Questions for the panel

- Are the AutoML challenges of relevance for improving AutoML systems in practice or is industrial interest a better driver for innovation?+

- Can we deal with the causation problem (e.g. causal inference, causal discovery) by using the AutoML methodology? How about EDA? +

- Given a few rounds of BO, what is the best practice for suggesting the best solution? In other words, how different approaches such as the best observed value and the posterior mean value compares to each other.
- For a given few-shot classification task, what is the best way to determine if we should be using transfer learning/fine-tuning or architecture search?
- Cost-sensitive Bayesian optimization penalizes the acquisition function with the mean of a GP fitted to the log of the (evaluation and overhead) cost function. However, this approach does not consider the potentially different range of the acquisition function and the cost function and also requires fitting an additional GP at each iteration of BO, which can be computationally expensive. What is your take on this current way of modeling the cost in the BO literature and in how this can be improved?
- What are the advantages of using surrogate models besides Gaussian Processes? What are some interesting directions to explore when it comes to development of new surrogate models?
- Why are not there autoML tools for time series as the popular autoML tools auto-sklearn, TPOT,....?
- Could we better (semi)automate the process of defining the intentions expected from ml-solutions?
- Do you think that AutoML will result in more models being created that have concerning biases like Amazon's Rekognition?

Discussed Questions

- To what extent will AutoML replace machine learning practitioners? How do existing and rising ML practitioners (who do not perform cutting edge research and publish papers) best improve their competency in the era of AutoML?+++++++

- Will AutoML lead to job loss / to machine learning not creating as many jobs as it would otherwise? What can we do to avoid that? +

- There were only few submissions on classical hyperparameter optimization; can we consider hyperparameter optimization a solved problem? +++++

- Given how expensive benchmarking automl is and in light of recent papers raising concerns on reproducibility of findings in the domain, do you think building benchmarking simulators like NAS-Bench-101 should be the main focus of the Automl community at the moment? ++++

- Is continuous relaxation for discrete variables the best way to go forward?++++

- We have ML-pipeline optimization and NAS. Is there any other part that we missed so far that requires automatization? ++

- How do we expand the search space more generally? What algorithm could discover convolution or attention layers?+++
- Search space design (designing possible variation of hyperparameters and neural architectures) is still not automated. Do we need to automate it? Can we automate it?++

- Can you comment on the energy consumption of AutoML? Is AutoML being the ultimate global warming tool?+++

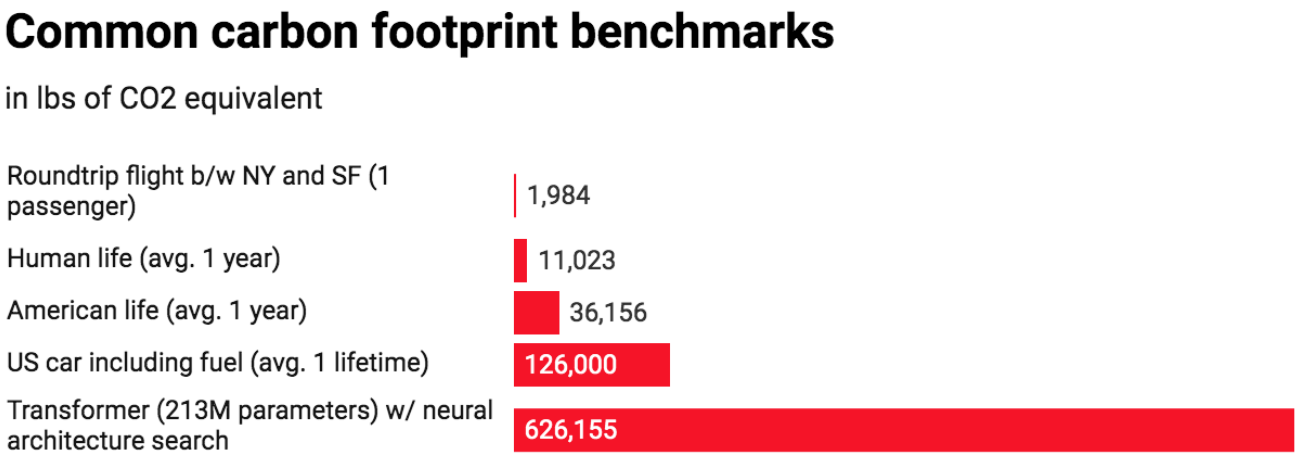

Consumption	CO_2e (lbs)
Air travel, 1 passenger, NY↔SF	1984
Human life, avg, 1 year	11,023
American life, avg, 1 year	36,156
Car, avg incl. fuel, 1 lifetime	126,000
Training one model	
SOTA NLP model (tagging)	13
w/ tuning & experimentation	33,486
Transformer (large)	121
w/ neural architecture search	394,863

(https://twitter.com/strubell/status/1129408199478661120)

- What are the best and worst settings to apply meta-learning to AutoML? +++
- There recently has been a lot of attention towards model robustness and stability. To what extent does autoML currently cover this area? More importantly, how do you think autoML and model robustness/stability will interact in the near future?+