AI for Decision Making under Uncertainty

Application to Vision-Based Landing,

Aerospace Manufacturing & Earth Observation tasks

Emmanuel Benazera *, Guillaume Infantes,

Antoine Jacquet, Pierre Pereira, Nicolas Meuleau, Louis Jean

04/04/2024

* speaker

Company & Activities

Open Source Software
DeepDetect: Open Source Deep Learning Server
joliGAN: Automated data domain adaptation
Wheatley: Reinforcement Learning for Planning & Scheduling
Parasol: Big Data Visualization
LibCMAES: Black-Box Optimization

Datasets
100To+ of ready-to-use datasets (images, texte, cyber, ...)
Pre-trained models for a range of industries

Compute
50+self-hosted GPUs with ~1Pflop, cloud as supplement
Embedded models, hardware and software

Services
ML/DL in Production: Text, Image, Data, Audio, ...
Deep Learning Cloud & Appliances
AI Consulting

Applications
Content & Media
Cyber-Security
Transportation
High tech & Hardware
Medical & Sat Imagery
Art, Luxury, ...

R&D
Adversarial & Generative Models
Neural Combinatorial Optimization
False Positives Minimization
Neural Search Engines
Small Print Neural Nets
Homomorphic Encryption

Customers

Personnels
2023: 12 persons (7 seniors, 5 juniors)
Artificial Intelligence: 6 PhDs with > 10 to 20y XP, top track record, AAAI, IJCAI, ECAI, IEEE, JAIR.
Maths/Physics: 1 PhD with 10y XP
Interns: ISAE, Polytechnique, INSA, Dauphine, Paris8, ENSEIIHT, Epitech, TSE, ..
Academic & Industrial Pool: NASA, Inria, CNRS, Onera, IRIT, Motorola, Apple, RedHat

Academia & Ecosystem
Collaboration with top-most institutions: ISAE, ONERA, INRIA, CNRS, IRT-SystemX, Confiance.AI, IRT St-Ex
International: NeurIPS, ICML, HN, ML Reddit, ...
Community: Toulouse&Paris, Open Source, CCC
Chart: R&D, no competition, commodifying ML, information sharing, excellence & openness

Legal: 77 rue Pargaminières, 31000 Toulouse / SIRET: 818 893 414 00016

Jolibrain Intervention tracks

Critical & Efficient AI Applications

- Critical precision & speed applications
- Defects & rare events detection
- Multi-models AI Applications

Airbus Vision-based Landing (2023)

Airbus painting cracks scanner detection (2023)

Generative AI Industrial Applications

- Dataset smart augmentation
- Generate rare & corner cases
- Augmented reality
- Enhances simulation to real

SNCF sim2real catenary laser-based metrology data (2020)

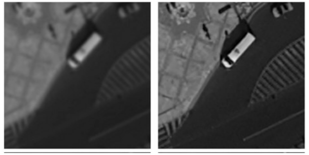

Airbus DS super-resolution (2023)

Autonomous car test set & corner cases generation (2023)

AI for Automatic Decision Making under Uncertainty

- Finds near-optimal robust decision sequences
- Custom schedulers fitted on Airbus/suppliers supply chain, factories, pulse lines…
- Handles uncertainty (time, resources, fail, ...)

Airbus Pulse Line Scheduling (2023)

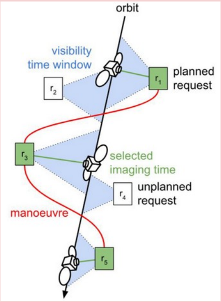

Airbus Earth Observation Scheduling (2023)

AI for Decision Making under Uncertainty – Airbus Applications

Most AI Applications
Data point(s) → Decision = Detection, Localization, Prediction

1-step Decision Making

~ anything a human can achieve in a few seconds gets automated within this category (System 1)

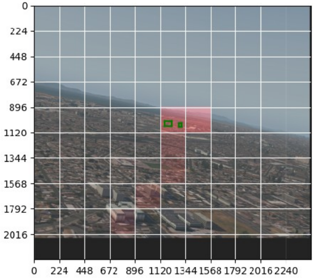

Gaze-like Object Detection
Navigate an image, optimized compute trade-off

AI search solutions
Problem Description → Sequence of decisions over time (e.g. play Go, job-shop, ...)

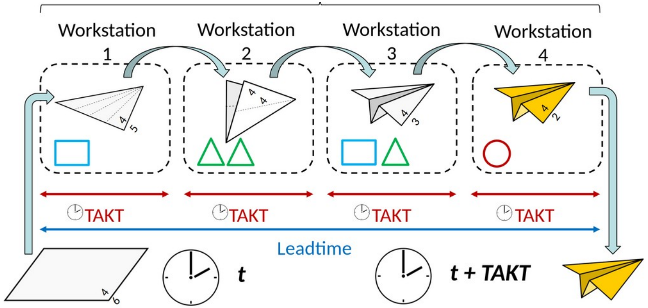

Manufacturing Scheduling & Optimization
- What piece goes on what machine/workstation under resource constraints
- Minimizes makespan / Takt / Costs / Energy / Carbon emission

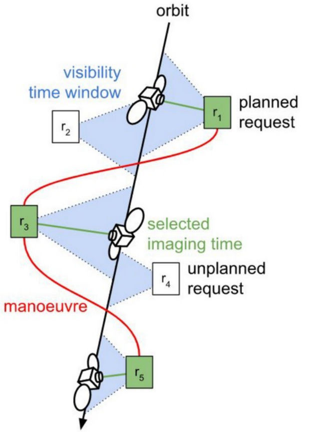

Earth Observation Scheduling
- What requests to execute among too many
- Maximizes priorities and returns

N-steps Decision Making

Humans cannot find good solutions easily
→ cannot be calculated in your head

Team

Dedicated team to yield the best chances of unlocking decision under uncertainty on many tasks.

- Planning & Scheduling senior experts with 20y+ XP
- AI/DL R&D engineers with XP in Neural Combinatorial Optimization (NCO)
- AI/DL developers

Emmanuel Benazera
P&S/DL Researcher with 20y+ XP

Nicolas Meuleau
RL/P&S Researcher with 25y+ XP

Guillaume Infantes
RL/P&S/DL Researcher with 20y+ XP.

Antoine Jacquet
#2 Coding Games worldwide (!)
MuZero expert

Louis Jean
RL/DL for compute-adaptive vision and much more

Pierre Pereira
RL/DL for NCO, generalization, complexity

Along with “Magic Stephanie” from ONERA :)

Application to Compute-Adaptive Object Detection

Detection in Large Images has issues

- Akin to finding a needle in a 4k x 4k+ pixels haystack
- The larger the image, the lower the model size to fit in (GPU) memory
- Model FLOPS increases non-linearly with the size of the image
- Small objects have low accuracy due to image resizing
- Sliding windows consume compute linear to the image size
- **Most compute is wasted in areas of no interest**
  - e.g. planes → mostly sky, boats → water, ring → finger
- **Tracking is separated from detection**
  - Requires two steps, with error accumulation

- Human gaze (basically) works in two steps
  - Overview of the full image field
  - Rapid target specialization
  - Gaze ‘walk’ is **planned** through the image

- Human gaze trajectories are not random
  - **Can we plan an object detector to ‘walk’ through large images ?**

- (As a side-note: Pre-DL computer vision rarely worked on the full image at once)

Episode length (auto-regressive)

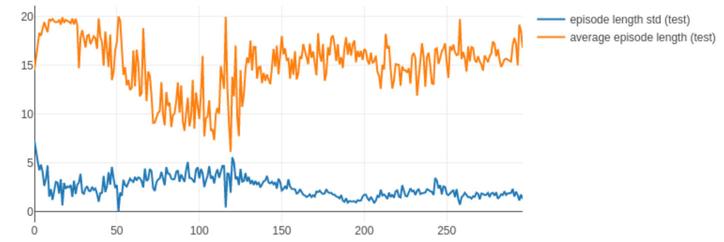

% patches found	97.3%
MAP	N/A
Episode length	16.7
% stop action used	99.1%

BBox patches metrics (test) (auto-regressive)

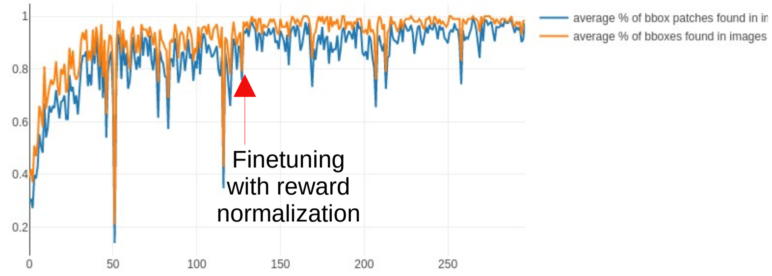

Stop action metrics (auto-regressive)

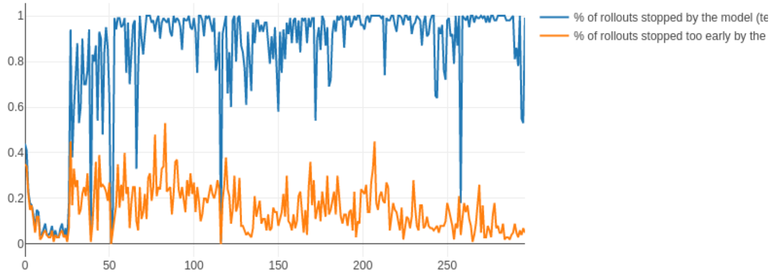

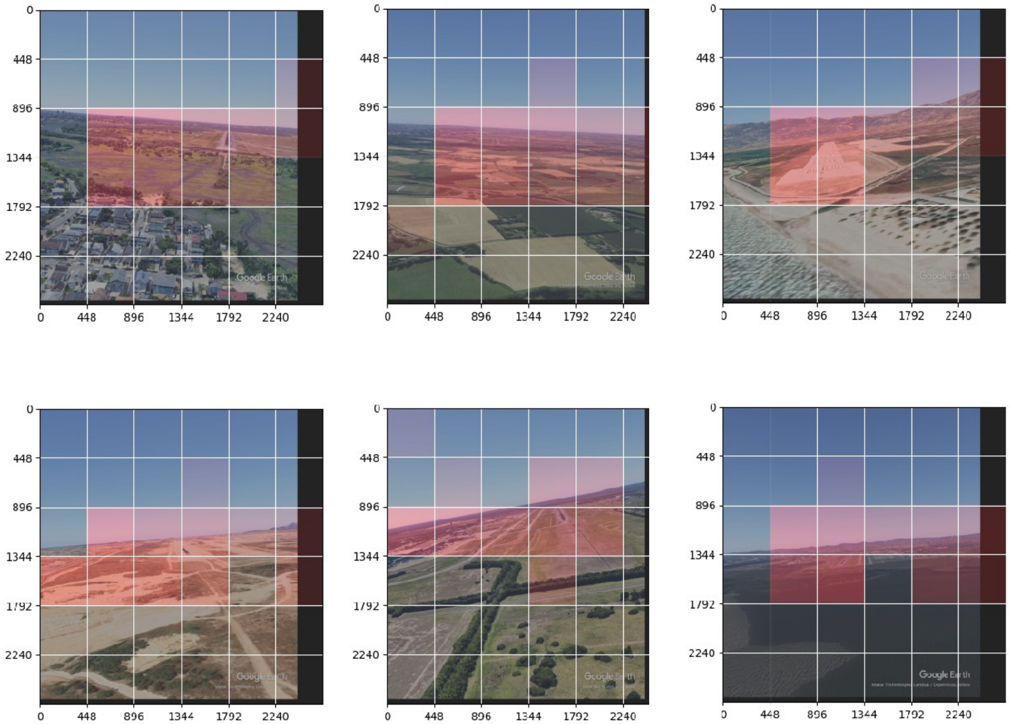

Application to Supply Chain Scheduling

Difficulties for Automating Decision Making under Uncertainty

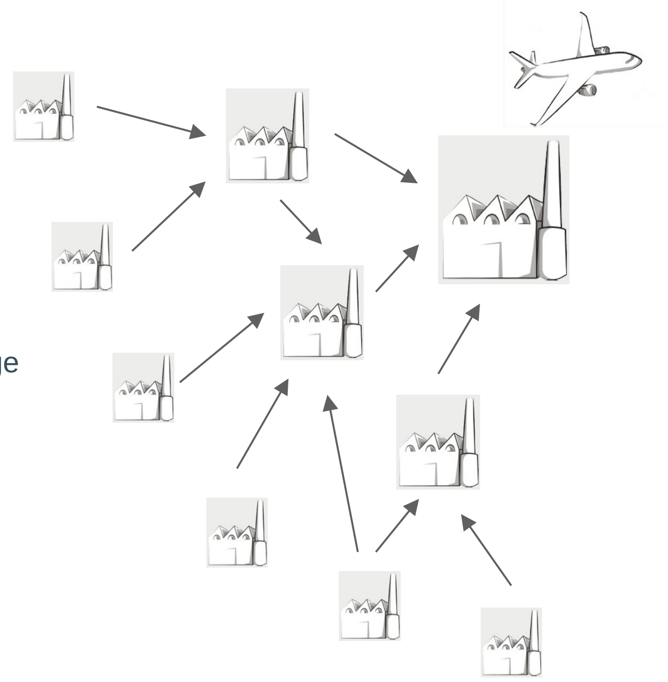

State of the existing software products & solutions

- Scheduling problem has complexities and exploding time points
    - Huge solution search space!
- Software solvers are “universal” and costly → hours to days of compute
    - Compute is exponential the size of the problem → Prevents “solving” large problems
- Solver use custom heuristics that are costly to craft / write by hand (code)
- **Solvers don’t handle uncertainty**
    - Schedule for average uncertainty or worst uncertainty (pessimistic)
    - Consequence: most “work” takes place off-schedule
    - Hard / too long to re-schedule every time uncertainty drives execution off-schedule
    - *Uncertainty builds across the supply chain*

WHEATLEY : an Open Source Product for Next Generation Scheduling

wheatley

Next-generation scheduling problem solver based on GNNs and Reinforcement Learning

reinforcement-learning scheduling planning graph-neural-networks jssp

No time discretization

Use ML to learn a **robust digital twin** from simulated pulse line or satellite, with uncertainties

Searches for best solutions / schedules with RL while learning the twins under uncertainty

Generalization: learns from small problems, apply to large problems

Consequences

- Learns a custom solver once (long compute a single time)
- Fits precisely to the factory / satellite at hand
- Executes / Finds schedules in seconds
- **Schedules are robust to uncertainties**, e.g. no re-scheduling in most cases

Results on Airbus Pulse Line Scheduling: beating optimal average solution

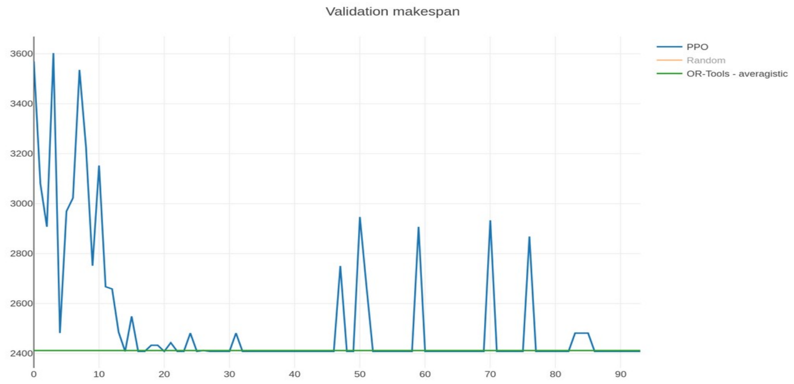

Airbus Pulse Line schedule under uncertainty
→ better makespan than average-uncertainty optimal solution

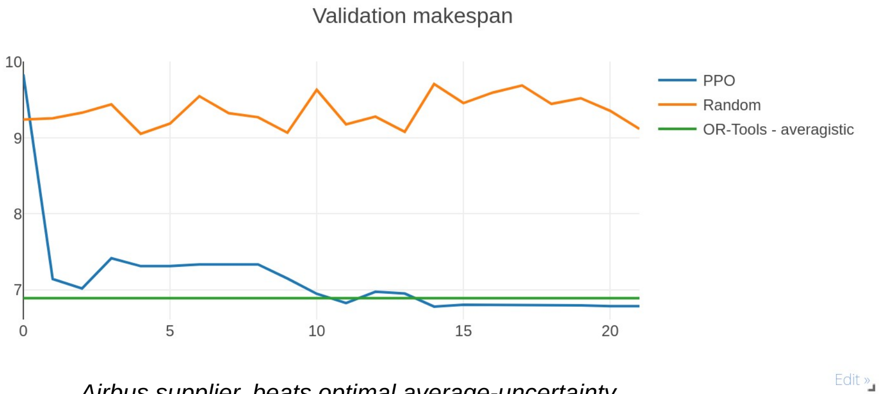

Airbus supplier, beats optimal average-uncertainty schedule.

Wheatley solves for scheduling problem under both uncertainty and resource constraints

- Beats optimal solution based on average uncertainty
- Scales from suppliers to Airbus-size problems

- **Jolibrain initiated and leads a regional project (backed by Occitanie) with Onera and Agilea Group**
    - Airbus has an ‘observer’ status

Results on JSSP

Evaluation	W-10x10	L2D	Best PDR	OR-Tools
6 × 6	521 (7.4)	571 (17.7)	545 (12.4)	**485 (0)**
10 × 10	890 (9.6)	993 (22.3)	948 (16.8)	**812 (0)**
15 × 15	1389 (17.2)	1501 (26.7)	1419 (19.8)	**1185 (0)**
20 × 15	1583 (16.9)	-	1642 (21.3)	**1354 (0)**
20 × 20	1959 (24.9)	2026 (29.2)	1870 (19.3)	**1568 (0)**
30 × 10	1829 (5.5)	-	1878 (8.9)	**1725 (0)**
30 × 15	2043 (14.5)	-	2092 (17.3)	**1784 (0)**
30 × 20	2377 (22.0)	-	2331 (19.7)	**1948 (0)**
50 × 15	3060 (8.3)	-	3079 (9.0)	**2825 (0)**
50 × 20	3322 (14.9)	-	3295 (14.0)	**2891 (0)**
60 × 10	3357 (1.7)	-	3376 (2.3)	**3301 (0)**
100 × 20	5886 (6.9)	-	5786 (5.1)	**5507 (0)**

Table 3: Results on *deterministic* Taillard instances

Evaluation	W-10x10	Wd-10x10	MOPNR	CP-stoc	OR-Tools	
					mode	real
6 × 6	714 (16.3)	817 (33.1)	699 (13.8)	**669** (9.0)	728 (18.6)	*614 (0)*
10 × 10	1217 (21.5)	1464 (46.1)	1252 (25.0)	**1177 (17.5)**	1262 (25.9)	*1002 (0)*
15 × 15	1889 (29.3)	2406 (64.7)	1988 (36.1)	**1872 (28.1)**	1925 (31.8)	*1461 (0)*
20 × 15	**2181 (30.5)**	2729 (63.3)	2314 (38.5)	2222 (33.0)	2244 (34.3)	*1571 (0)*
20 × 20	2643 (36.4)	3511 (81.2)	2708 (40.0)	**2631 (35.8)**	2619 (35.1)	*1938 (0)*
30 × 10	**2425 (14.1)**	3511 (65.2)	2532 (19.1)	2476 (16.5)	2598 (22.2)	*2126 (0)*
30 × 15	**2792 (26.7)**	3251 (47.5)	2964 (34.5)	2892 (31.2)	2943 (33.5)	*2204 (0)*
30 × 20	**3305 (36.9)**	4186 (73.3)	3390 (40.4)	3355 (39.0)	3299 (36.6)	*2415 (0)*
50 × 15	**4043 (16.5)**	4413 (27.1)	4262 (22.8)	4239 (22.1)	4435 (27.7)	*3472 (0)*
50 × 20	**4520 (26.8)**	5351 (50.1)	4679 (31.2)	4682 (31.3)	4758 (33.4)	*3566 (0)*
60 × 10	**4315 (6.3)**	4475 (10.2)	4451 (9.6)	4442 (9.4)	4579 (12.8)	*4061 (0)*
100 × 20	**7591 (11.8)**	8377 (23.3)	7956 (17.1)	8203 (20.8)	8188 (20.5)	*6793 (0)*

Table 4: Results on *stochastic* Taillard instances

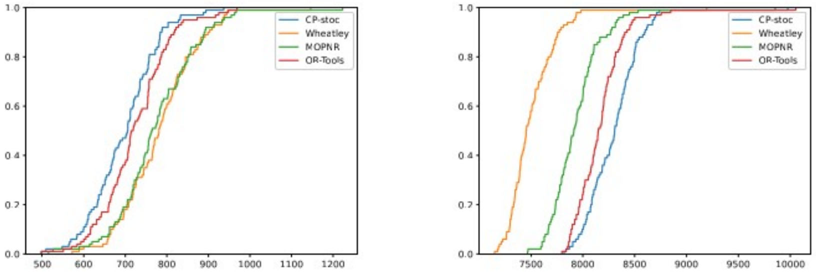

(a) 6 × 6 instance

(b) 100 × 20 instance

Fig. 4: Cumulative makespan of W-10x10 and *CP-stoc* for 100 duration scenarios.

« Learning to solve Job Shop under Uncertainty »
G.Infantes et al. CPAIOR 2024

Application to Earth Obvervation Scheduling

EOSP continuous approach

Discrete graph is the internal state of the Environment.
Continuous graph is the observation provided to the RL Agent.

Continuous graph nodes are acquisition requests.
Continuous graph nodes attributes contain min/max/mean of corresponding discrete nodes attributes.
A continuous graph edge exists if the is at least one discrete edge for corresponding pair of discrete nodes.

Acquisitions	Discrete Nodes	Continuous Nodes	Ratio	Discrete Edges	Continuous Edges	Ratio
100	10297	100	97	835566	9273	90
300	52020	300	169	12598738	81244	155
500	46589	500	92	14842035	225398	66
800	59583	800	74	28015753	447945	63
1100	94071	1100	88	58343397	741634	79

Discrete graph

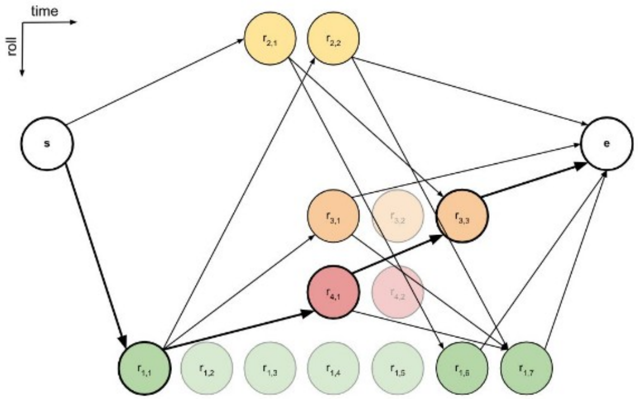

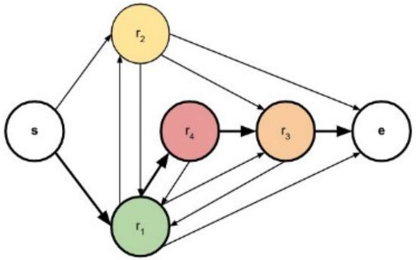

Continuous graph

Results on Earth Observation Scheduling: beating greedy operational solution

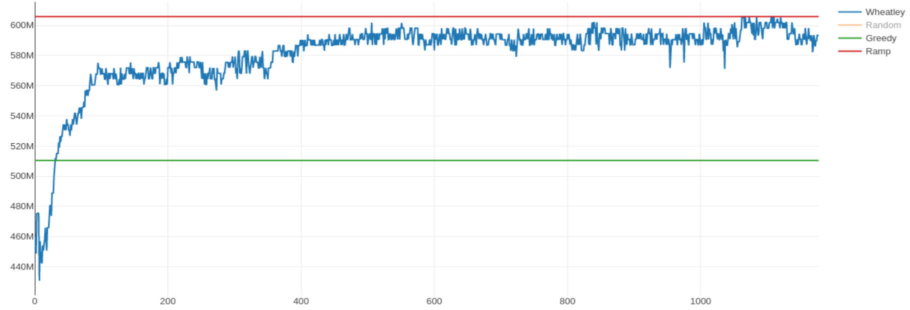

Training on 639 medium problems and evaluating on 27 unseen medium problems

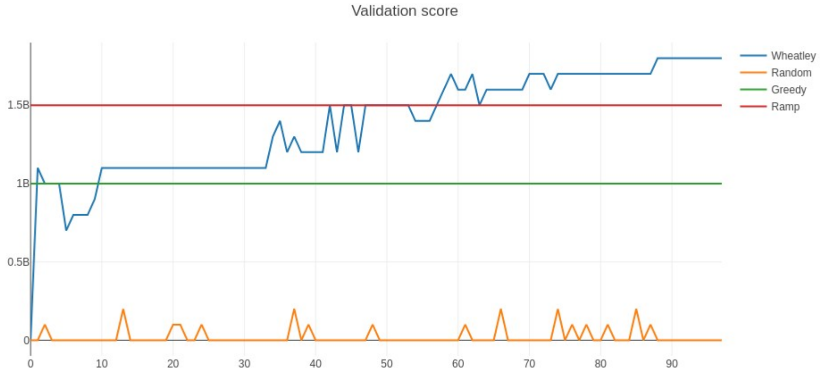

STN + learning on single problem

- Training set : generated hundreds of scheduling problems to learn from;
- Wheatley beats operational greedy heuristic easily on medium sized problems (larger problems under study);
- Generalizes to 5x/10x larger problems;
- Wheatley allows solving the problem from its continuous formulation, contrary to the pitfall of most solutions in this space;
- New STN formulation nears best solutions without learning;
- Learning amplifies the gains.

Conclusion

- Jolibrain is an AI A-team
- Dedicated AI Decision making team with 20y XP + pool of external scientists
- **Decision under uncertainty is Jolibrain's focus as it generalizes other approaches**

- Collaborate with us
  - We're **looking for the most challenging problems** in decision making
  - Many so-called impossible endeavours can be solved through careful reformulation
  - Visit us, best solutions and formulations emerge from live exchanges

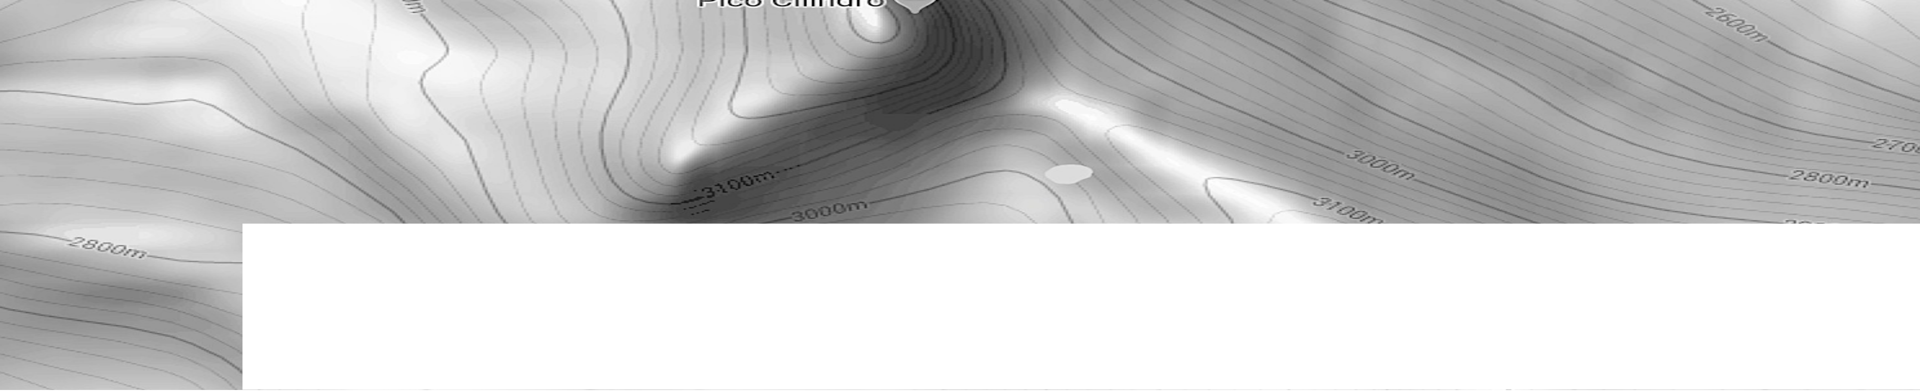

Emmanuel Benazera

emmanuel.benazera@jolibrain.com

+33 6 29 76 47 56

jolibrain

WHEATLEY : an Open Source Product for Next Generation Scheduling

Mix of Jolibrain-copyrighted Open Source Software and custom-developed additions for Airbus applications

Problem simulator

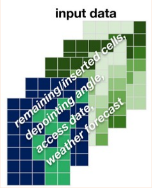

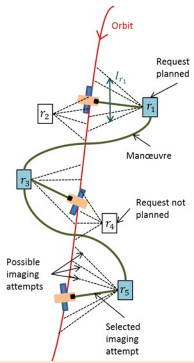

Simulator outputs to tentative schedule

①

Custom problem input reader

②

Custom Application Scheduling API

Airbus Operational Software

Custom Application Scheduling Components

Heuristics baselines
Optimal baselines
for comparison

Wheatley

Neural Scheduling Model

Wheatley configuration for Application

④

③

Origin
Open source components uner Jolibrain Copyright
Already developed by Jolibrain (-> license transfer)
Covered by Jolibrain for Airbus (-> with IP transfer)
Covered by Airbus

Process	
1:	*Problem modeling & acquisition*
2:	*Problem in Wheatley format*
3:	*Scheduling model training*
4:	*Evaluations, test-beds, comparisons*

Operational score / fixed problem

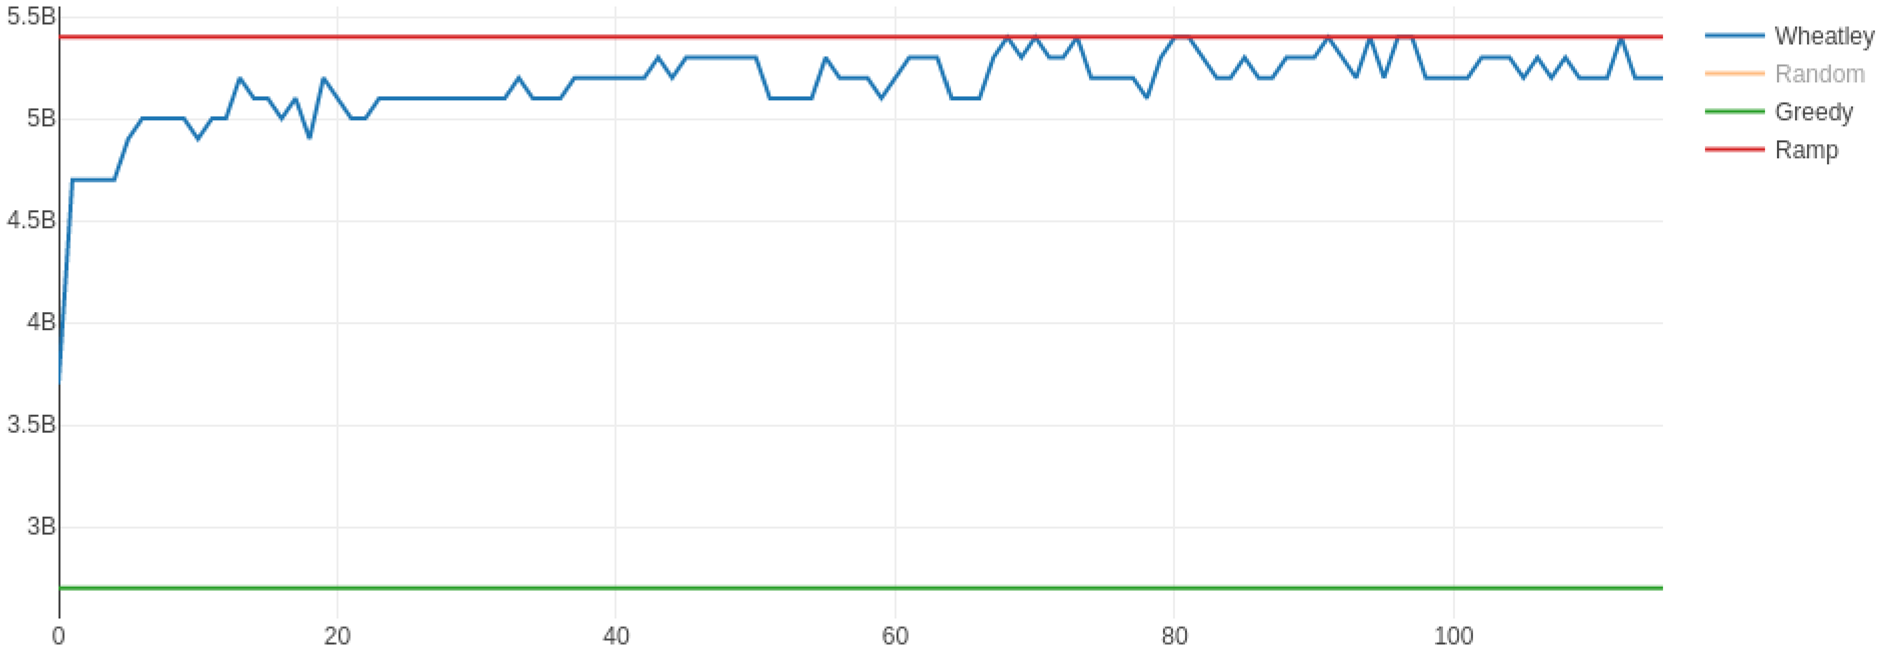

Training and evaluating on the “eosp_88acqr_53ramp_40greedy_DTCyUhAw” problem
Overfits and reaches the Ramp solution

Operational score / generalization / score at each checkpoint

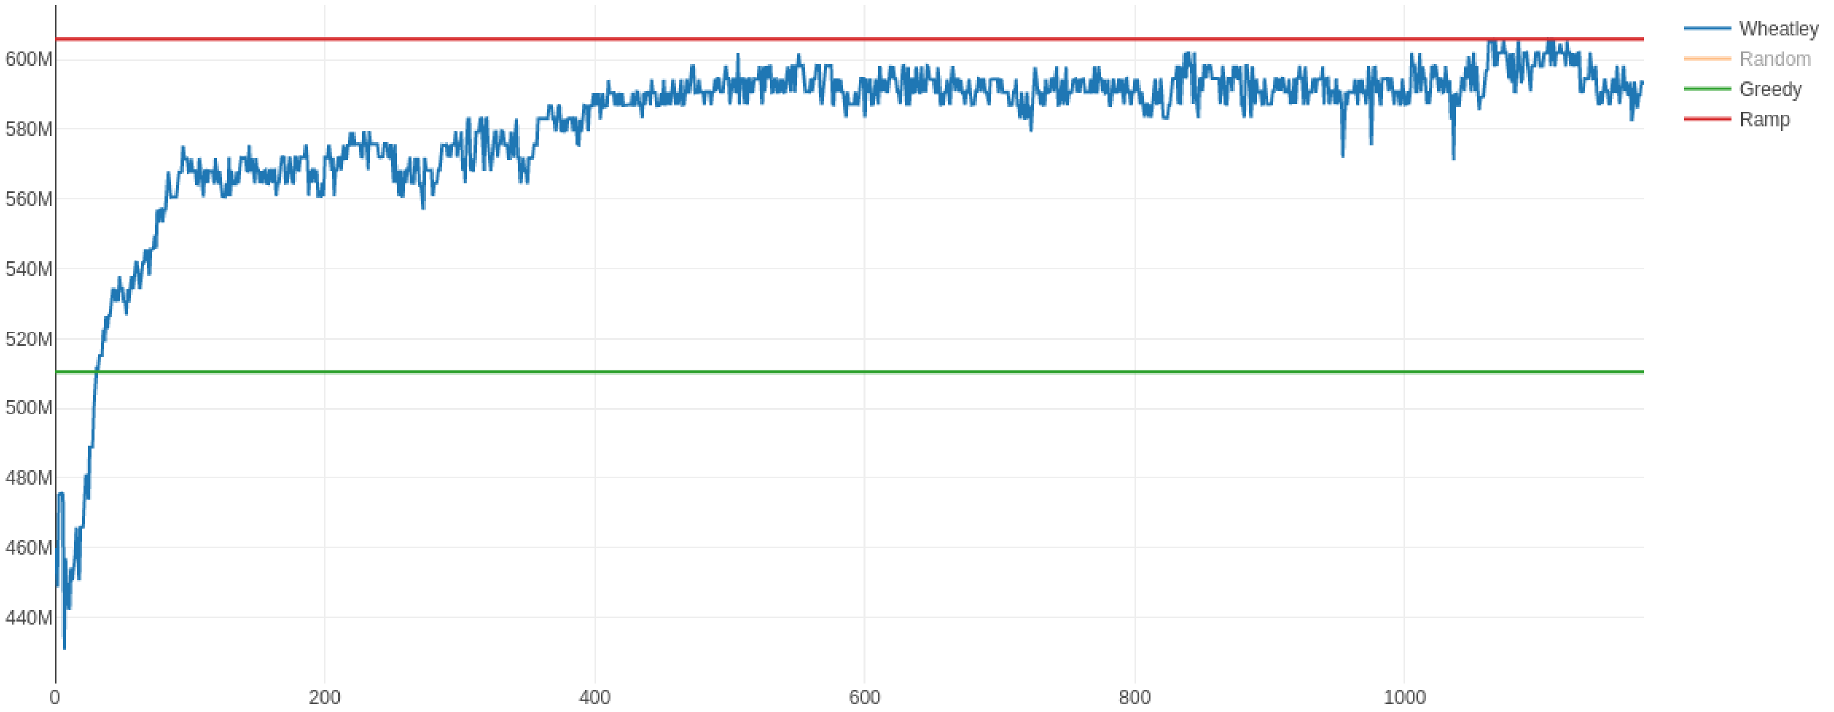

Training on 639 medium problems and evaluating on 27 unseen medium problems

Operational score / generalization / comparison at each checkpoint

Wheatley vs Greedy

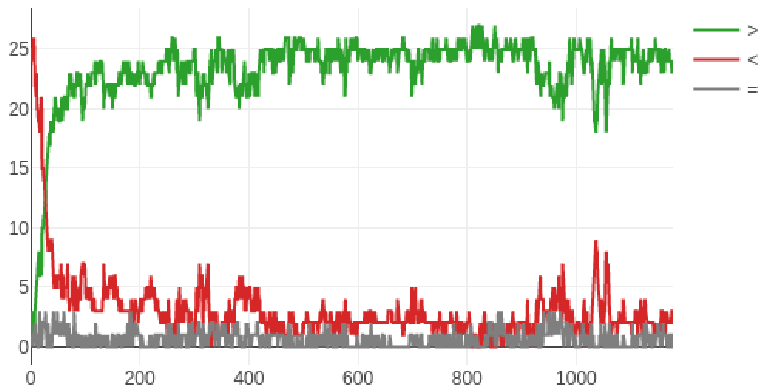

Wheatley vs Ramp

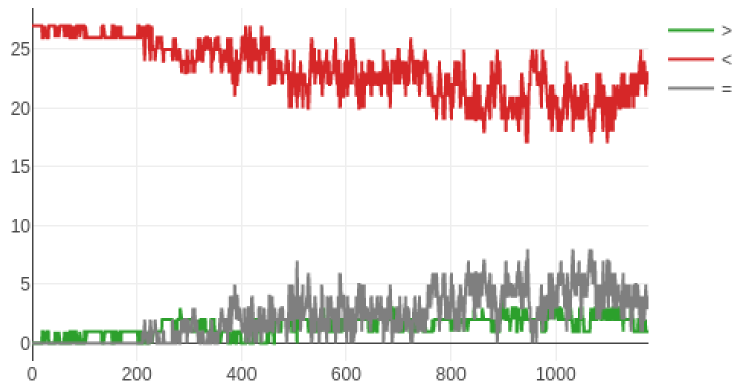

Number of test problems where Wheatley is above, equal or below target algorithm

Operational score / inference exploration / pre-training

We trained a model with Wheatley to solve ~100 acquisitions problems in operational mode.
During inference, Wheatley uses the argmax of the policy to select the best action.
Under the hood, Wheatley uses PPO which has a critic and an actor network head.

We can use these heads to explore during inference using:
- MCTS (like AlphaZero) with 40 sims
- Beam search with a window size of 8 candidates (according to critic) and expanding the 2 best actions (according to actor)

Algorithm	Ramp	Wheatley	Wheatley/Ramp	Time
Wheatley argmax	605894711	605732913	0.9901	~ 7 minutes
Wheatley MCTS(40)	605894711	598283522	0.9888	~ 15 hours
Wheatley Beam(8, 2)	605894711	609511049	0.9983	~ 6 hours

We also tried pre-training the actor head using Ramp solution in supervised learning.
This allows Wheatley to start with better solutions, but did improve our best results.