Neural Networks for Decision AI

Risto Miikkulainen

The University of Texas at Austin and
Cognizant Technology Solutions

Joint work with Olivier Francon, Santiago Gonzalez, Babak Hodjat,
Elliot Meyerson, Xin Qiu, and Hormoz Shahrzad

1

Cognizant

How to Make Good Decisions?

Organizations have lots of data
• Can build predictive models of patients, customers, students…

Such models do not specify how to make decisions
• It is a different learning problem

2

Cognizant

Challenges in Learning to Make Decisions

Optimal decisions not known; domain partially observable; interactions nonlinear
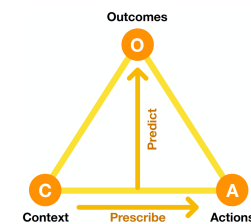• Direct methods like LP, gradient descent are ineffective

Need to search for decision strategies
• But testing strategy candidates in the real world is costly
• Simulators often not available, inaccurate, or costly

3

Cognizant

Surrogate Optimization Approach

Use a predictive model as a surrogate for the world
Train model with historical data: Context+Actions → Outcomes
• Phenomenological model (based on data)
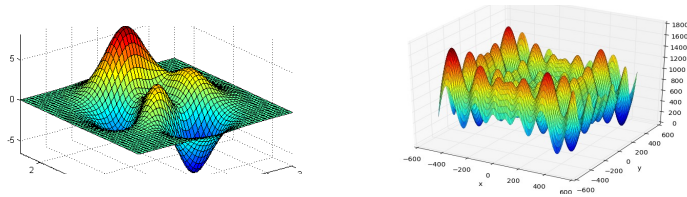• Not a simulation from first principles
Search for a good decision strategy (i.e. policy): Context → Actions
• Use the model to evaluate strategies

4

Cognizant

Evolutionary Surrogate Optimization

Search space is nonlinear, deceptive, large, high-dimensional, multiobjective
- Difficult for Kriging, Bayesian optimization
- Difficult for RL as well, e.g. DQN, A3C, PPO

Surrogate-assisted evolution might work
- But need to evolve a strategy, not a single point
- Evolve neural networks to represent the strategy

Cognizant

ESP: Evolutionary Surrogate-Assisted Prescription

Combine a neural-network surrogate with neural-network strategy
- Train the surrogate, evolve the strategy
- Predictor and Prescriptor neural networks

(Variations: RF, SVM, Linear Predictors; Rule-set Prescriptors)

Cognizant

ESP: Evolutionary Surrogate-Assisted Prescription

Based on components in existing applications:

Cyberag: Evolving growth recipes for vertical farming
- Millions evaluated with surrogate, hundreds planted
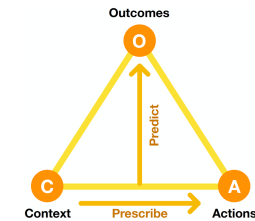- Discovered a 24-hr light preference, size/taste tradeoff

Ascend: Evolving auto-segmented web designs
- A neural network evolved to create different pages for different users
- Discovers and utilizes surprising interactions (e.g. day/provider/urgency)
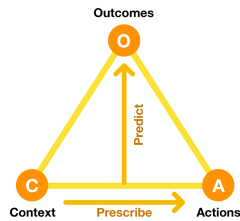
Cognizant

ESP Method: Prescriptor

Find a decision policy that optimizes outcomes
- Expressed as a Prescriptor model Pr (e.g. a NN, ruleset)
- Map contexts to actions: Pr(C) = A
- Optimal A not known: Pr needs to be evolved
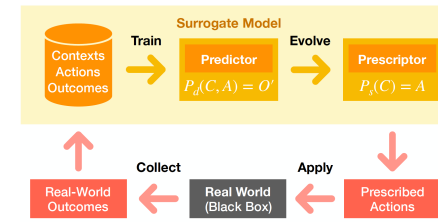- How can we evaluate Pr candidates?

Cognizant

ESP Method: Predictor

Predictor model Pd trained with historical data
- $Pd(C,A) = O$
- Can be used to evaluate each Prescriptor candidate
- Can be multiobjective
- Can be a neural net, random forest, SVM etc.
  (in special cases: a simulator, or the real world)
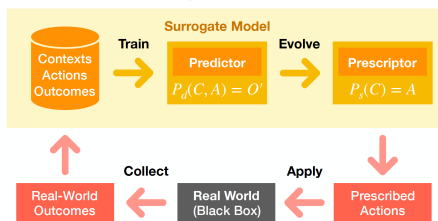
Cognizant

Static ESP Process

Can be learned in one step, or in an incremental outer loop
- Train Predictor with historical data
- Evolve Prescriptors with the Predictor
- Apply the best Prescriptor to the world
- Collect new data
- Repeat

Cognizant

ESP Process on Sequences

Many iterations of the outer loop
- Apply Prescriptors to the world; Collect C, A, O
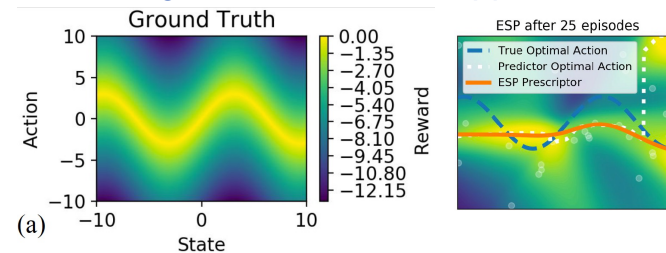- Train the Predictor; Evolve Prescriptors
- Repeat

Surprising synergy emerges:
- Incremental co-learning regularizes both models!
- Results in automated curricular learning

Cognizant
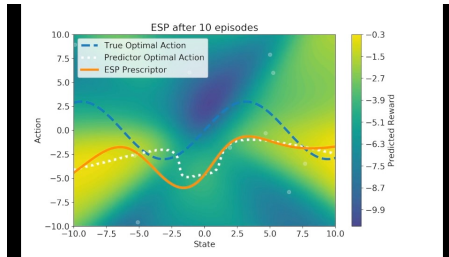

Illustrating ESP in Function Approximation

Ground truth known; easy to visualize
- Samples
- Predictor
- Prescriptor
- Optimal actions for Predictor
- Optimal actions for ground truth

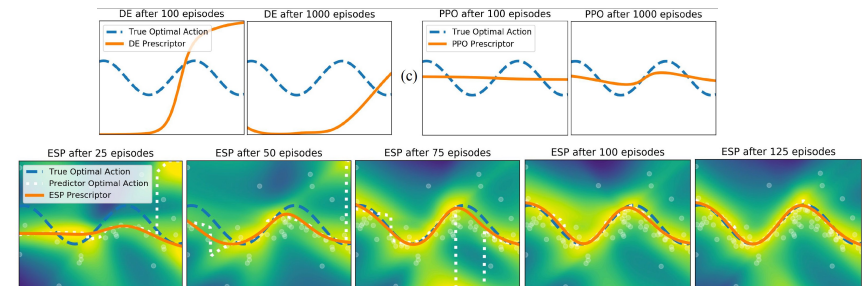**Cognizant**

ESP Learning Process

ESP after 10 episodes

Predictor approaches ground truth
Prescriptor approaches ground-truth-optimal actions
Prescriptor is closer to ground-truth-optimal than Predictor-optimal
• Prescriptors regularize!
• By ensembling approximate Predictors
• (Demo)

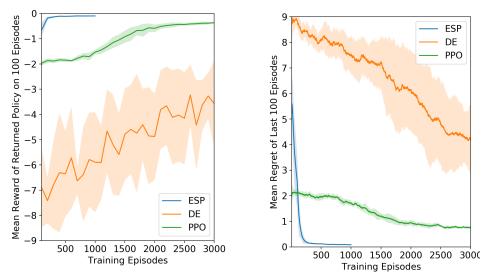**Cognizant**

Qualitative Comparison of ESP, DE, and RL

Direct Evolution (DE) using real-world evaluations instead of the surrogate
PPO is SotA in continuous RL
DE, PPO need many samples in the real world
• Even with 10x samples, still have not learned well

Cognizant

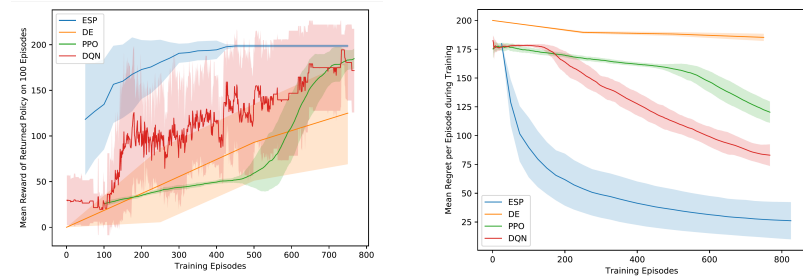# Quantitative Comparison of ESP, DE, and RL

ESP is faster, more accurate, has lower variance
• i.e. is more reliable
ESP has lower regret
• i.e. has lower cost, is safer
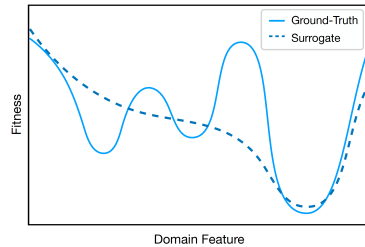
Cognizant

Standard RL Benchmark: CartPole-v0

OpenAI Gym benchmark with known DQN and PPO implementations

Similar results as in function approximation
- ESP learns faster, has better solutions, lower variance, lower regret

Cognizant

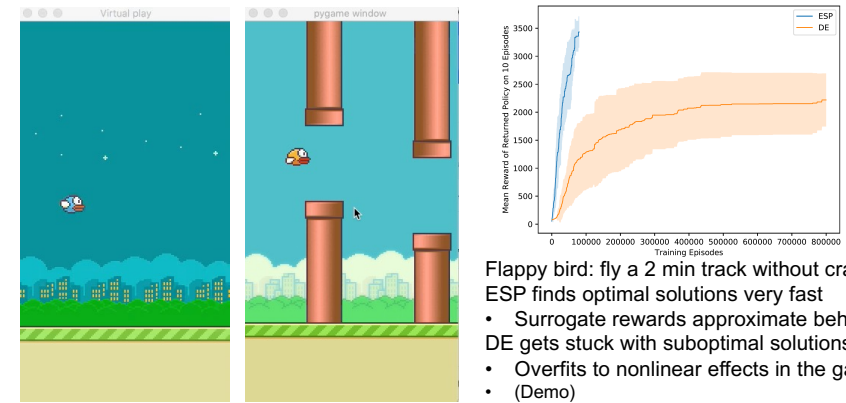## Regularization Through Surrogate Modeling

1. Ensembling of approximate Predictors -> regularized Prescriptors

2. In more complex domains, Predictors regularized as well
   - Incomplete learning of complex landscapes
   - Results in automatic incremental evolution, i.e. a curriculum

Cognizant

Demo of Surrogate Regularization
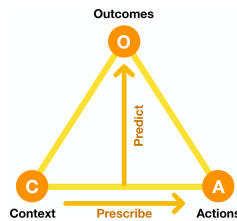

Flappy bird: fly a 2 min track without crashing
ESP finds optimal solutions very fast
- Surrogate rewards approximate behaviors
DE gets stuck with suboptimal solutions
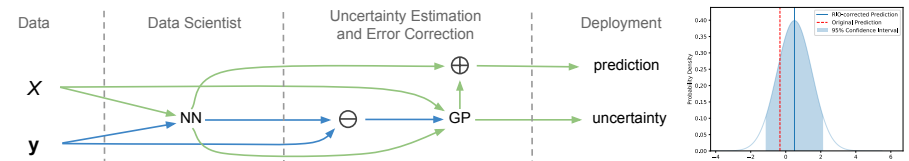- Overfits to nonlinear effects in the game
- (Demo)

Cognizant

Making Decisions Trustworthy: Estimating Uncertainty

- ESP predicts value, cost, side effects…
- How can we trust it?
- It needs to estimate uncertainty of its predictions

Cognizant

RIO: Residual Estimation with Input, Output Kernel

- Train a Gaussian Process model to predict residuals (i.e. signed errors)
  - Based on model input, output, and labels
- Provides estimates of uncertainty in model output
- Provides corrections to model output

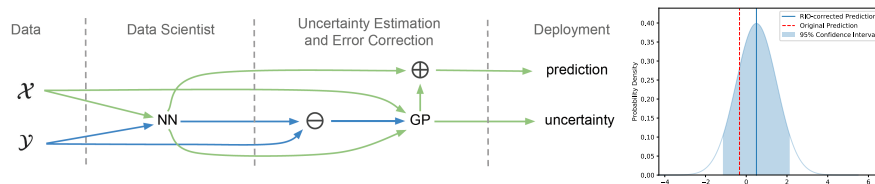- Applies to ESP Predictor (and any model trained with labeled data)

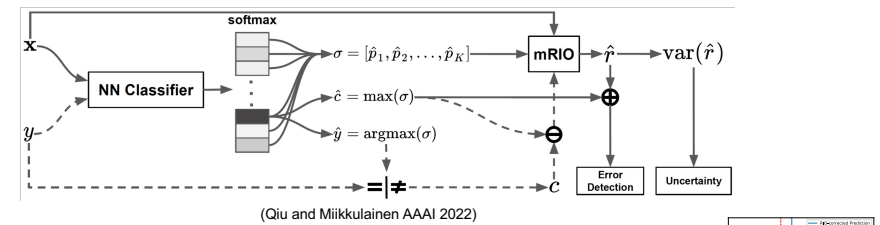**Cognizant**

Why does it work?

- Why is RIO better than NN alone or GP alone?

- NN is expressive (i.e. high variance)
  - Learns structure that GP would treat as noise
- Remaining structure is easier to learn
  - GP can capture part of it
  - GP is more regular than NN
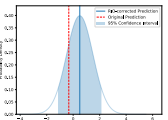
Cognizant

RED: Residual-based Error Detection

(Qiu and Miikkulainen AAAI 2022)

Extension of RIO to classification tasks
- Softmax output of the NN estimates class probabilities
- Form a GP model of softmax residual errors: $\hat{c}'_* \sim \mathcal{N}(\hat{c}_* + \bar{\hat{r}}_*, \mathrm{var}(\bar{\hat{r}}_*))$
- Use the model to identify likely classification errors:
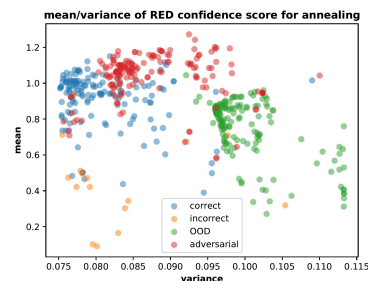 - Does the model have a different argmax?
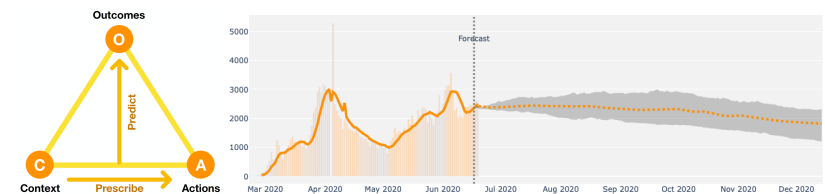
Cognizant

RED Analysis of Error Types

- Mean and variance of confidence score is different for correctly classified, misclassified, OOD inputs, adversarial inputs

Cognizant

Example 1: ESP for Optimizing NPIs for COVID-19

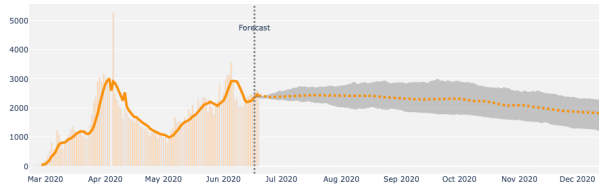From prediction to prescription: A new role for AI
- Current models predict outcomes of given NPIs
- Given a desired health/cost balance, ESP prescribes NPIs

Not just what will happen, but what we should do about it

Cognizant

## First build a predictor with available data

- Oxford dataset of cases and NPIs
  - (Hale, Webster, Petherick, Phillips, Kira (2020). Oxford COVID-19 Government Response Tracker)
  - Daily case numbers, NPIs (eight kinds with 2-4 levels of stringency)
- Standard compartmental models
  - Based on epidemiological assumptions on S, I, R
  - Require setting unknown interaction parameters
- Data-driven modeling is phenomenological
  - Requires no assumptions; includes all the interactions
  - Requires sufficient data

Cognizant
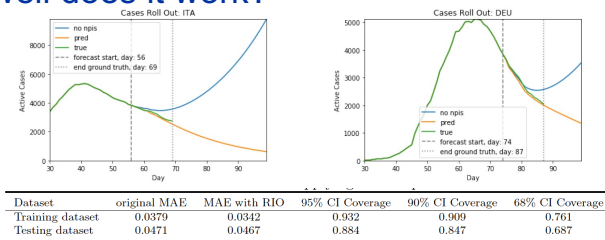

Predictor Design
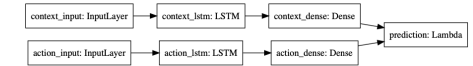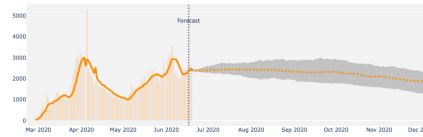
- Train a Deep Learning (LSTM) model to predict new cases
  - Predict transmission rate; 7-day moving avg; scaled by susceptible population size
  - Separate paths for context (cases) and action (NPI stringency)
    - Possible to ensure stringency has a monotonic effect
  - Predict based on 21-day past data; rollout 180 days into the future
- Estimate confidence with RIO
  - Monte Carlo rollouts from RIO-determined distribution
  - Plotted as quartiles (i.e. box-plot base over time)

Cognizant

How well does it work?

| Dataset	original MAE	MAE with RIO	95% CI Coverage	90% CI Coverage	68% CI Coverage
Training dataset	0.0379	0.0342	0.932	0.909	0.761
Testing dataset	0.0471	0.0467	0.884	0.847	0.687

- Surprisingly accurate rollouts even with limited data
  - 20 countries with highest number of cases
  - Several metrics over 14 days
  - More accurate than MLP, RF, SVR, Linear baselines (e.g. MAE 0.42 vs. 2.47, 0.95, 0.71, 0.64)
- Reasonable confidence intervals
  - Randomly selected 14 days across 20 countries as test data
  - Allows correcting the predictions slightly

Cognizant

Then evolve a Prescriptor with the Predictor as a surrogate
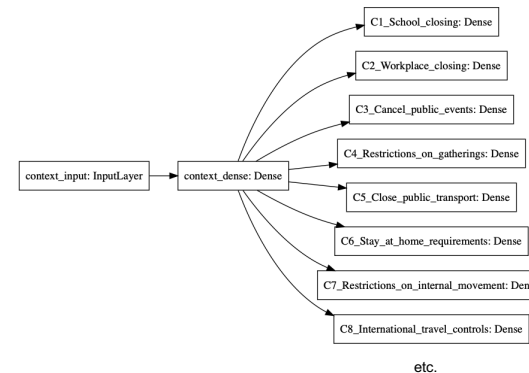
No gradients!
- Need to search for a good model

Evolve neural network prescriptors
- Population-based search
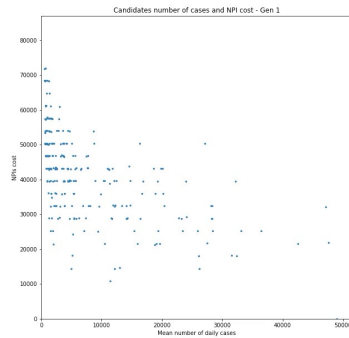- Use Predictor to evaluate each one

Multiobjective
- Minimize future cases
- Minimize stringency of NPIs (a proxy for economic cost)

Cognizant

Solutions represent different tradeoffs

Candidates number of cases and NPI cost - Gen 1

Evolution minimizes both objectives
Results in a Pareto front:
- Some minimize cases
- Some minimize cost
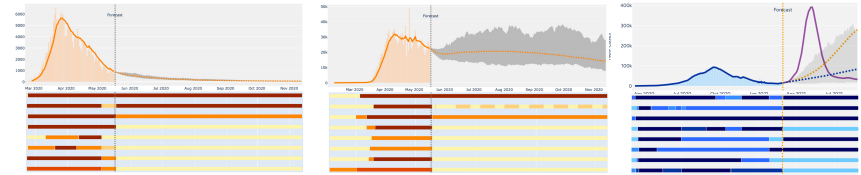- Others balance the two to different degrees

Given a desired balance, best possible solutions
- Empowering human decision makers
- (Demo)

Cognizant

Discoveries on Prediction and Prescription

Highlights (often 2 weeks in advance):
- May 2020: Focus on schools and workplaces (i.e. indoors); alternation
- Sept 2020: Focus on gatherings, travel restrictions; open schools
- Delta surge: India (March 2021); others with low rates (July 2021)
- August 2021: Recommendations for schools (Iceland)
- Dec 2021: Missed omicron surge; it happened everywhere at once
- March 2022: Impact of masking

Cognizant

Interactive Demo
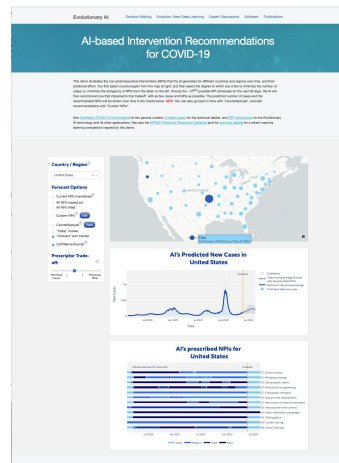
https://evolution.ml/demos/npidashboard
Updated daily May 2020 – Dec 2022

Can be used to obtain data-based recommendations
- Select country
- Select date (current, or past)
- Select health/economy tradeoff (cases vs. NPIs)
- Obtain NPI recommendations and case predictions

Trustworthiness:
- Obtain confidence bounds
- Design custom NPIs with a scratchpad
- Evolve explainable rulesets (instead of neural networks)

Cognizant

Obtaining Recommendations

Select
- Country
- Date (current/past)
- Tradeoff

Observe cases
- Actual
- With current NPIs
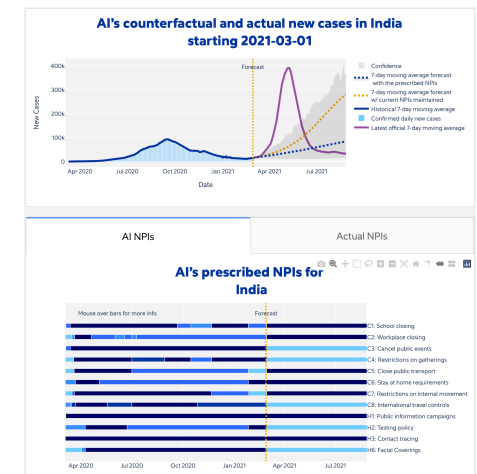- With AI NPIs
And confidence bounds
- GP model
- Multiple rollouts

Observe NPIs
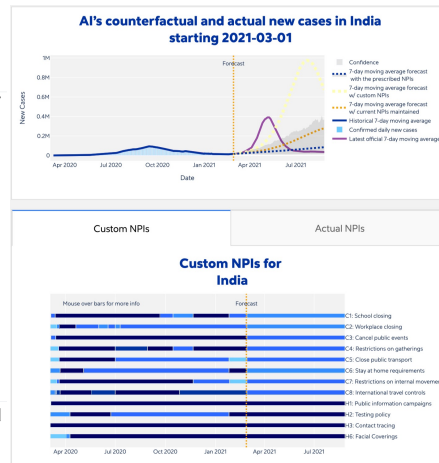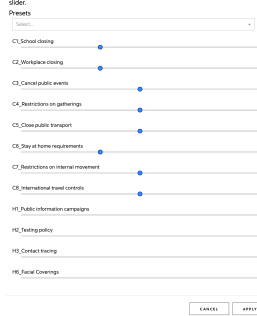- Actual
- Prescribed

Cognizant

Explore alternative NPIs

Start from
- AI's prescriptions
- Current
- Maxed-out
- No NPIs

For instance
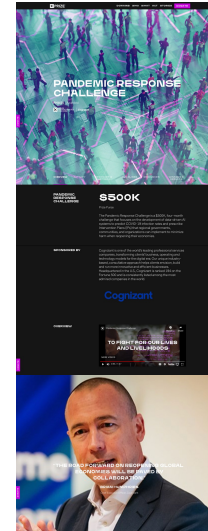- Less school, work, home, transport
- More masks, tests, internatl. travel

Observe results
- A different way to achieve the same result?
- Not quite ☹

Cognizant

XPRIZE Pandemic Response Challenge

November 2020-March 2021
- Phase 1: Prediction Accuracy; 100 teams narrowed down to 50 finalists
- Phase 2: Prescription Effectiveness; 2 winners (Valencia, Slovenia) and 8 runner-ups

Many Machine Learning approaches as well as conventional ones
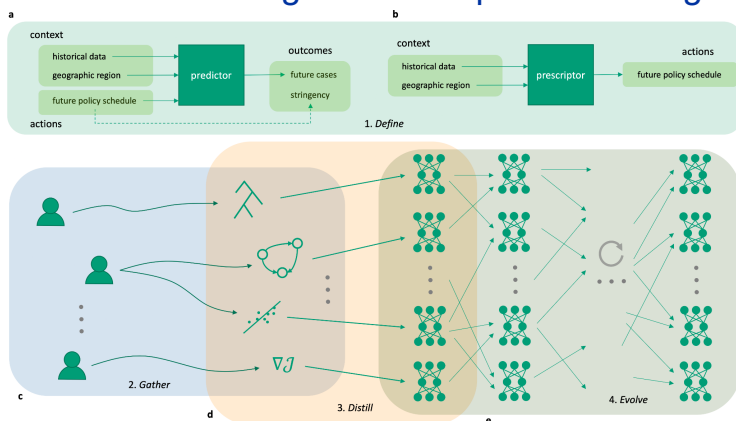- Evaluation and analysis: https://evolution.ml/xprize

Significant impact
- 1.2B views; Informed policy in Valencia, Iceland
- UN/ITU, GPAI/OECD follow-ups

169 expert-designed solutions
- Can we combine them and find even better solutions?
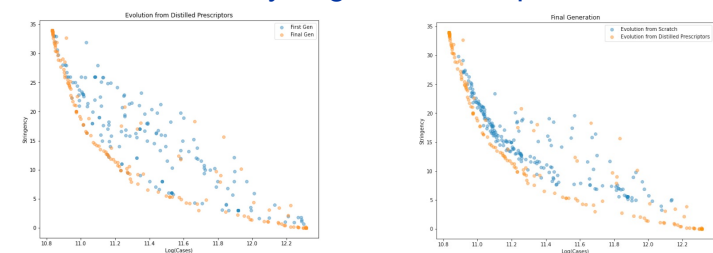- Can neuroevolution improve upon them further?

Cognizant

RHEA: Realizing Human Expertise Through AI

1. Define the problem formally

2. Gather expert solutions

3. Distill them into neural networks

4. Evolve the neural networks

(Meyerson et al. 2023)

Cognizant

RHEA Discovers Synergies in the Expert Solutions

XPRIZE entries have many useful, diverse ideas
- Neuroevolution can be used to recombine and refine them
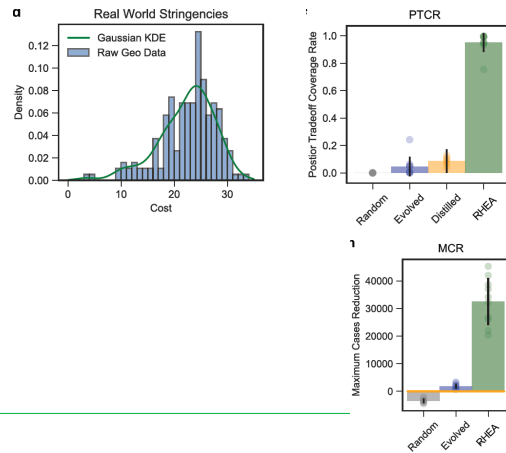
Neuroevolution can also discover new ideas

RHEA improves upon XPRIZE entries
RHEA improves upon evolution from scratch

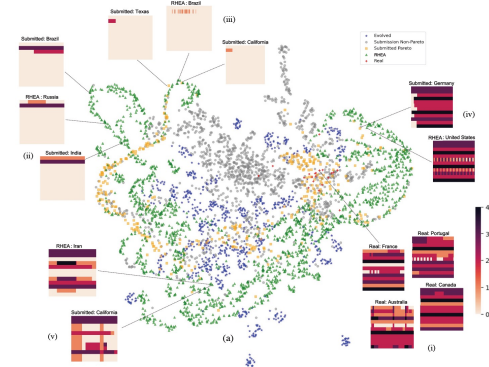**Cognizant**

How good are RHEA's solutions?

- Given human preferences for certain stringencies, the best tradeoffs come from RHEA

- RHEA reduces number of cases maximally for the same stringency level (compared to Distilled)

Cognizant

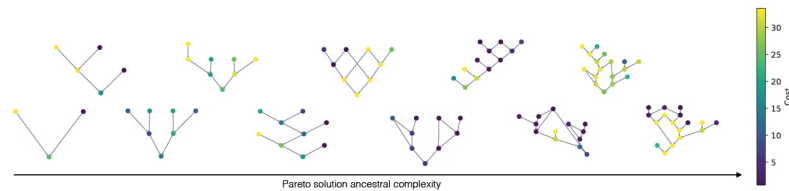## Visualizing the Solutions

A UMAP projection of 90-day policies

Expert Pareto is on a 1-D manifold (yellow)
- Organized by stringency
RHEA (green) forms arcs around it
- Elaborates, interpolates, expands on them
Evolved from scratch (blue) are scattered
Real (red) clustered in the middle
- Some are agile, periodic (i); France, Portugal

Example interpolation: (ii), (iii)
Trade swing+separability for
  agility+periodicity: (iv) and focus (v)

Cognizant

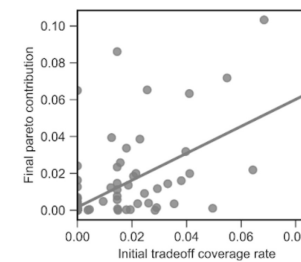## Where do the solutions come from?

Visualizing the ancestry (leaves) of Pareto-front solutions (root)
- A variety of experts used, with different cost
- A variety of structures, from few experts to many
- Child cost usually between those of parents
Evolution is working as expected

Cognizant

Realizing the Potential of Diversity

How much "DNA" does each solution contribute to the Pareto front?
- Good solutions generally contribute more
- Some poor solutions make outsize contributions
Diversity is fundamental in problem solving
- Can realize latent potential hidden in poor entries
- Technology to bring the community effort together

Cognizant

Example 2: Land-Use Optimization

(Globe Observer, 2023)

Global Carbon Budget Imbalance (Friedlingstein et al. 2023):

$$B_{IM} = E_{FOS} + E_{LUC} - (G_{ATM} + S_{OCEAN} + S_{LAND})$$

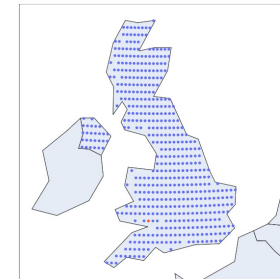Emissions due to land-use change (ELUC) is a major factor
- How much allocated for forest, crops, pasture, range, urban…
- Different amount of carbon release/capture

Optimize to balance carbon emissions vs. economy

The first project of Project Resilience
- A UN/ITU backed "AI for Good" initiative
- Follow-up from XPRIZE Pandemic response

AI and Data Commons
Cognizant

Data Source 1: Historical Land Use

Region	United Kin... ▾
Lat	51.625 × ▾
Lon	-3.375 × ▾
Year	2021 ▴▾

Primary: Vegetation that is untouched by humans
- primf: Primary forest
- primn: Primary nonforest vegetation

Secondary: Vegetation that has been touched by humans
- secdf: Secondary forest
- secdn: Secondary nonforest vegetation

Urban
- urban: Urban areas

Crop
- c3ann: Annual C3 crops (e.g. wheat)
- c4ann: Annual C4 crops (e.g. maize)
- c3per: Perennial C3 crops (e.g. banana)
- c4per: Perennial C4 crops (e.g. sugarcane)
- c3nfx: Nitrogen fixing C3 crops (e.g. soybean)

Pasture
- pastr: Managed pasture land
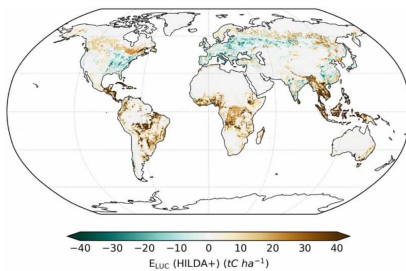- range: Natural grassland / savannah / desert / etc.

Land-use harmonization project (LUH2) (Hurtt et al. 2020)
- Cells with 0.25x0.25 degree resolution
- Annually 850-2022

AI and Data Commons
Cognizant

Data Source 2: Carbon Emissions from Different Uses

E_{LUC} (HILDA+) ($tC\ ha^{-1}$)

Bookkeeping of Land Use Emissions (BLUE) (Hansis et al. 2015)
- High-fidelity simulation
- Estimates long-term emissions resulting from land-use change (ELUC)
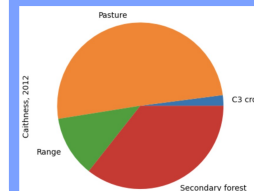
Too slow to run directly as a surrogate
- Prepared a dataset for 1850-2022 by sampling the simulator

AI and Data Commons
Cognizant

Setting Up ESP for Land-Use Optimization

For a given cell and year, what are the smallest changes we can make to reduce emissions as much as possible?

Context
- Cell, area, year
- Land use

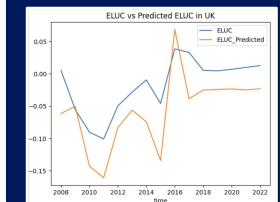### Actions
- Changes in land use

C3 crops diff	0.000283
Pasture diff	-0.003074
Range diff	-0.000801
Secondary forest diff	0.003570
Urban diff	0.000000

Outcomes
- Emissions
- Change amount

Cognizant

Training Predictive Models

Model	Time (s)	EU	SA	US	Global
LinReg (EU)	0.047	0.033	0.172	0.169	0.206
LinReg (SA)	0.457	0.137	0.153	0.061	0.110
LinReg (US)	0.331	0.139	0.146	0.035	0.073
LinReg (Global)	4.644	0.139	0.150	0.035	0.074
RF (EU)	17.697	0.064	0.211†	0.161†	0.218†
RF (SA)	209.688	0.133*†	0.071*†	0.074†	0.126†
RF (US)	111.701	0.163	0.185†	0.032*	0.094†
RF (Global)	417.647	0.041*†	0.076*†	0.028*	0.045*†
NeuralNet (EU)	10.711	0.025*†	0.277	0.286	0.334
NeuralNet (SA)	103.696	0.248	0.100*	0.562	0.399
NeuralNet (US)	73.141	0.136†	0.225	0.024*†	0.150
NeuralNet (Global)	1649.193	0.046*	0.110*	0.025*†	0.050*

Bold: the best model in each region; *: 99% confidence the model outperforms LinReg;
†: 99% confidence that RF outperforms NeuralNet, or vice-versa
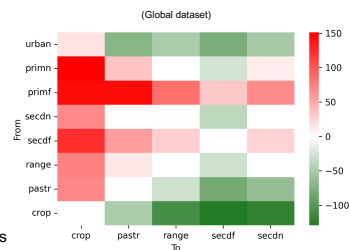
Evaluated linear regression, random forest, neural network models
To keep computations feasible:
- Separate models for Global, Europe, South America, US
- NN, LinReg trained with 1851-2011, RF 1982-2011; tested with 2012-2021

LinReg not sufficient
- Apparently a nonlinear problem
RF does not extrapolate well
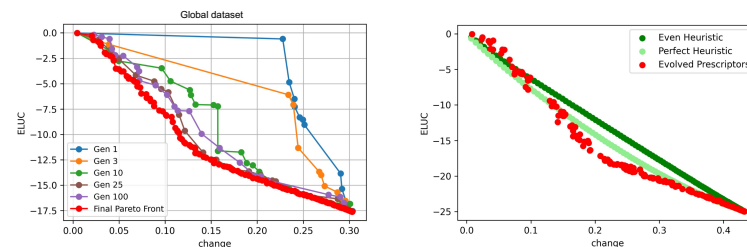Neural networks are the most accurate
- Learn positive and negative changes
- Modulated nonlinearly based on location, area, year

Evolving Prescriptive Models

A population of neural networks recommending land-use change for a particular cell and year
- Evolved against the neural-network predictor
- Trained with a random subset of 1851-2011, tested with 2012-2021

Results in a Pareto front for ELUC/Change tradeoffs
- Better than changing to forest equally, or linearly optimally
- Although worse in almost every location???
- Discovered a better principle: Pick your battles!

Interactive Demo

Available at https://landuse.evolution.ml
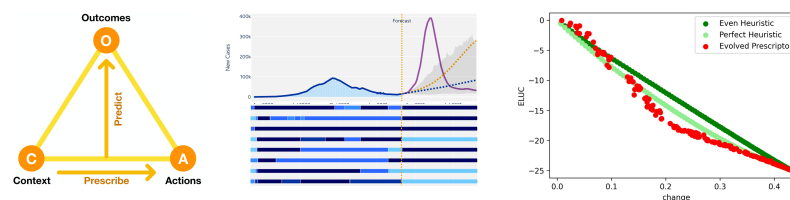- Explore different locations and time periods
- Observe actions and modify them
- See their outcomes

Future Opportunity: An AI-Enabled Society

Apply ESP to decision making in society
- Humans set the goals; ESP determines how to best achieve them
Can optimize productivity, cost, environmental impact, equality
- Free of human biases and agendas
- We can make society what we want it to be

A concrete first step: Project Resilience
- An AI-for-Good ecosystem with open source and data
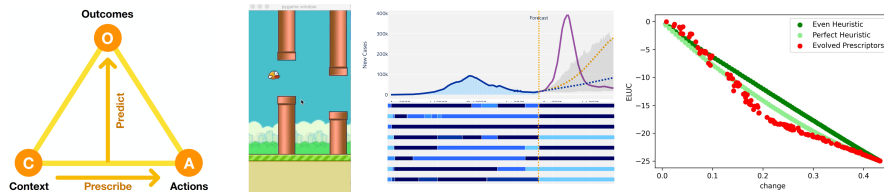- Backed by UN/ITU, GPAI/OECD, AI Commons, Cognizant
- Volunteers needed!

Conclusion

ESP utilizes a surrogate, evolution to do well in decision-making tasks
- Sample-efficient, reliable, low-cost and safe
- Intrinsic regularization in both in the Predictor and the Prescriptor

A demonstration of the power of data and machine learning
- In building Predictors, i.e. forecasting outcomes
- In building Prescriptors, i.e. making recommendations

Trustworthy: Estimating uncertainty, exploring alternatives; explaining decisions with rules

Creative decision making in healthcare, business, education, society…

Cognizant