

Welcome from all the chairs




Frank Hutter
General Chair
**University of Freiburg &
Bosch Center for Artificial
Intelligence**



Mihaela van der Schaar
PC Chair
University of Cambridge



Marius Lindauer
PC Chair
Leibniz University Hannover



Isabelle Guyon
PC Chair
**INRIA, University of Paris-
Saclay**



Joaquin Vanschoren
Tutorial Chair
**Eindhoven University of
Technology**



Holger Hoos
Journal Track Chair
Leiden University



Wei-Wei Tu
Competition Chair
**4Paradigm Inc., China and
ChaLearn, USA**




Alexander Tornede
Review Workflow Chair
University of Paderborn



Raman Arora
Local Chair
Johns Hopkins University



Colin White
Local Chair
Abacus.AI



Katharina Eggensperger
Social Chair
University of Freiburg



Matthias Feurer
Social Chair
University of Freiburg

The foundation of a new conference

- Preceded by a lot of workshops we organized
 - Yearly AutoML workshops at ICML since 2014
 - BayesOpt workshops at NeurIPS
 - MetaLearning workshops at NeurIPS
 - NAS workshops at ICLR
 - Many additional related workshops at many different conferences
- Risk of so many workshops: fragmentation of the community
- AutoML-Conf can become the overarching conference for all of these

Benefits of a more focussed conference

- Higher density of papers that feel relevant
- More-informed reviews
- Exploiting synergies between different parts of AutoML
 - E.g., advances in Bayesian optimization can directly help NAS
 - Meta-learning is relevant to HPO and NAS
 - Gradient-based NAS can inform gradient-based HPO
 - ...
- Building a tighter network
 - You will see many familiar faces at AutoML-Conf every year

Let's build a happy community

- Be open, welcoming, inclusive, polite
 - Try to show interest in others' work
 - Think of your own first conference experience
 - The junior community member you don't know might become your next awesome collaborator
 - Criticism is important, but the tone makes a big difference
 - The same applies to writing reviews

- For any issues,
please reach out to our diversity chair



Theresa Eimer



Katharina
Eggensperger

COVID-19

- We'd like to avoid becoming another superspreader event
 - We have the extra risk of following ICML
 - Thanks to everyone for working together and being supportive
- CVPR had very similar counter-measures as we do
 - Mandatory vaccinations & masks
 - Also gave out free antigen tests daily
 - But vigilance often ends at dinners, in the bars, etc
- Let's be smart about this
 - Use outdoor spaces whenever possible
 - Choose wisely for lunch/dinner/drinks locations
 - E.g., <https://automl.cc/restaurants/> lists restaurants with shaded outdoor seating
 - E.g., our social evening events will be on the patio, on a yacht, and in a rooftop bar

Schedule

All times are in time zone EDT (Baltimore local time)			
	Monday, July 25th	Tuesday, July 26th	Wednesday, July 27th
8:00-8:30	Registration & Coffee & Bagels (Glass pavilion)	Coffee & Bagels (Glass pavilion)	Coffee & Bagels (Glass pavilion)
8:30-9:00			
9:00-9:30	Intro (Hodson 110)	Keynote by Jeff Clune <i>AI-generating algorithms: the fastest path to AGI?</i> (Hodson 110)	Keynote by Julie Josse <i>Missing data: from inference to imputation and prediction; an overview of the main challenges</i> (Hodson 110)
9:30-10:00	Keynote by Alex Smola <i>AutoGluon: Recent Advances on AutoML for Tabular Data</i> (Hodson 110)	Poster spotlights (Hodson 110)	Poster spotlights (Hodson 110)
10:00-10:30			
10:30-11:00	Coffee break (Glass pavilion)		
11:00-11:30	Tutorial I: NAS 1 by Colin White & Debadeepta Dey (Hodson 110)	Tutorial II: Learning Curves by Felix Mohr & Jan van Rijn (Glass pavilion)	Conference posters & coffee break (Glass pavilion)
11:30-12:00			Workshop posters & coffee break (Glass pavilion)
12:00-12:30			
12:30-13:00		Lunch (city, on your own)	Roundtable lunch (Glass pavilion)
13:00-13:30	Roundtable lunch (Glass pavilion)		
13:30-14:00			
14:00-14:30		Keynote by Anima Anandkumar <i>Trinity of Explainable AI: Calibrated, Verifiable, and User-friendly AI</i> (Hodson 110)	Competition reports (Hodson 110)
14:30-15:00	Tutorial III: AutoRL by Aleksandra Faust (Hodson 110)	Tutorial IV: NAS 2 by Tejaswini Pedapati & Martin Wistuba (Glass Pavilion)	Poster spotlights (Hodson 110)
15:00-15:30			
15:30-16:00			
16:00-16:30	Coffee break (Glass pavilion)	Conference posters & coffee break (Glass pavilion)	Workshop posters & coffee break (Glass pavilion)
16:30-17:00	Keynote by Chelsea Finn : <i>Meta-Learning for Education</i> (Hodson 110)		
17:00-17:30		Panel Discussion : <i>AutoML in the age of large pretrained models</i> (Hodson 110)	Keynote + panel discussion, lead by Timnit Gebru : <i>On the relation of AutoML and fairness research</i> (Hodson 110)
17:30-18:00	Best paper talk (Hodson 110)		
18:00-18:30	Break	Transfer to banquet (on your own; we meet at 555 Light St, Baltimore, MD 21202 at 6:45pm, on time)	Closing (Hodson 110)
18:30-19:00	Research speed dating (Glass pavilion)		
19:00-19:30			Dinner (city, on your own)
19:30-20:00			
20:00-20:30		Banquet (on a yacht in the Baltimore Inner Harbour)	
20:30-21:00	Reception (Glass pavilion)		
21:00-21:30			Informal social event (in a bar outside)
21:30-22:00			

Schedule elements

- Lots of time for networking
 - Roundtable lunch
 - Research speed-dating
 - Reception
 - Banquet
 - Social event in a bar
- Extended, interleaved poster sessions (always with coffee)
 - Tuesday: conference & journal track posters – posters are up all day
 - Wednesday: workshop & competition posters – posters are up all day
- Monday: tutorials
- Tuesday & Wednesday: panels
 - AutoML in the age of pretrained networks
 - AutoML and fairness
- 2 keynotes per day

Keynote Speakers




Anima Anandkumar
Caltech & NVIDIA

Trinity of Explainable AI: Calibrated, Verifiable, and User-friendly AI



Jeff Clune
**University of British
Columbia & OpenAI**

AI-generating algorithms: the fastest path to AGI?



Chelsea Finn
Stanford University

Meta-Learning for Education




Timnit Gebru
**Distributed Artificial
Intelligence Research
Institute**

On the Relationship between Fairness and AutoML (panel discussion)



Julie Josse
INRIA Ecole Polytechnique

Missing data: from inference to imputation and prediction; an overview of the main challenges



Alex Smola
Amazon Web Services

AutoGluon: Recent Advances on AutoML for Tabular Data

Virtual Attendees

- All parts in yellow will be live streamed
- For questions on Zoom, either
 - use the chat (the session chair will read out questions)
 - raise your hand & we'll signal you to unmute
- Next week on Monday, August 1st
 - Additional virtual poster session in Gather.Town
 - Tutorial 5 on “Statistical Analyses for Multi-objective Stochastic Optimization Algorithms” with Q&A



All times are in time zone EDT (Baltimore local time)

	Monday, July 25th	Tuesday, July 26th	Wednesday, July 27th
8:00-8:30	Registration & Coffee & Bagels (Glass pavilion)	Coffee & Bagels (Glass pavilion)	Coffee & Bagels (Glass pavilion)
8:30-9:00	Intro (Hodson 110)	Keynote by Jeff Clune <i>AI-generating algorithms: the fastest path to AGI?</i> (Hodson 110)	Keynote by Julie Josse <i>Missing data: from inference to imputation and prediction; an overview of the main challenges</i> (Hodson 110)
9:00-9:30	Keynote by Alex Smola <i>AutoGluon: Recent Advances on AutoML for Tabular Data</i> (Hodson 110)	Poster spotlights (Hodson 110)	Poster spotlights (Hodson 110)
9:30-10:00	Coffee break (Glass pavilion)		
10:00-10:30	Tutorial I: NAS 1 by Colin White & Debadepta Dey (Hodson 110)	Tutorial II: Learning Curves by Felix Mohr & Jan van Rijn (Glass pavilion)	Conference posters & coffee break (Glass pavilion)
10:30-11:00			Workshop posters & coffee break (Glass pavilion)
11:00-11:30			
11:30-12:00	Roundtable lunch (Glass pavilion)	Lunch (city, on your own)	Roundtable lunch (Glass pavilion)
12:00-12:30			
12:30-13:00			
13:00-13:30			
13:30-14:00			
14:00-14:30	Tutorial III: AutoRL by Aleksandra Faust (Hodson 110)	Tutorial IV: NAS 2 by Tejaswini Pedapati & Martin Wistuba (Glass Pavilion)	Keynote by Anima Anandkumar <i>Trinity of Explainable AI: Calibrated, Verifiable, and User-friendly AI</i> (Hodson 110)
14:30-15:00			Competition reports (Hodson 110)
15:00-15:30			Poster spotlights (Hodson 110)
15:30-16:00	Coffee break (Glass pavilion)	Conference posters & coffee break (Glass pavilion)	Workshop posters & coffee break (Glass pavilion)
16:00-16:30	Keynote by Chelsea Finn <i>Meta-Learning for Education</i> (Hodson 110)	Panel Discussion: AutoML in the age of large pretrained models (Hodson 110)	Keynote + panel discussion, lead by Timnit Gebru: <i>On the relation of AutoML and fairness research</i> (Hodson 110)
16:30-17:00	Best paper talk (Hodson 110)	Transfer to banquet (on your own; we meet at 555 Light St, Baltimore, MD 21202 at 6:45pm, on time)	Closing (Hodson 110)
17:00-17:30	Break		
17:30-18:00	Research speed dating (Glass pavilion)		
18:00-18:30			
18:30-19:00			Dinner (city, on your own)
19:00-19:30			
19:30-20:00			
20:00-20:30		Banquet (on a yacht in the Baltimore Inner Harbour)	
20:30-21:00	Reception (Glass pavilion)		
21:00-21:30			Informal social event (in a bar outside)
21:30-22:00			

Welcome in the Name of the PC-chairs



Mihaela van der Schaar
PC Chair
University of Cambridge



Marius Lindauer
PC Chair
Leibniz University Hannover



Isabelle Guyon
PC Chair
INRIA, University of Paris-Saclay


Goals


1. High-quality papers
 - The conference scope is more niche, but the quality should be at least on the level of ICML, NeurIPS or ICLR
2. High-quality reviews
 - Reviewers are likely much more familiar with AutoML than the average reviewer for ICML, NeurIPS or ICLR
 - Each paper received 4.5 reviews on average (at least 3)
3. Embrace Open-Source
 - Modern science is strong because of the chance of collaborations and building upon others' work; open-source is mandatory
4. Ensure Reproducibility
 - Reproducibility crisis is again and again a major problem; not for AutoML-Conf!
 - Explicit reproducibility reviewers

AutoML-Conf in Numbers

- Main Track
 - 104 submissions
 - 22 desk-rejected
 - 17 papers withdrawn by the authors
 - 20 papers accepted for main track
 - 19,2 % acceptance rate wrt all submissions
 - (32,8% acceptance rate wrt submissions after desk-rejects and withdrawn)
- Workshop Track
 - 28 submissions
 - 20 papers accepted
 - 71 % acceptance rate
 - an additional 16 papers were invited from the main track to workshop track
- Journal Track
 - All 7 (previously reviewed) submissions were accepted
 - No additional reviewing!

Acceptance Distribution





Open Research Knowledge Graph (ORKG)

Venue: <https://orkg.org/venue/R197419>

Comparison: <https://orkg.org/comparison/R199173/>



View

Tools

About

Metrics

Resource

Search

+ Add new



Results and Footprints of AutoML-Conf22 Papers

Comparing empirical results and compute footprints of papers accepted at AutoML-Conf22; papers only included that provided meta-information

July 2022

Martin Lindauer

Use Shift + Mouse Wheel for horizontal scrolling in the table.

Properties	Automated Super Network Generation for Scalable Neural Architecture Search	BERT Sort-A-Zero-shot MLM Semantic Encoder on Ordinal Features for AutoML	AutoCats: A Unified Data Model Co Search Framework for Graph Neural Networks	Tackling Neural Architecture Search With Quality Diversity Optimization	YAHPO-Gem: An Efficient Multi-Objective Multi-Fidelity Benchmark for Hyperparameter Optimization	Non-Uniform Adversarially Robust Pruning	When, where, and how to add new resources to ANNs	ScaleNAS: Multi-Path One-Shot NAS for Scale-Aware High-Resolution Representation	SynE Tune: A Library for Large-Scale Hyperparameter Tuning and Reproducible Research	DHSA: Differentiable Automated Feature Engineering	What to expect of hardware metrics predictions in NAS	Differentiable Architecture Search for Reinforcement Learning
publication/evaluation	nan	nan	nan	nan	nan	nan	nan	nan	nan	nan	nan	nan
benchmark performance	nan	Ordinal Value Benchmark, 27% semantic ordinal accuracy Ordinal Value Benchmark, 55% accuracy	Actor 38.04 Texas 80.21 Wisconsin 80.39 Cornell 64.86 Computer 78.91 CS 93.05 Photos 85.16 Physics 93.28	nan	nan	nan	nan	Multi-Person Pose Estimation COCO test-dev, 71.6 AP Multi-Person Pose Estimation CrowdPose, 71.3 mAP @0.5:0.95	nan	nan	nan	Multi-Task/Game Program w/ Rainbow and Micro Search Space, Up to 250% performance over IMPALA-CNN baseline on selected games Single-Task/Game Program w/ PPO and Macro Search Space, RL-DARTS is competitive with strong IMPALA-CNN baseline Search for random search significantly Single-Task DM-Control w/ SAC and Micro Search Space, Up to 20% improvement over 4-layer baseline CNN encoder RL-DARTS in Random Search w/ Rainbow and Micro Search Space, Up to 10x search efficiency over random search
benchmark time	nan	Ordinal Value Benchmark, 1 AutoML Ordinal Value Benchmark, 76	nan	nan	nan	nan	nan	Multi-Person Pose Estimation COCO test-dev, 3 GPU days Multi-Person Pose Estimation CrowdPose, 3 GPU days	nan	nan	HW-NAS-Bench, 44.1 CPU days TransNAS-Bench-101, 4.8 CPU days	IMPALA-CNN Baselines, 72 GPU days Multi-Task/Game Program w/ Rainbow and Micro Search Space, 10 GPU days Single-Task/Game Program w/ PPO and Macro Search Space, 120 GPU days Single-Task DM-Control w/ SAC and Micro Search Space, 6 GPU days RL-DARTS vs Random Search w/ Rainbow and Micro Search Space, 50 GPU days
evaluation metrics	Yes	Yes	No	nan	Yes	No	No	Yes	Yes	Yes	No	Yes
resource consumption	Resource Consumption	Resource Consumption	Resource Consumption	Resource Consumption	Resource Consumption	Resource Consumption	Resource Consumption	Resource Consumption	Resource Consumption	Resource Consumption	Resource Consumption	Resource Consumption
cpu hours	1800.0	87.0	0.0	939.0	3349.0	0.0	11766.793	900.0	300.0	0.0	1200.0	6000.0
estimated co2 footprint	167.0	2.46	5.18	72.3	381.0	204.56	159.6	126.0	nan	nan	nan	24.0
gpu hours	2000.0	0.0	48.0	77.0	1080.0	0.0	4793.967	900.0	12.0	0.0	0.0	270.0
tpu hours	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

CPU hours / GPU hours / CO₂ Footprint

CPU Hours	GPU Hours	CO2 Footprint
0.3	0	N/A
1,2	0	N/A
67	0	2.46
0	48	5.18
4	270	24
939	72	72.3
1,8	2	167
500	12	N/A
900	900	126
11,766.703	4,792.967	159.6
0	960	204.96
3,349	1,08	381

Insights & Open Questions:

- AutoML research can (still) be crazy expensive
 - surrogate benchmark libraries hopefully will help here
 - How does this compare to other conferences?
- No correlation between CPU/GPU hours and CO₂ footprint
- Only few authors report this so far
- **We need better tooling for this!**

Best Paper Award

AutoML-Conf 2022

BEST PAPER AWARD

Automatic Termination for Hyperparameter Optimization

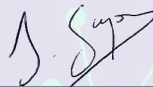
Anastasia Makarova, Huibin Shen, Valerio Perrone, Aaron Klein,
Jean Baptiste Faddoul, Andreas Krause, Matthias Seeger & Cedric Archambeau



Frank Hutter
General Chair



Marius Lindauer
PC-Chair



Isabelle Guyon
PC-Chair



Mihaela van der Schaar
PC-Chair

Tutorials

- Goal: ***easy-to-follow overviews*** on ***recent trends, fast-moving topics***, hands-on sessions or future challenges related to AutoML
- Submissions were very competitive
 - Sadly we could only squeeze in 4 tutorials (+1 virtual)
- Semi-plenary
 - 2 sessions of 2 tutorials each
 - Hudson 110 and Glass Pavilion in parallel
 - Live-streamed in 2 Zoom breakout rooms (recorded)
- 5th tutorial will be virtual-only
 - Will run in Zoom before the virtual poster session (1 Aug)

Off-Campus Events



See also
website!

Wifi: JHGuestNet (no password) or Eduroam

Tuesday Lunch on your own: see website for suggestions

Tuesday Dinner on a Yacht, **We meet at 6:45pm, on time**

Wednesday Dinner on your own: see website for suggestions

Wednesday Bar: The Chasseur, 8:30pm

Questions? Contact Raman Arora or Colin White (colin@abacus.ai, +1-847-828-3885)

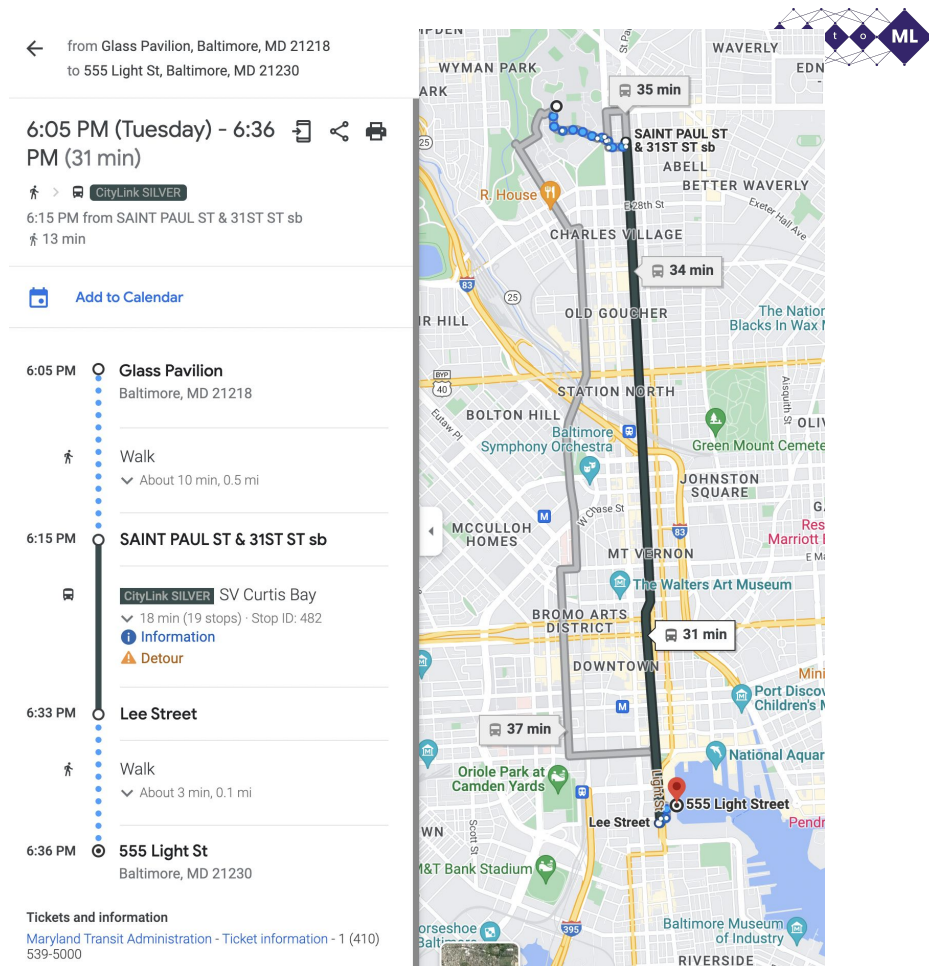
Tuesday Yacht

555 Light St


Baltimore, MD 21202

Arrive by 6:45pm, or it will leave without you!

- Bus (\$2 exact)
- Uber / Lyft



Social Events I: Roundtable Lunch



Where? Glass Pavilion

What? Enjoy your lunch while meeting people working in your domain


When? Monday, 12:30-14:30 and Wednesday 12:00-14:00

Topics

- Benchmarking
- Evolutionary Methods
- Reinforcement Learning
- Zero cost proxies for NAS
- Hyperparameter Optimization
- Open source
- Bayesian optimization
- Large (Language) Models
- Societal Impact of AutoML
- Meta-Learning

See also
website!

Social Events II: Research Speed Dating



When? Today, 6:30pm - 7:30pm

Where? Glass pavilion / patio

What?

- Get to know the AutoML community & extend your network
- Get to know other attendees of the conference that you might not otherwise meet

Join the AutoML community on Slack



Thanks to our Sponsors!

Gold sponsors



Silver sponsors



Bronze sponsors

