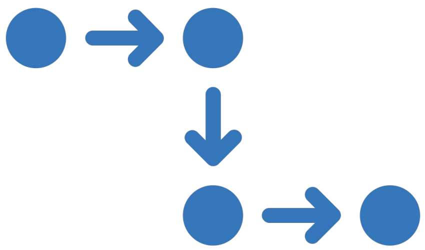

Large library of models and transformations on data provides a rich space to find novel computational/predictive paths.

Often with different properties, assumptions, errors, and **execution time**.

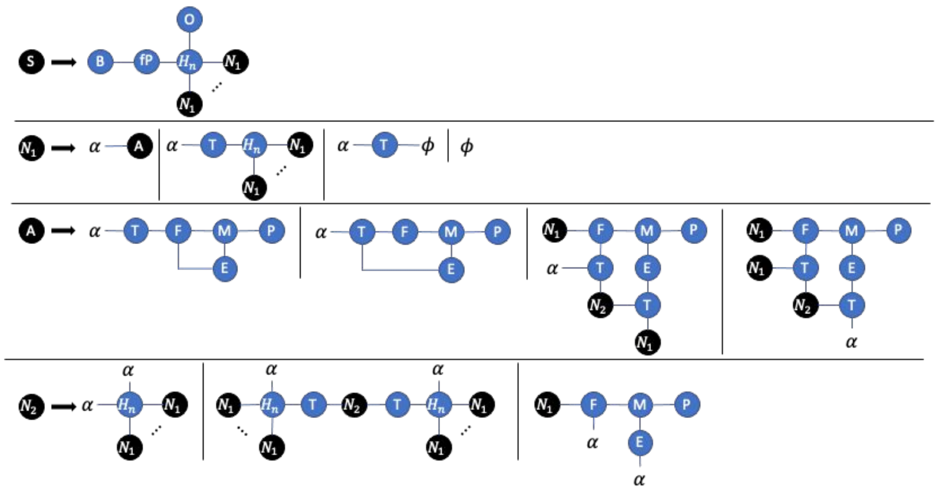

ORCA Framework

Using ORCA, we can optimize the models and functionality implemented by a given model, respecting its ORCA API.

Moving compute to the edge.

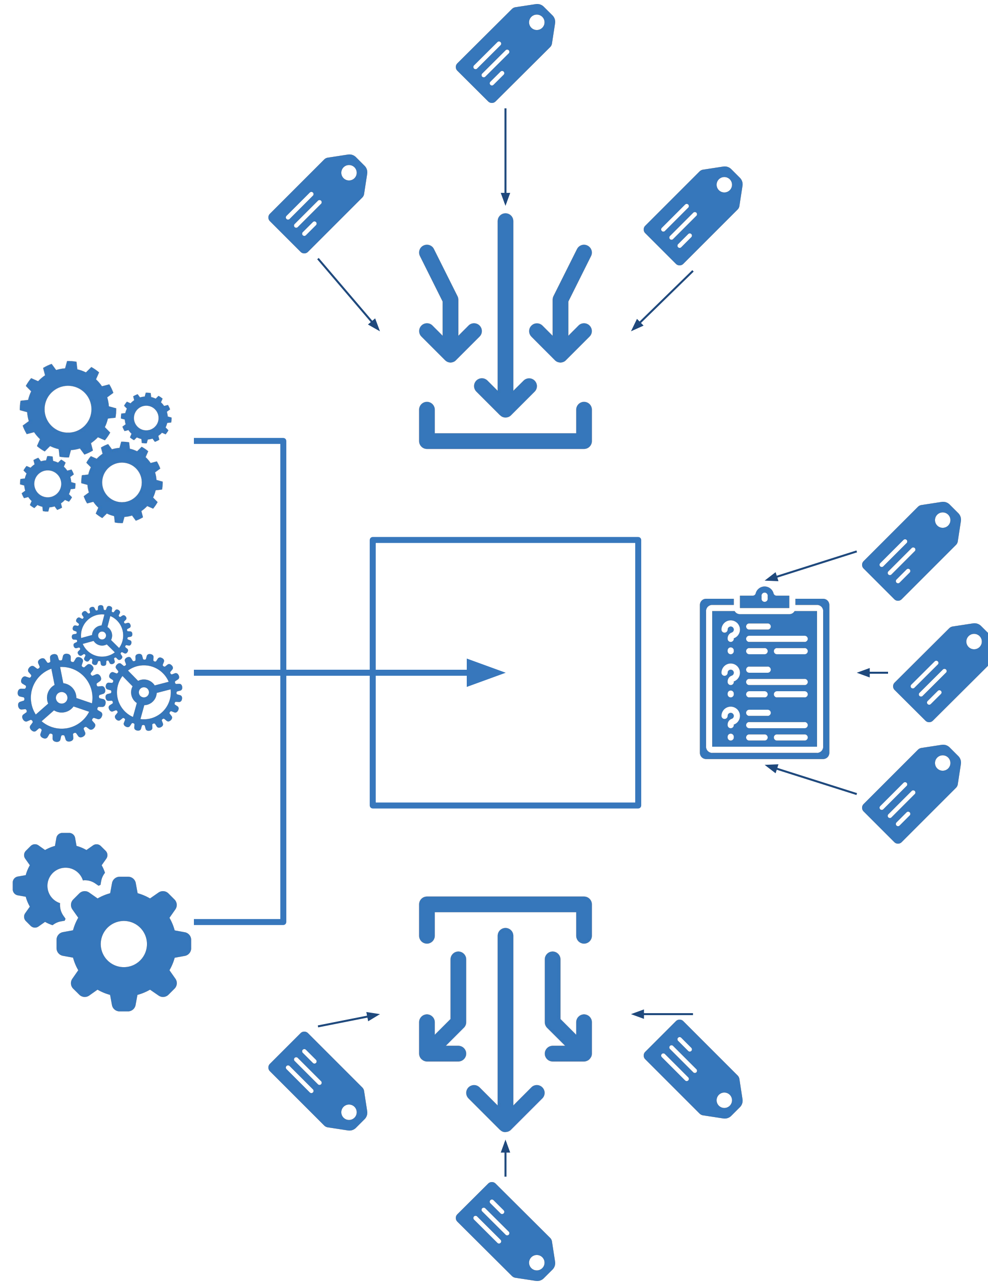

Edge Compute

Crisis response has diverse compute requirements and resources.

Some jobs can be executed on traditional hardware.

Other compute workloads need to be optimized or approximated for edge compute.

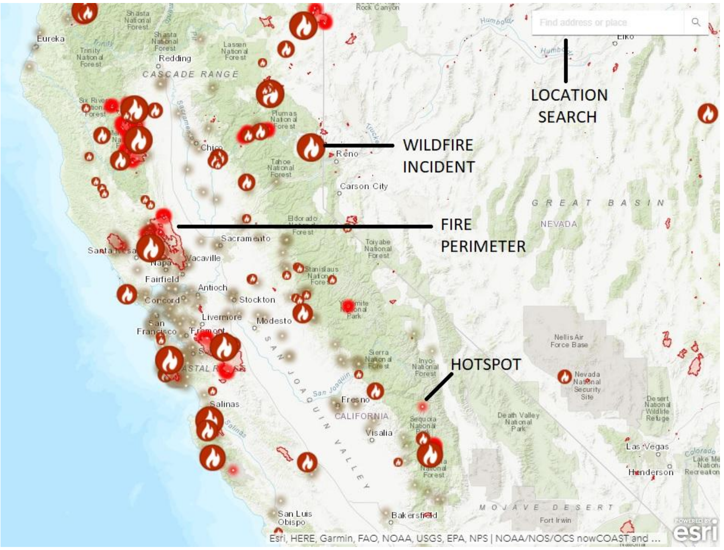

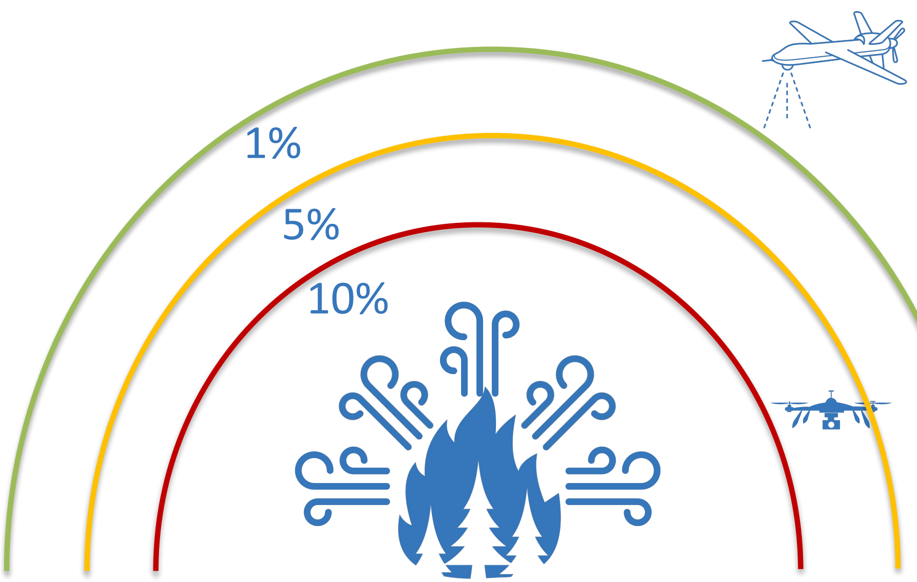

HARVARD
MEDICAL SCHOOL

COVID-19 Model Extraction

Mechanistic models are often large, describing an entire system.

Measuring the ***covid19_IL18_active*** pathway of the model.

Original Model **Extracted Model**

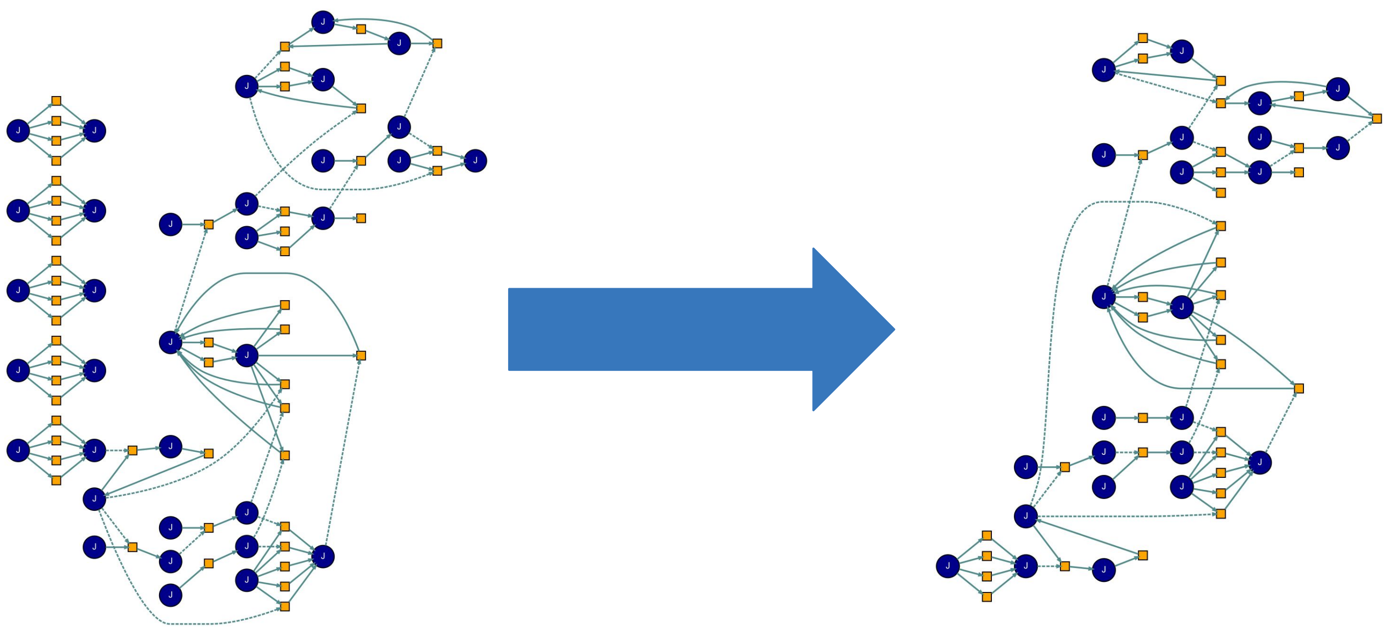

COVID-19 Model Extraction

Mechanistic models are often large, describing an entire system.

Measuring the ***covid19_IL18_active*** pathway of the model.

Original Model

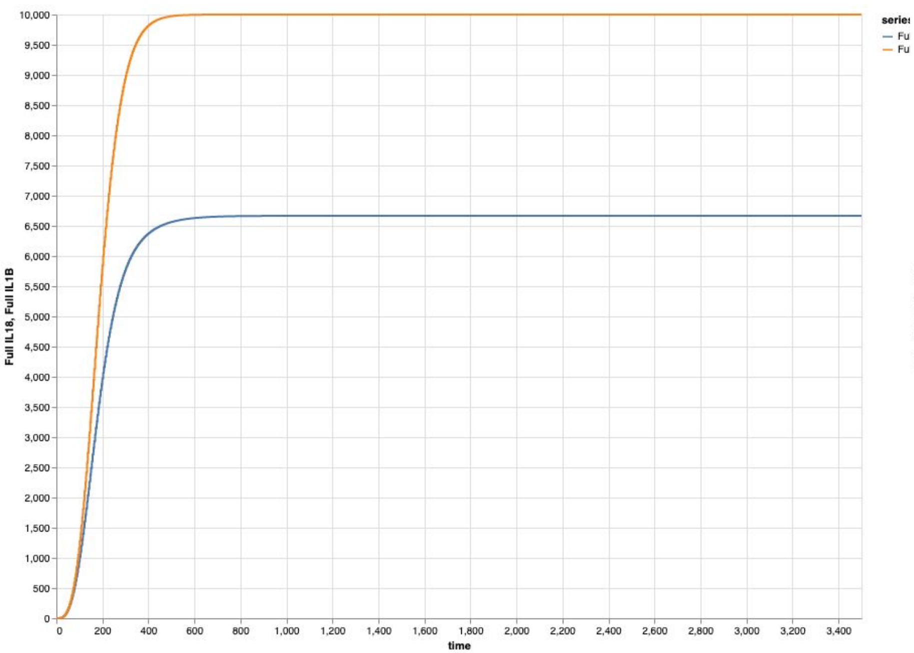

Extracted Model

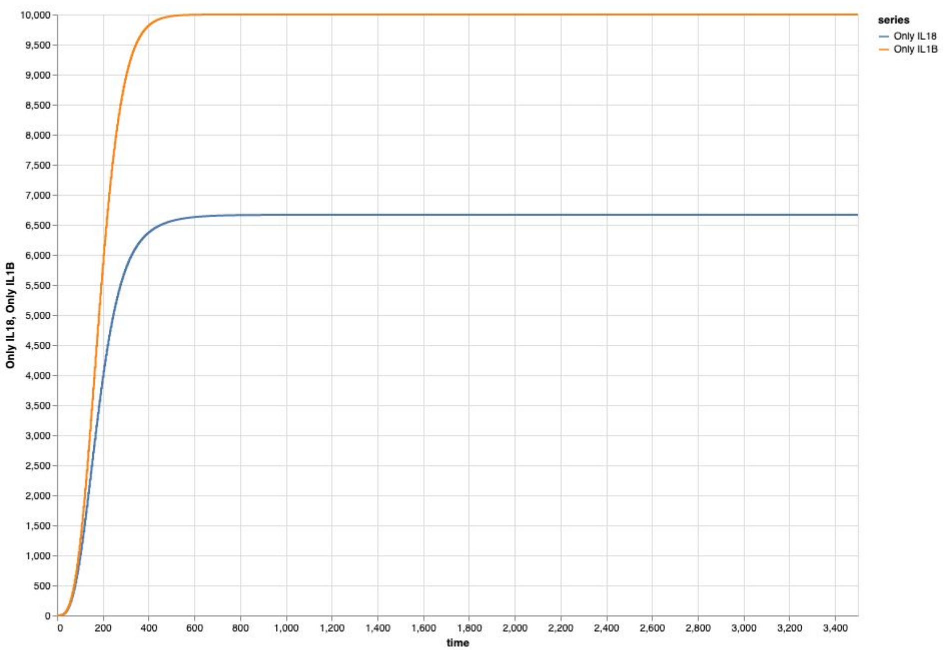

COVID-19 Model Extraction

Mechanistic models are often large, describing an entire system.

Measuring the ***covid19_IL18_active*** pathway of the model.

Method is **exact**.

Results in a model that runs 1.4x faster, and is 0.75 the size in state variables and 0.63 the size in events.

Removes many events which increased stiffness of the model, but did not interact with ***covid19_IL18_active***.

Extracted Model

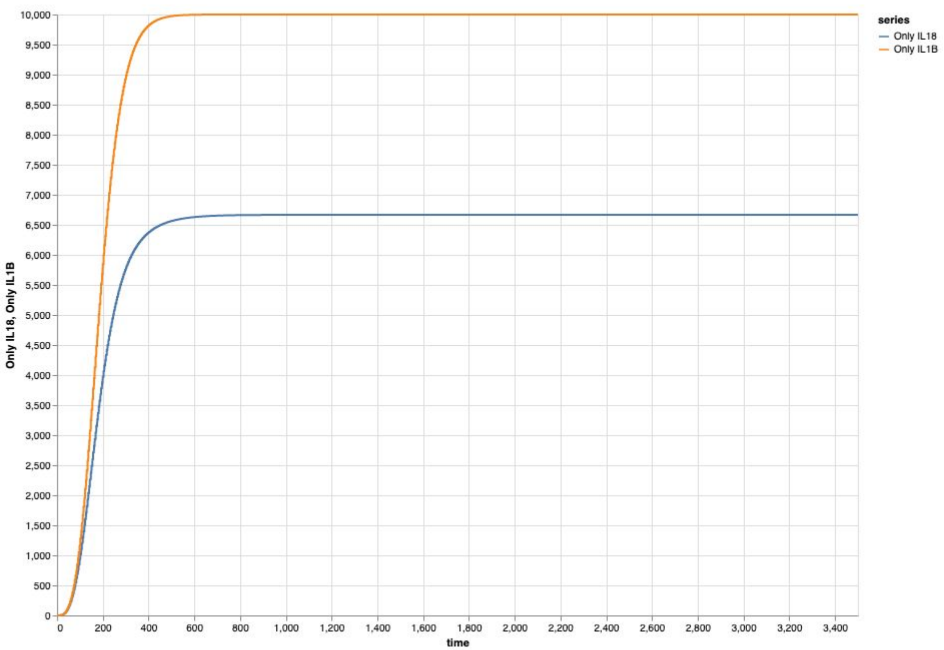

HARVARD
MEDICAL SCHOOL

Extracting the MAPK Pathway from the RAS Machine

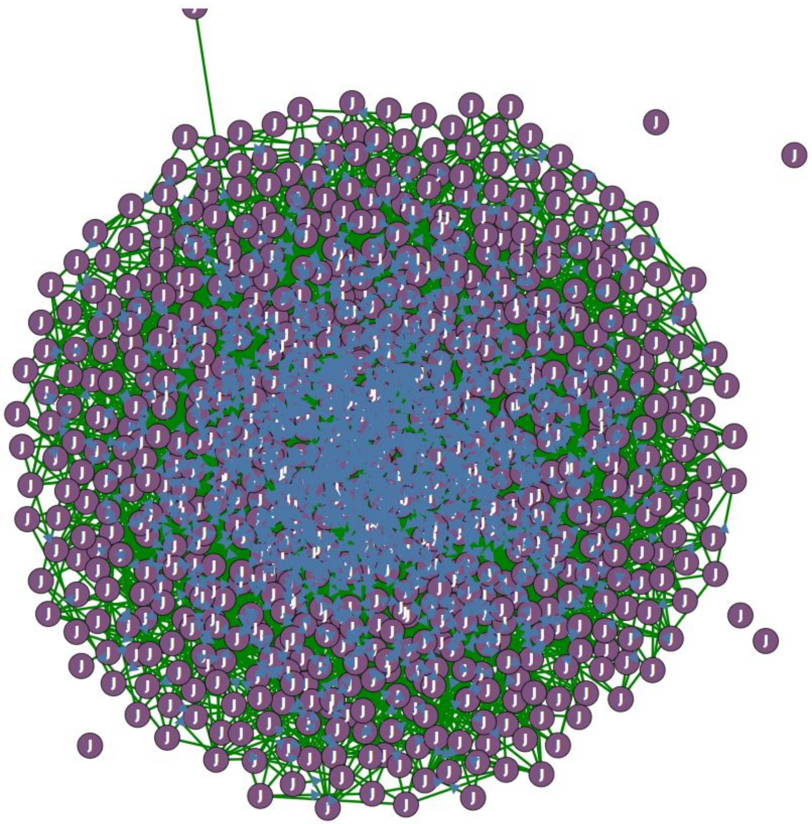

RAS-machine is an example of a larger biochemical system.

- 849 state variables
- 21,592 events

Extracting the MAPK Pathway from the RAS Machine

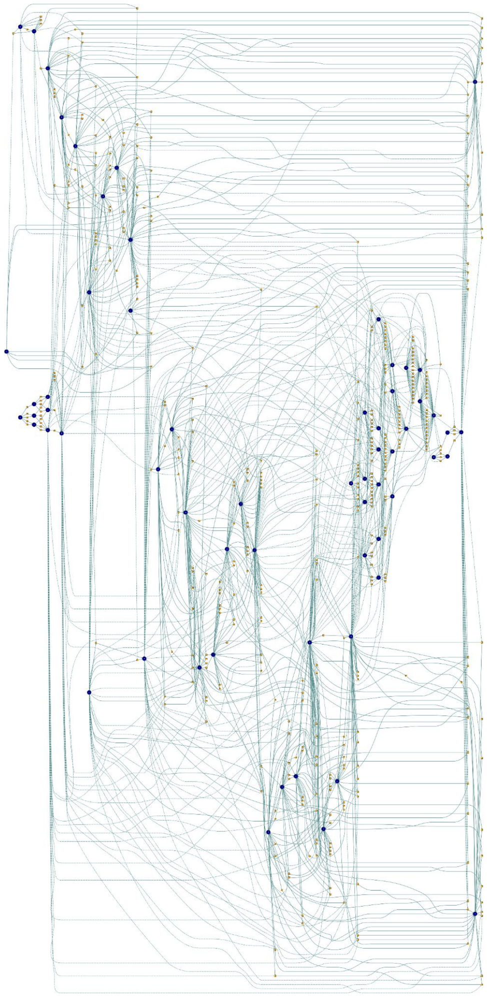

RAS-machine is an example of a larger biochemical system.

Extract the MAPK pathway:
[MAPK1_kinase=inactive_cytoplasm, MAPK1_kinase=active_cytoplasm, MAPK1_kinase=inactive_cytoskeleton, MAPK1_kinase=active_cytoskeleton]

- 8.4% of state variables
- 2% of the events
- 44x speed up, including simplification time
  - 112x speed up in the limit (achieved when fitting the model)

Extracting the MAPK Pathway from the RAS Machine

- 8.4% of state variables
- 2% of the events
- 44x speed up, including simplification time

Original Model

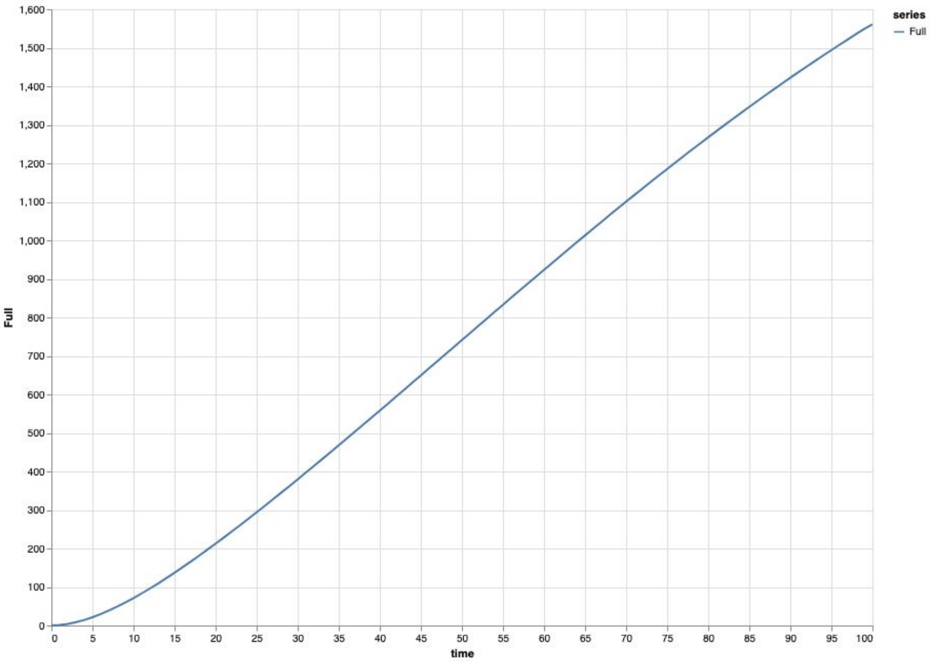

Extracted Model

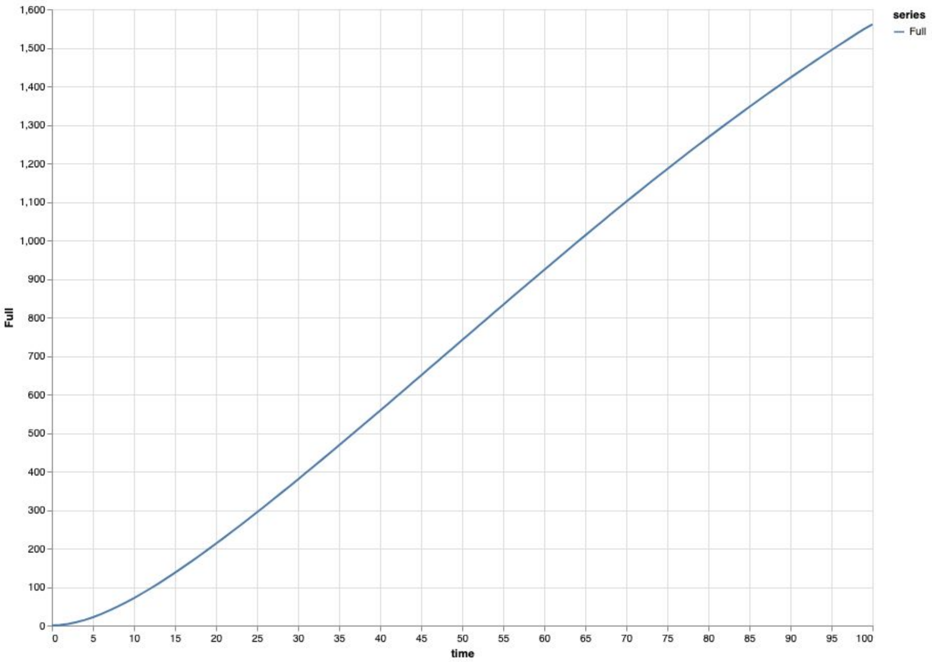

General Surrogate Models

What about systems where our original model is unavailable, or cannot be reduced?

What about systems where our models are prohibitively expensive for our use case?

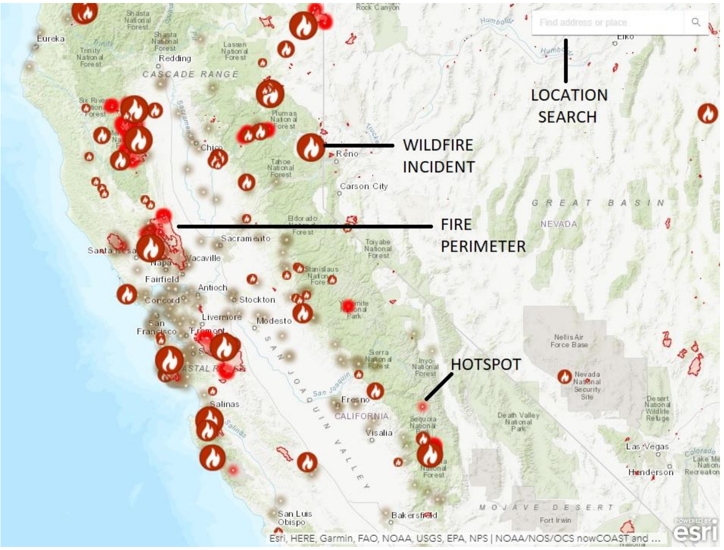

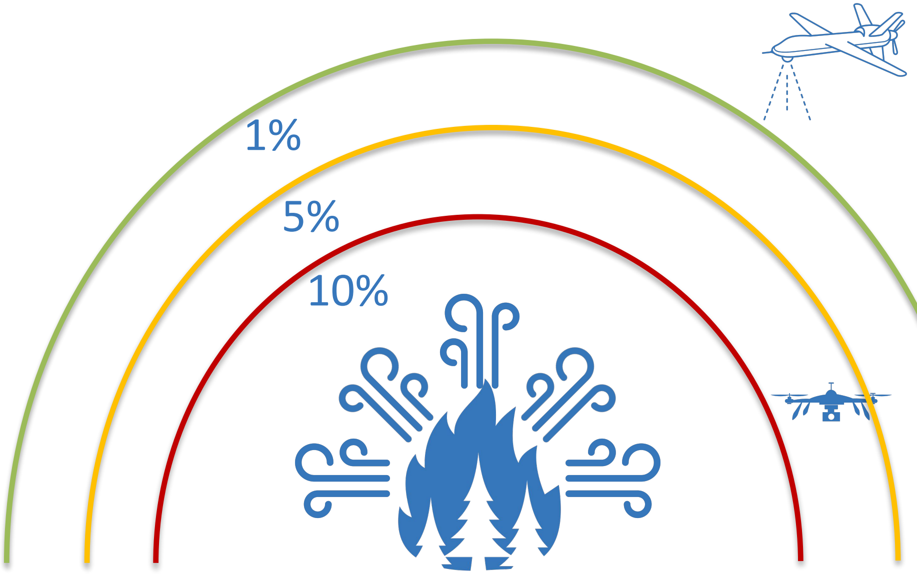

Design Space Exploration

- Build a *surrogate model* of a vehicle design's *analysis pipeline*, mapping *component assemblies* to useful *design metrics*.
- The designs are richly structured (property graphs)
- The design metrics are often informed by the dynamics of the vehicles.
- Techniques that utilize known physics and design information are needed to overcome the generalizations imposed by the design variability.
    - Bonus points if the model can be decomposed, reusing some stages for other design tasks.

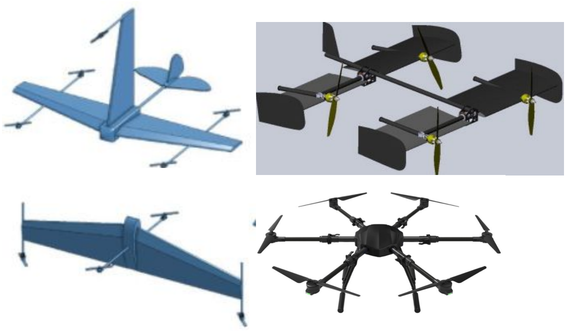

Train Set

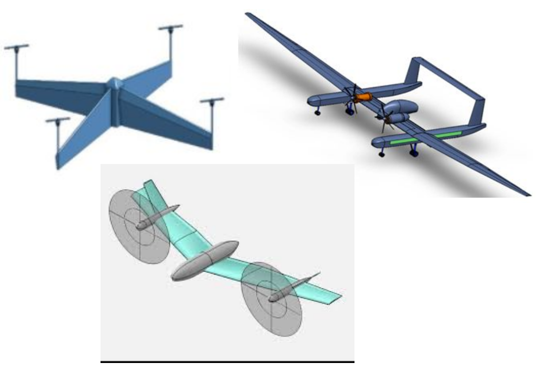

Test Set

$\hat{f}\Bigg($  $\Bigg) = \begin{bmatrix} 75.2 & \text{max trim state (m/s)} \\ 3050 & \text{flight time (s)} \\ \vdots & \\ 0.1 & \text{max path deviation (m)} \end{bmatrix}$

Vehicle Structure as Graphs

Assembled UAVs are labeled property graphs, with the vertices being physical components and the edges being structural connections.

- [SwRI] JanusGraph holds a detailed property graph, where design components are also property subgraphs
- [SwRI] Autograph describes a UAV as a program that builds a graph incrementally
- CyPhyPy takes an assembly graph that is closer to the Autograph graph description
    - Hierarchy creates *subassembly abstractions*
    - Typed nodes create *component abstractions*
    - Flattening refines a design to parts that can be simulated
    - Vertices have *continuous* and discrete properties

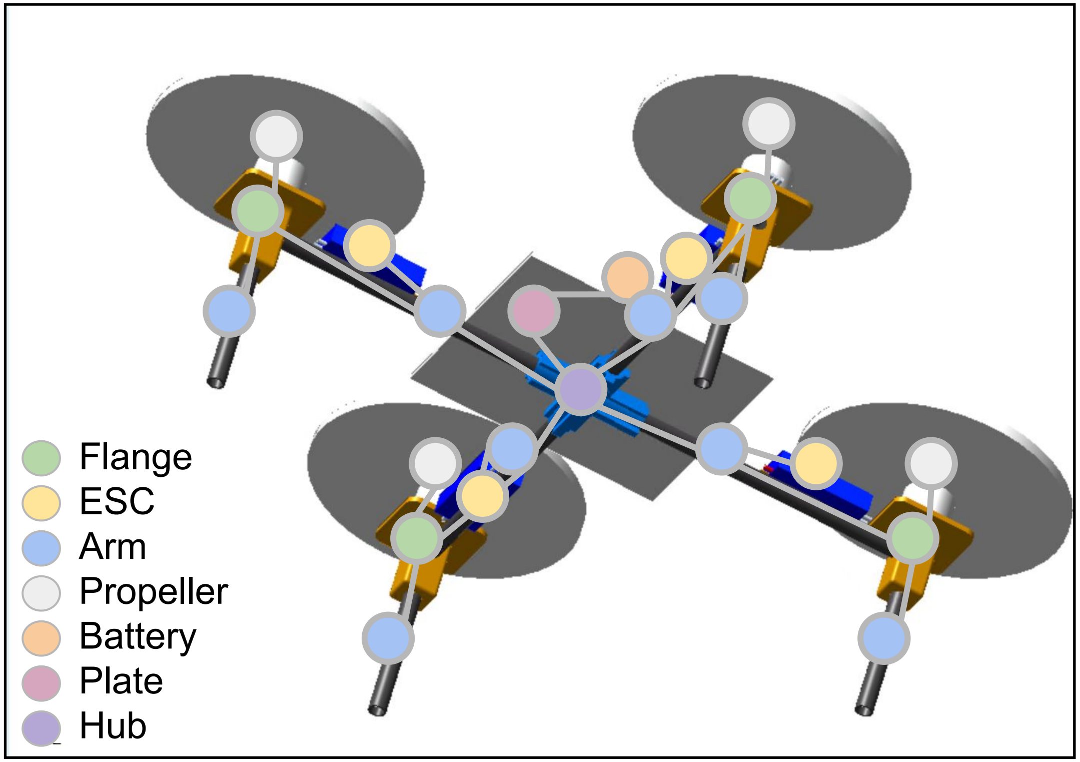

Surrogate modeling

- The problem:
  - Exact design simulations take significant time to run (full CAD and aerodynamic models).
  - The total search space that can be explored gets reduced.

- Our solution:
  - Use machine learning to train surrogate models using small amounts of exact simulation data.
  - Once trained, use the surrogate model to evaluate large numbers of designs very rapidly.

Essentially, surrogates incur a one-time initial cost and then become “free” to evaluate any number of new designs.

Equations

Equal torque and angular velocity

Motor

Equal voltage

Battery and control

Propeller

$$\omega = 2\pi n$$

$$J = \frac{\mathcal{V}}{nD}$$

$$\eta_P = \frac{C_t J}{C_p}$$

$$T = C_t \rho n^2 D^4$$

$$\tau = \frac{C_p \rho n^2 D^5}{2\pi} = \frac{C_p \rho \omega^2 D^5}{(2\pi)^3}$$

$$P_P = \tau\omega = C_p \rho n^3 D^5 = \frac{C_p \rho \omega^3 D^5}{(2\pi)^3}$$

$$I = \frac{V - \frac{\omega}{K_V}}{R_w}$$

$$\tau = K_T(I - I_0) = \frac{K_T}{R_w}\left(V - \frac{\omega}{K_V} - I_0 R_w\right)$$

$$P_M = VI = \frac{V^2 - \frac{\omega V}{K_V}}{R_w}$$

$$\eta_M = \frac{P_P}{P_M} = \frac{C_p \rho n^3 D^5}{VI}$$

$$\frac{R_w C_p \rho D^5}{K_T (2\pi)^3}\omega^2 + \frac{1}{K_V}\omega + (I_0 R_w - V) = 0$$

$$V = \frac{1}{2}\left(\frac{\omega}{K_V} + \sqrt{\left(\frac{\omega}{K_V}\right)^2 + 4 P_M R_w}\right)$$

$$\frac{C_p \rho \omega^2 D^5}{(2\pi)^3} = \frac{K_T}{R_w}\left(-\frac{\omega}{2K_V} + \sqrt{\left(\frac{\omega}{2K_V}\right)^2 + P_M R_w} - I_0 R_w\right)$$

$$I_B = \frac{I}{\eta_E}$$

$$P_B = V I_B$$

Deep Koopman

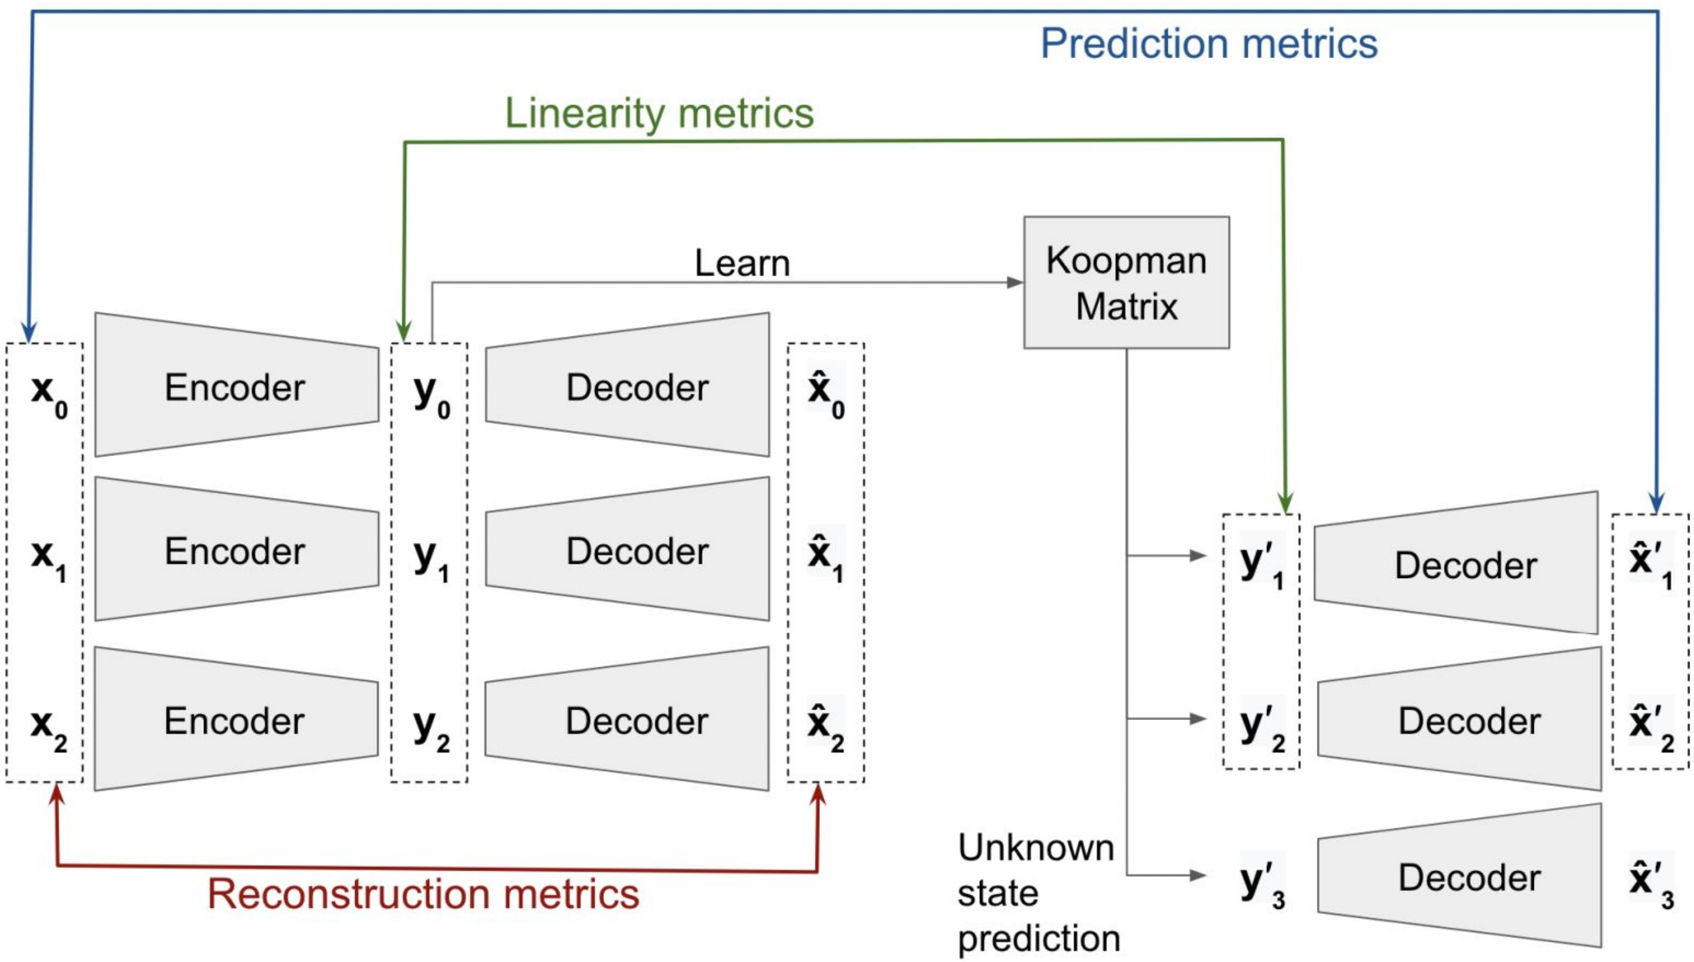

Deep Koopman

Individual component data present in corpora.

Equations connecting components obtained from flight dynamics model.

Train deep neural net Koopman models to predict any dependent variable for new values of the independent variable propeller RPM.

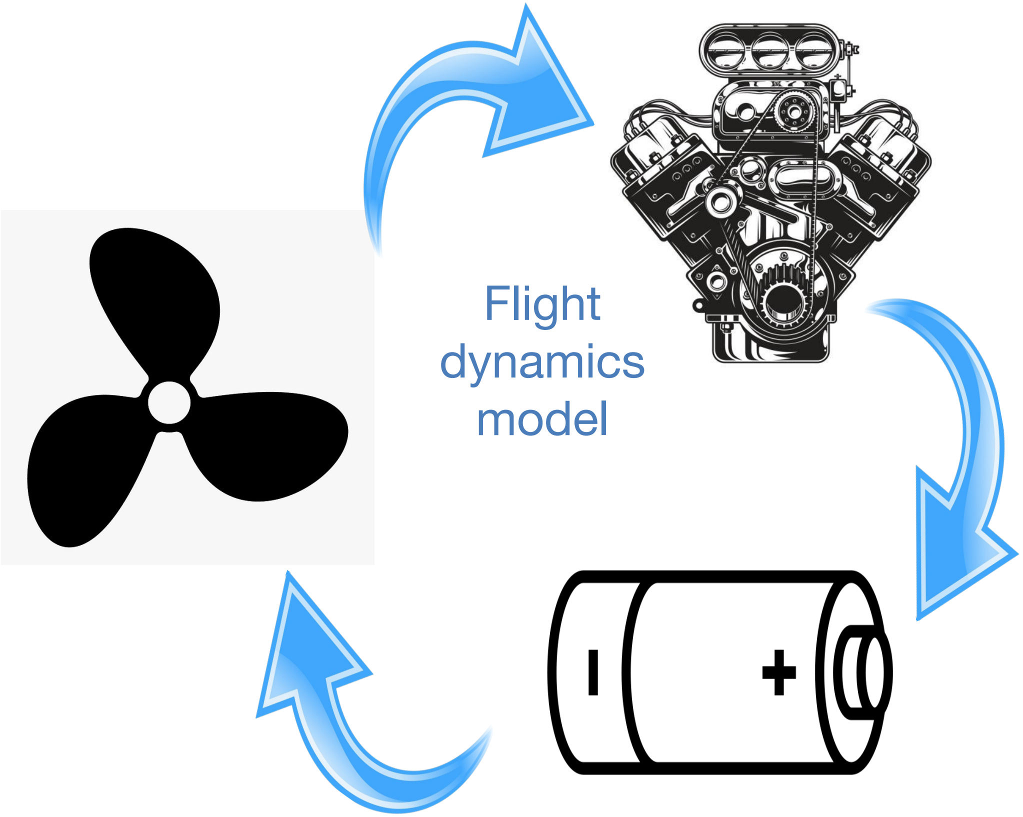

Inexact, But Fast

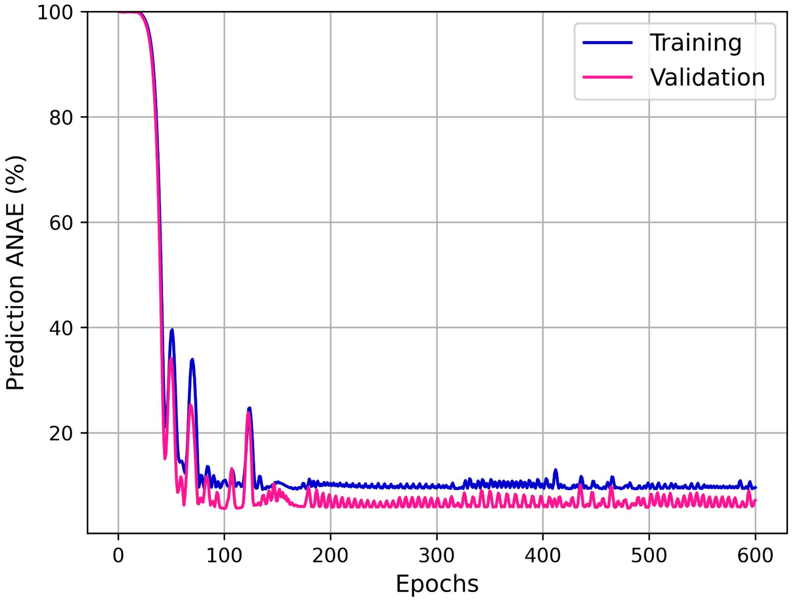

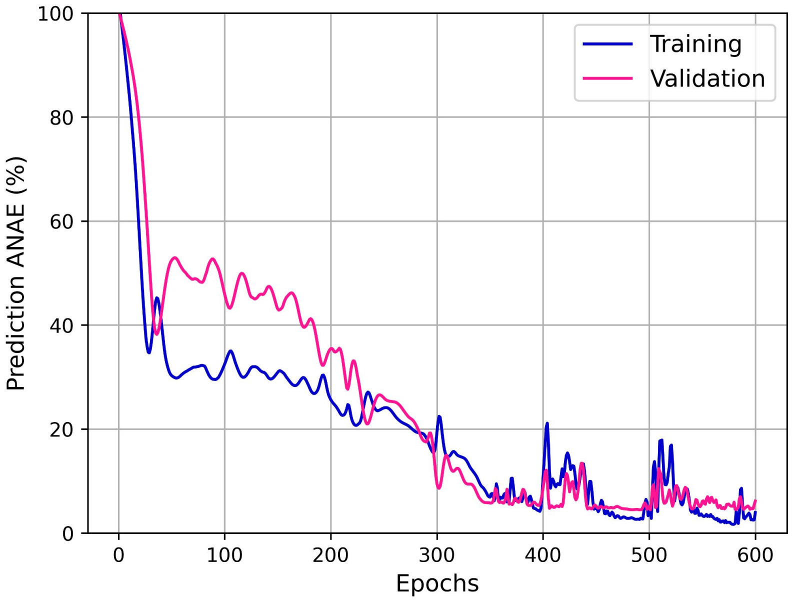

<10% error

Model thrust across all velocities of propeller `83x48_5_400_51_390` as a function of RPM.

Model current of motor `MAGiDRIVE500` connected to propeller `83x48_5_400_51_390` across all its velocities as a function of RPM.

Performance Across the Design Space

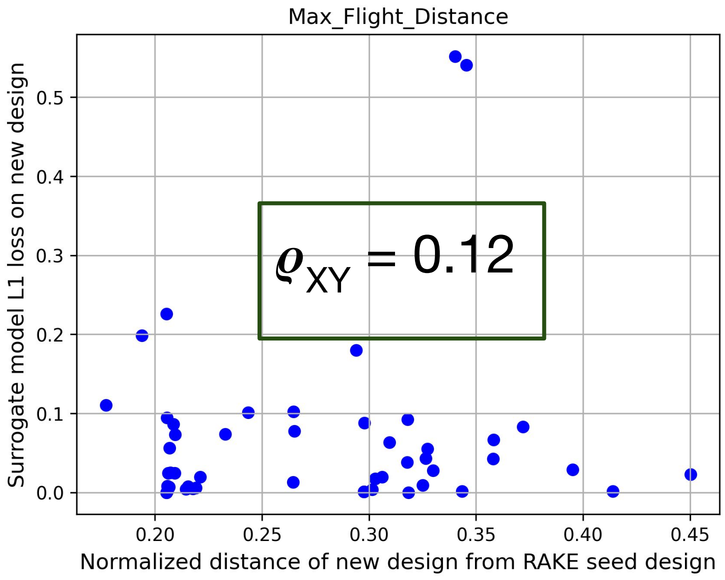

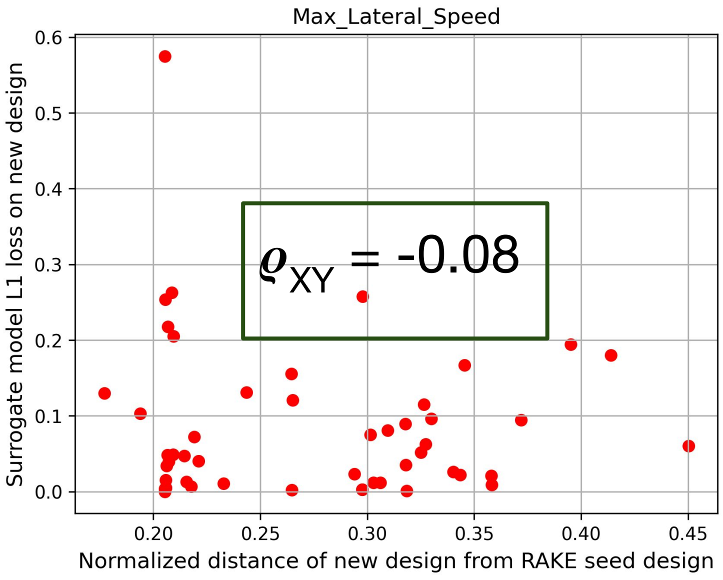

There is almost zero correlation ϱ between X-axis = "**distance** of a new design from the seed design" and Y-axis = "surrogate **performance** on the new design".

Surrogates perform robustly in the 'corners' of the design space.

Compositionality and Prediction

All components affecting performance are considered, i.e. the approach naturally achieves composition

We answer "How does the design (i.e. components) evolve with performance?"

Example: Learn DeepK model on given data,

then predict

Hover time = 50s		Hover time = 100s		Hover time = 150s
• Prop diameter = 10m • Motor current = 1.5A • Battery # cells = 5 • Total cost = $250	→	• Prop diameter = 7m • Motor current = 2.5A • Battery # cells = 7 • Total cost = $400	→	• Prop diameter = 5m • Motor current = 4A • Battery # cells = 10 • Total cost = $600

Required Hover time = 80s
- *Prop diameter = 9m*
- *Motor current = 2A*
- *Battery # cells = 6*
- *Total cost = $320*

Novel component synthesis

Required Hover time = 500s
- *Prop diameter = 3m*
- *Motor current = 7A*
- *Battery # cells = 14*
- *Total cost = $950*

Explainability and exploration

Design Candidate Analysis: Clustering

- Experiment Setup
    - The experiment used 10k design sketches from Purdue.
    - **Spectral Clustering** was used to label the designs
        - Any technique that uses a distance matrix can use a kernel via an affinity relationship (Gaussian RBF in this case)
        - Number of clusters were chosen via a gap statistic
    - Kernel PCA was used to visualize the data in a 3D plot (n_components=3)

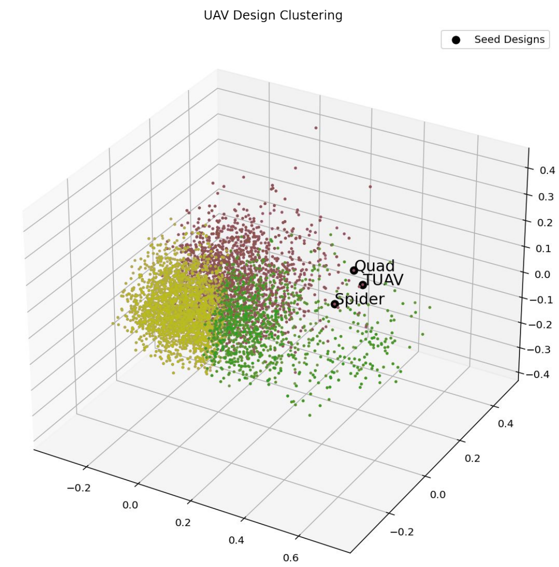

G0: i = 0

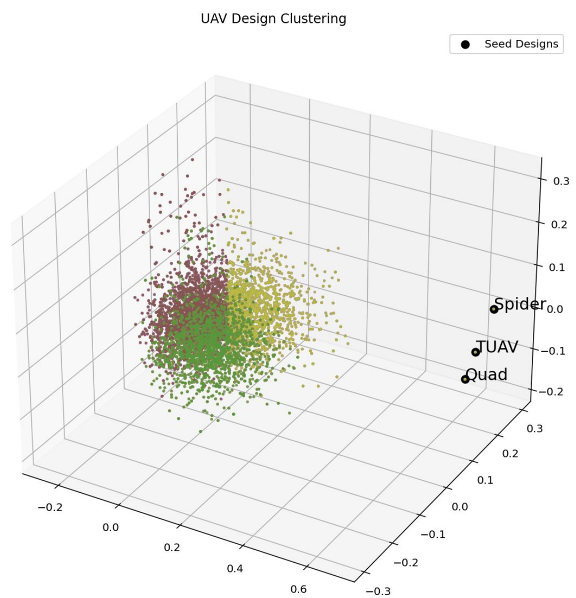

G1: i = 1

Case Study

DARPA DSO CompMODS

Model synthesis for NASA heat shield during reentry. Vehicle undergoes significant heating during aerobraking, critical to model accurately for Mars payloads.

Prediction of the surface heating is necessary to design the thermal protection system. LINK allows composition from individually verified models, tracking accuracy and correctness requirements in the composed model.

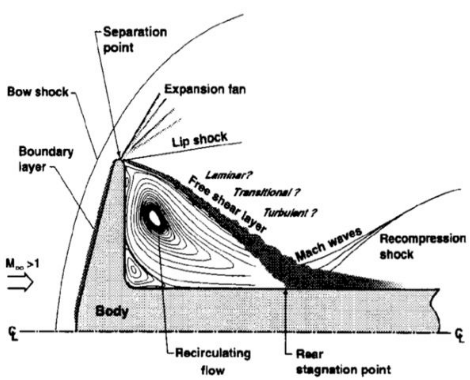

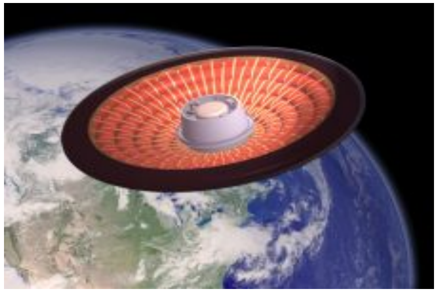

Illustration of a NASA's inflatable heat shield; LOFTID

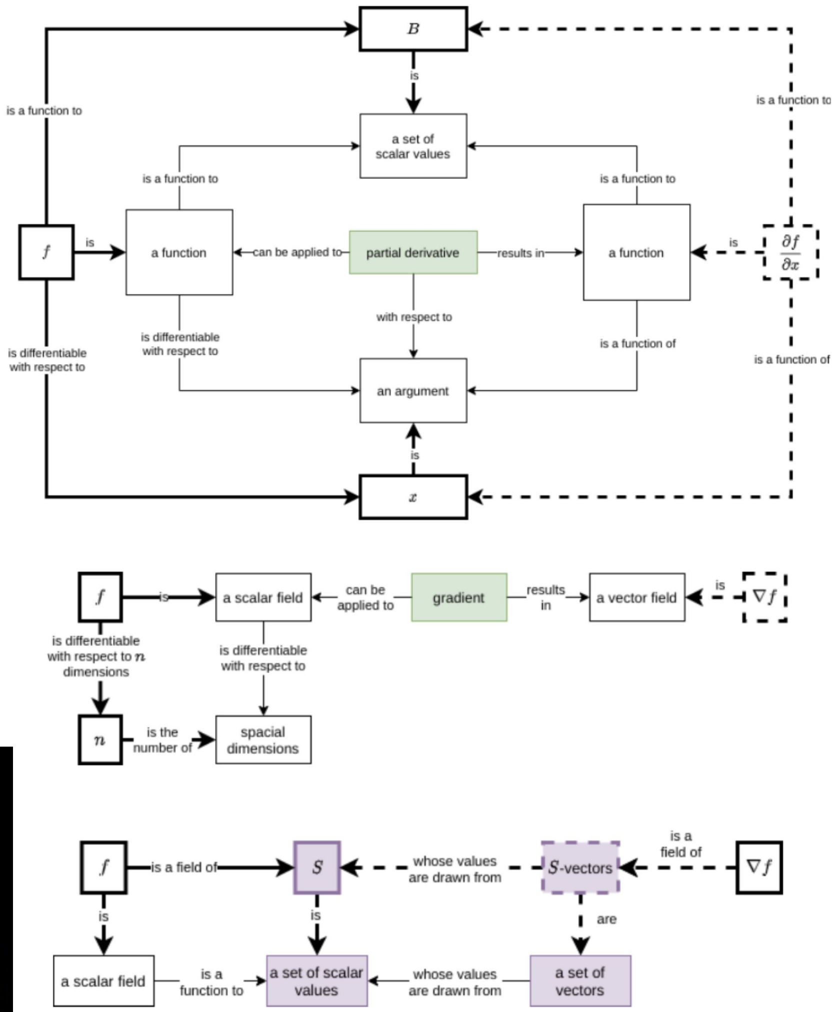

Commander’s Intent and Risk

Give the AI trade-offs in a risk calculus.

- Provide opportunities for incorrect choices, or unknown value comparisons.
  - Strategic guidance specified in terms of desired state over time, and acceptable risk to mission segments.
- Allow for more diversity in responsive design, with real or simulated human-in-the-loop to evaluate alternatives.
  - Partially hidden objective functions can create space for quantitative assessment of symbiosis.

P Sulu creates safety margin that satisfies risk bounds and maximizes expected utility

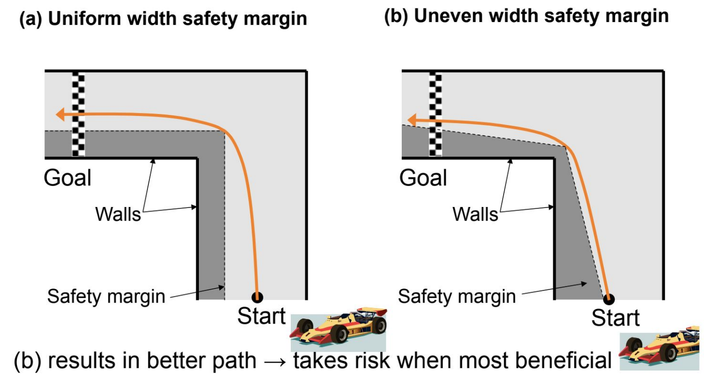

[Ono & Williams, AAAI 08]

Monterey Bay Mapping Example

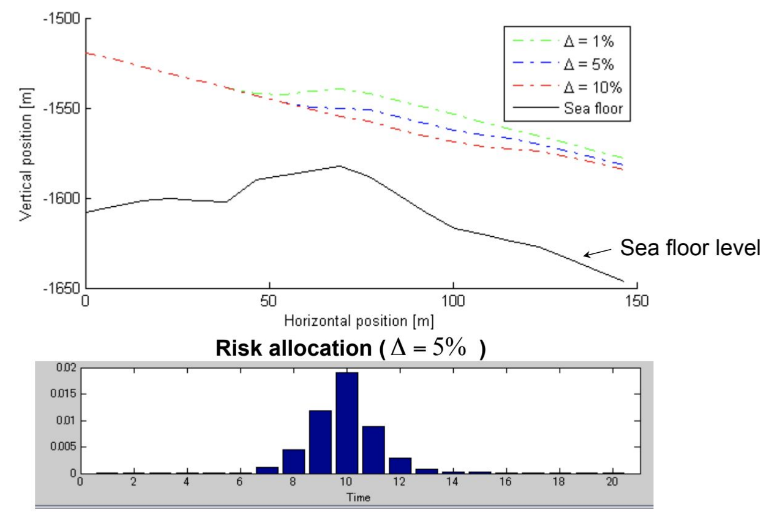

Ono & Williams, AAAI 08

Semantic Tags

Deriving useful semantic tags, labels, and other metadata for use in ORCA can be challenging.

Used existing ontologies (SNOMED), bootstrapped from the domain, establishment of common ground with human co-performers.

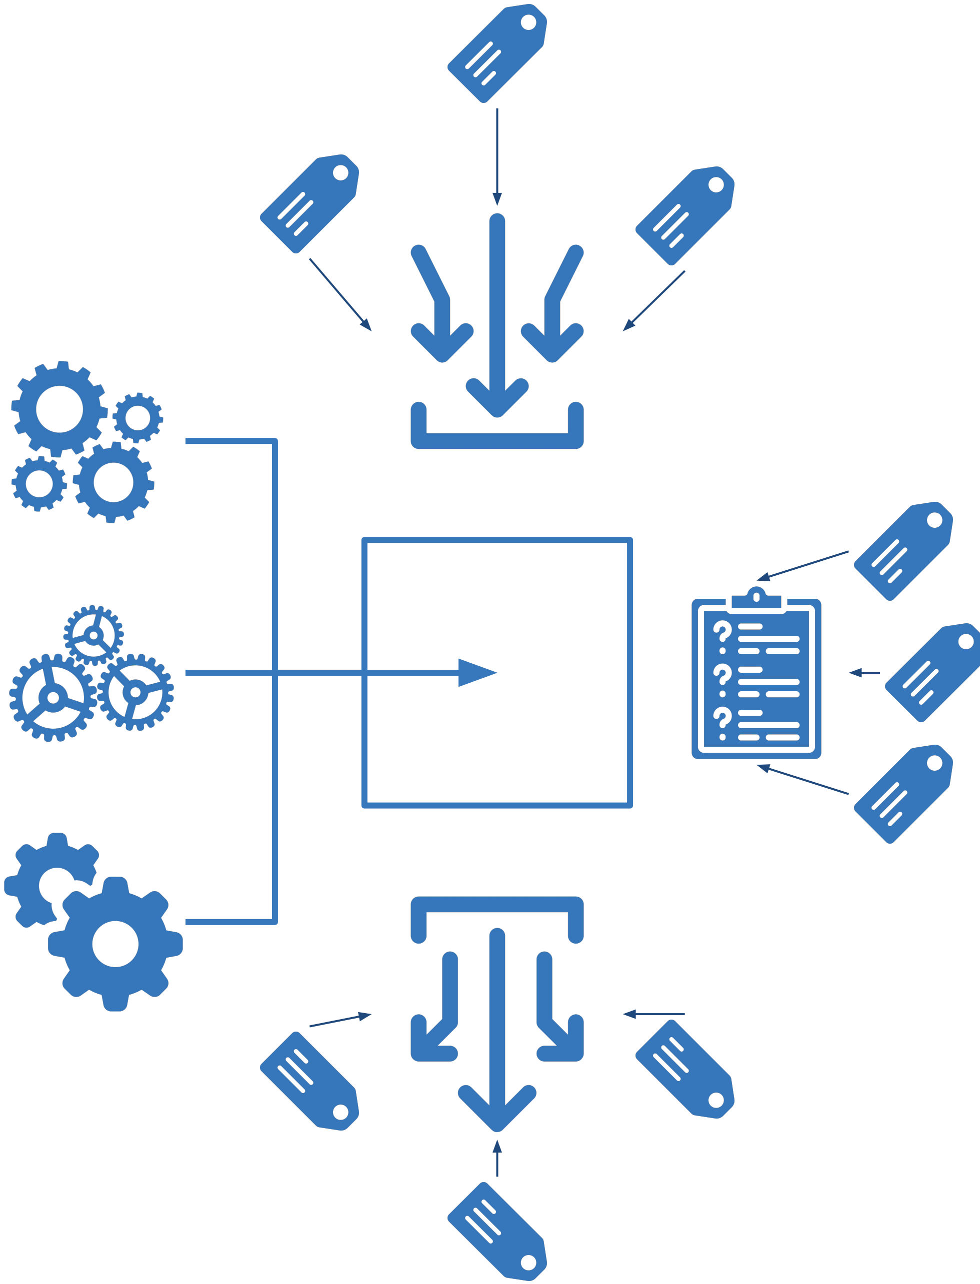

Director/Matcher Experiment

Currently, we have implemented and tested the right hand side of this experiment with a **human** (simulated) director and an **MCP** matcher.

- Supported by corpus (simulate the director through its corpus).
- MCP is not allowed to ask clarifying questions (no data) but sees all director utterances.
- Need to expand experimental framework to support human interaction.

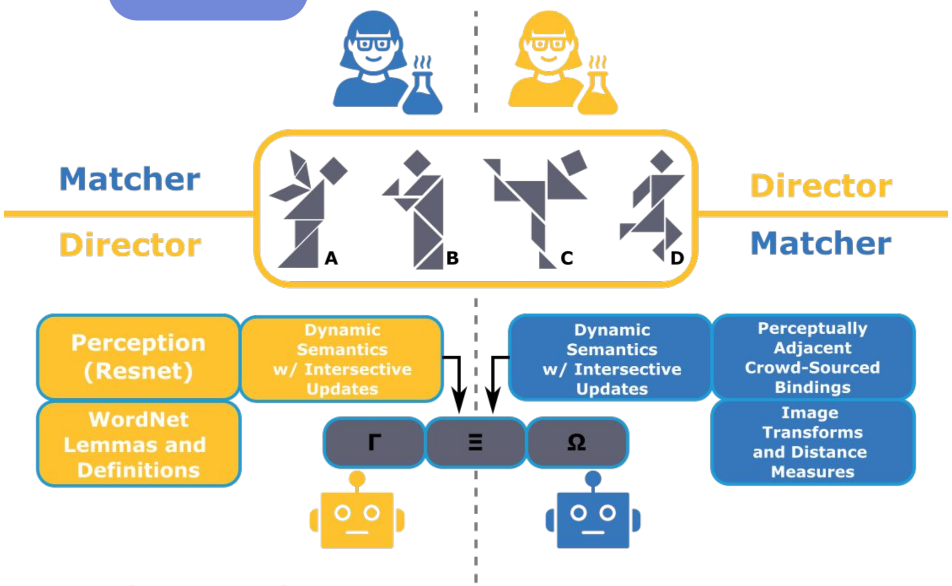

Common Ground

“Lexical entrainment is the phenomenon in conversational linguistics of the process of the subject adopting the reference terms of their interlocutor.”

Common ground is about managing/estimating the Wasserstein barycenters of mutual alignment in a symmetric simplicial set of common ground.

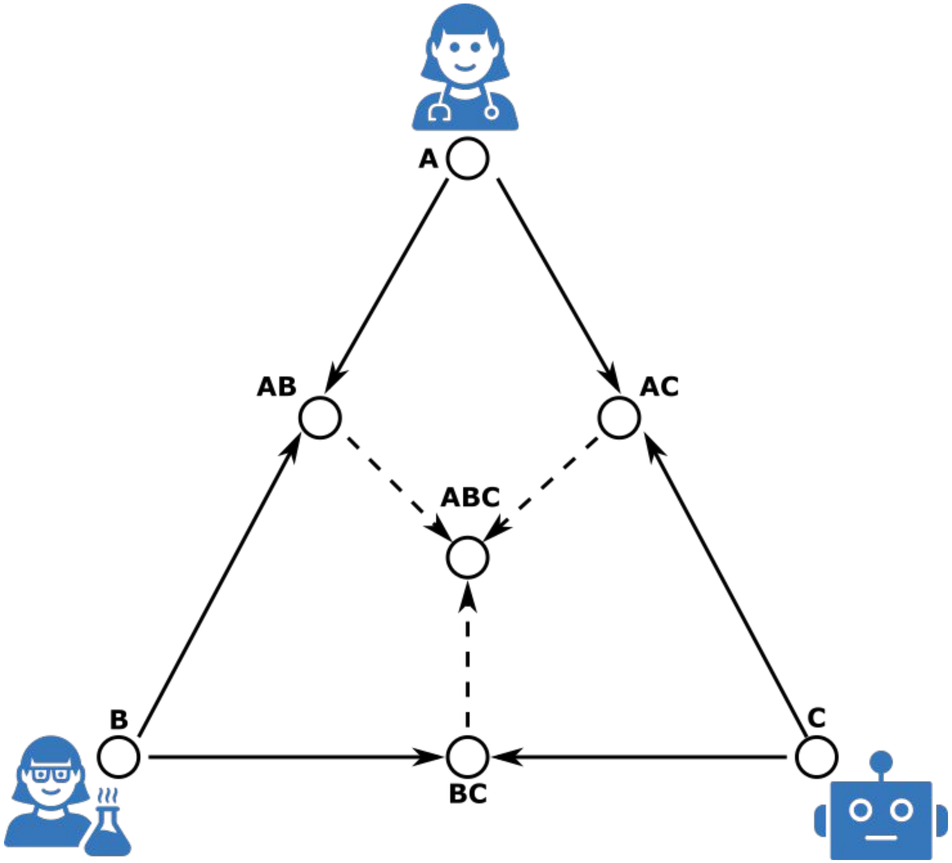

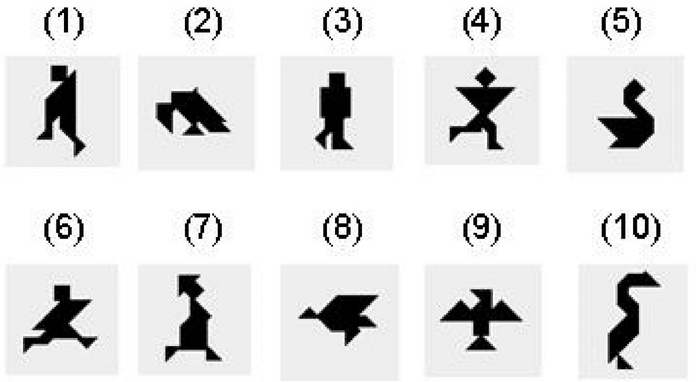

Common Ground

Our initial formalism models the space as follows:

Ξ - The current working ontology of all referents.

Γ - The accepted common ground of object referents.

Ω - Rejected referents.

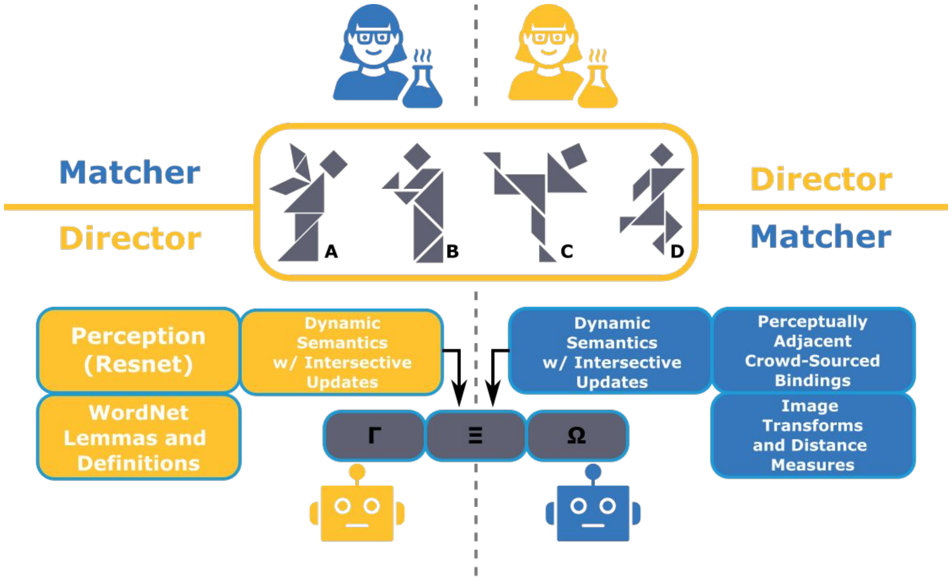

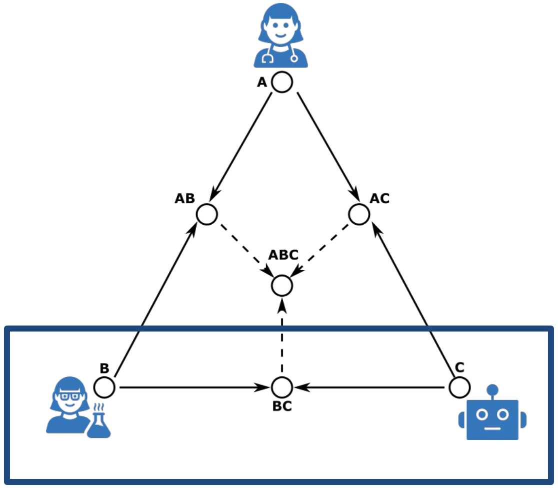

Results and Metrics

Successfully entrains the entirety of the Stanford Corpus Experiment (12,000 independent examples).

Lexically entrains 65% of the time in the first utterance.

top_k	Human	matcher
$k = 1$	0.00	41.66%
$k = 3$	N/A	63.01%
$k = 5$	N/A	83.56%

	Time (ms)		Utterances Needed	
Tangram ID	Human	matcher	Human	matcher[1]
A	31737	1.2	2.5	1
B	21156	7.8	3.75	1
C	15311	3.3	2.5	2.3
D	27794	0.4	2.4	1
E	16614	2.9	2.4	1
F	50496	14.1	2.5	2.3
G	21756	2.1	2.4	2.5
H	26559	1.8	2.4	1
I	37634	2.4	2.4	1
J	37392	2.2	2.4	2.3
K	60380	2.9	4.8	1
L	42110	5.1	2.3	5
Average	32411.58	3.9	2.73	1.78

Common Ground and Ontology Bootstrapping

Text reading of MCO and other similar model files.

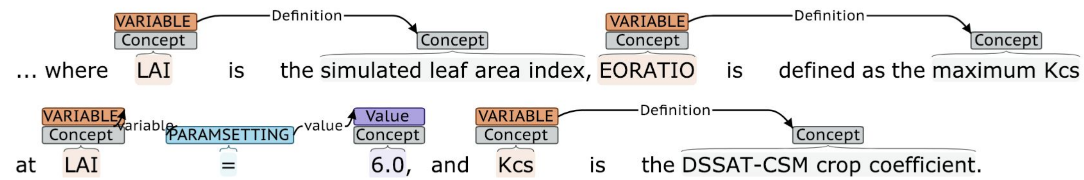

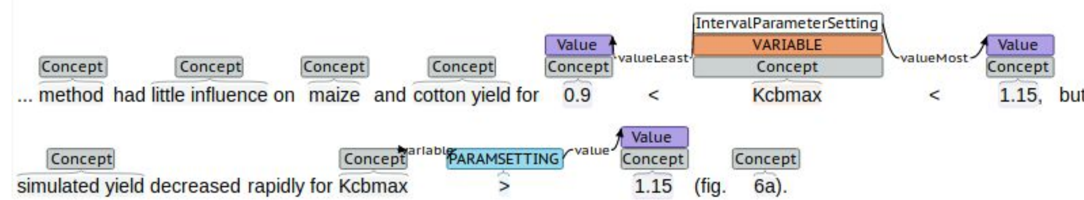

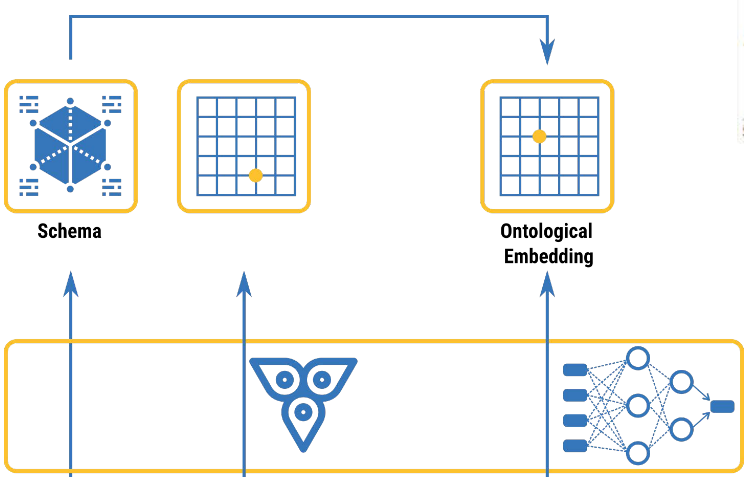

MCP::Subject Engagement

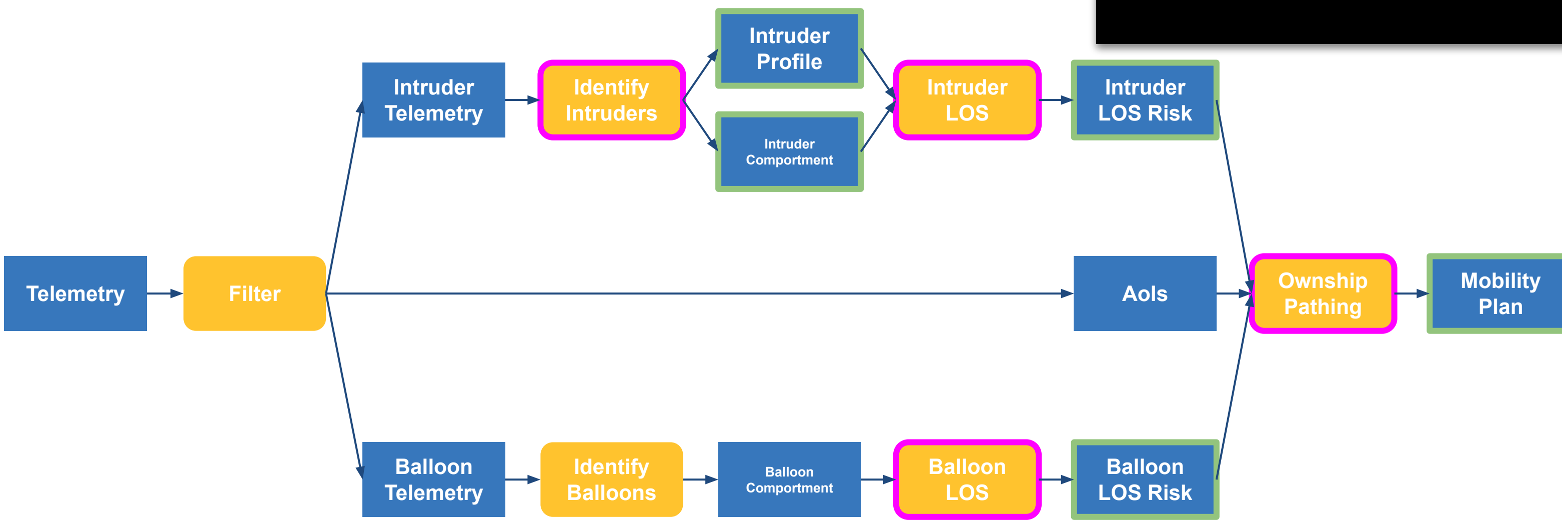

Shift the view to a more comprehensive, interdependence-aware plan.

Allows for delegation in certain areas.

MCP::Subject Engagement

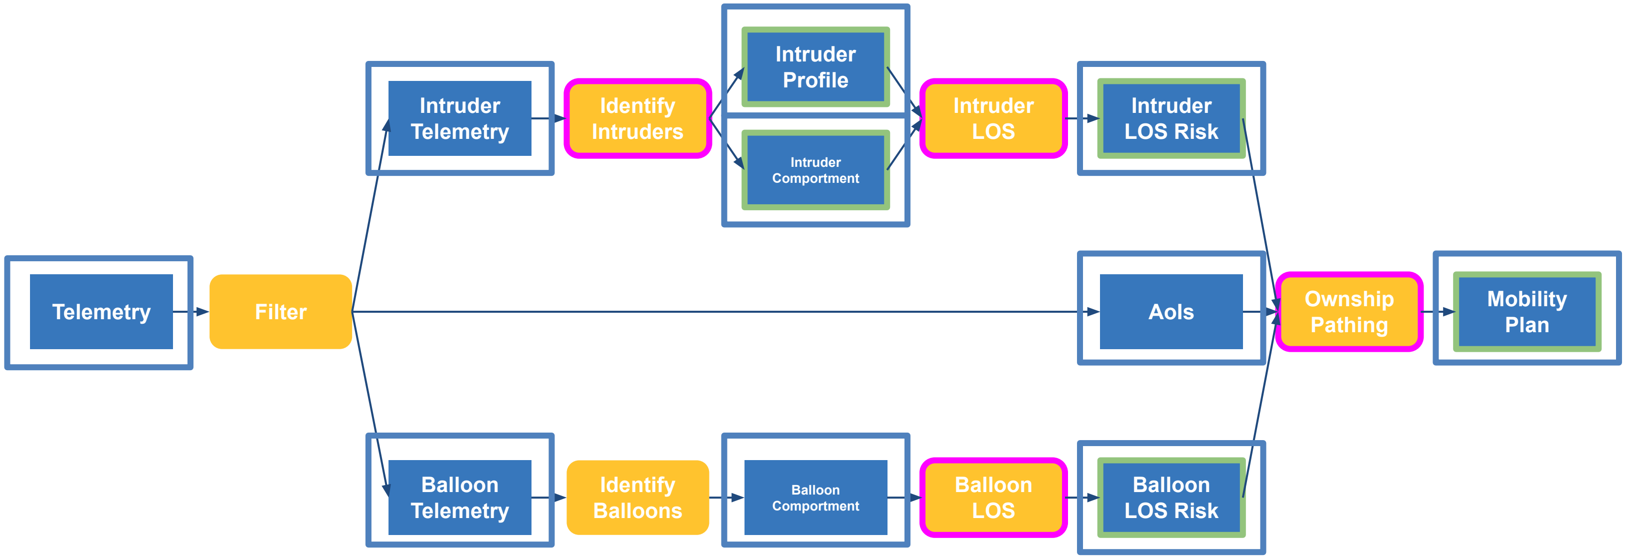

Shift the view to a more comprehensive, interdependence-aware plan.

Allows for delegation in certain areas.

MCP::Subject Engagement

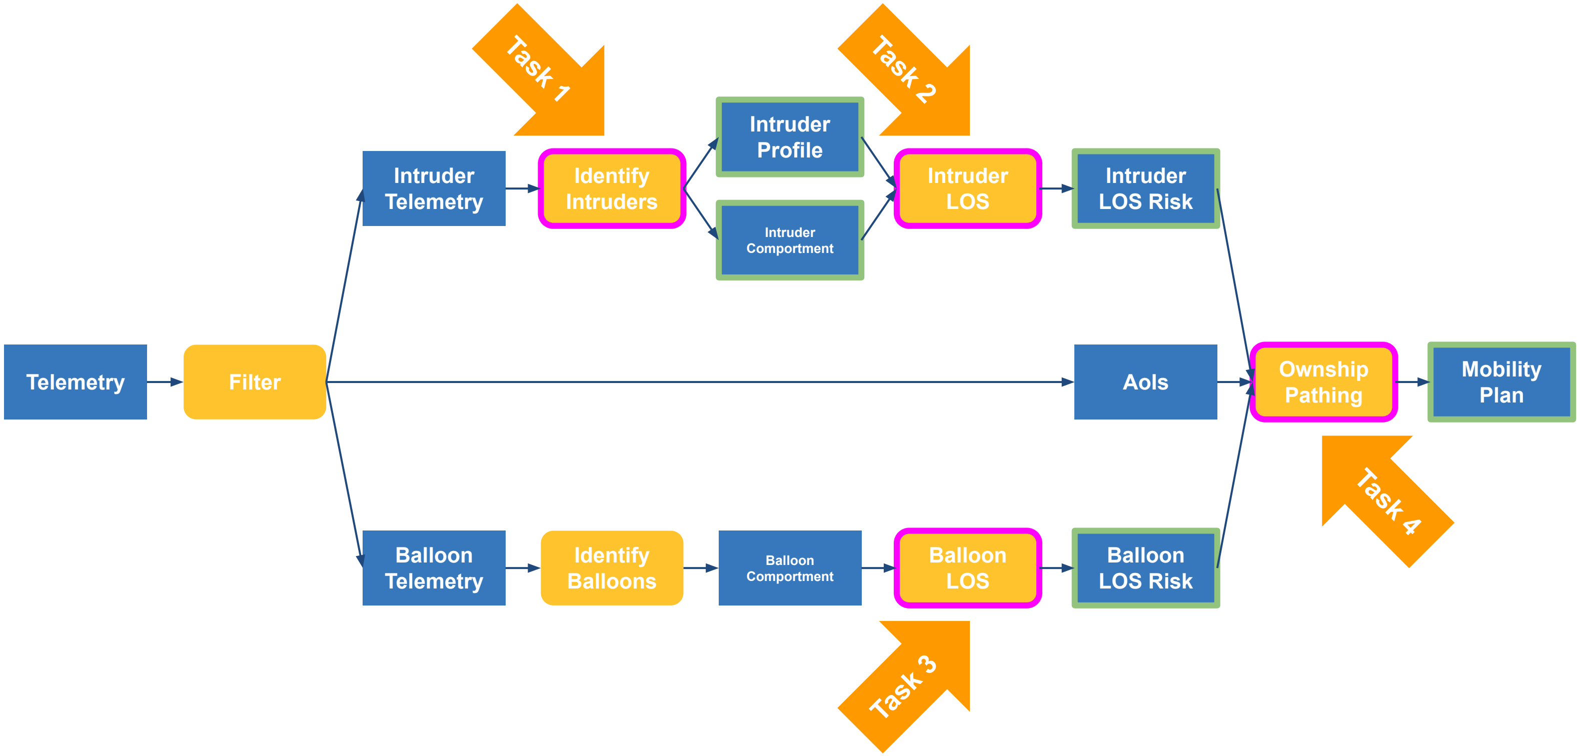

Four tasks in the workflow are under delegation control by the subject.

Delegation turns the task over to the MCP
Intervention turns the task over to the subject

ORCA: Orchestrating Symbiotic Intelligence for Agile and Adaptable, Crisis Response Decision Making

Sam Cowger, Sourya Dey, Matt LeBeau, Ethan Lew, Panchapakesan Shyamshankar, Ted Hille, *Eric Davis**

ewd@galois.com

https://github.com/GaloisInc/deep-koopman

https://github.com/GaloisInc/AMIDOL

This work funded in part by:

DARPA: DARPA-PA-18-02-AIE-FP-039
HR001119S0083-SDCPS-FP-009
HR001118S0057-Office-Wide-BAA-FP-057

AFRL: FA8650-22-C-6420