Impressions from NeurIPS 2019 in Vancouver

Yann Coadou[a], Emille Ishida[b]

[a]CPPM Marseille, [b]Université Clermont Auvergne

IN2P3/IRFU Machine Learning workshop
CC-IN2P3, 22 January 2020

Depuis 80 ans, nos connaissances bâtissent de nouveaux mondes

CENTRE DE PHYSIQUE DES PARTICULES DE MARSEILLE
CPPM

Conference size

- Last year (2018): sold out in 12 minutes
- This year: lottery to win the right to pay \$350 fee + \$250 for workshops + \$150 for tutorials
- 4 parallel tracks
- 51 workshops, 9 tutorials
- 79 official NeurIPs meetups in over 35 countries on 6 continents
- 9,185 papers submitted (6,743 after "filtering"), 1,428 accepted (36 orals, 164 5min-spotlights, so mostly posters)
- More than 20,000 reviews written by 4,543 reviewers

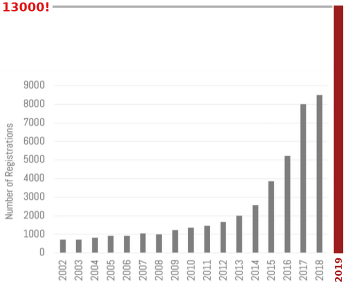

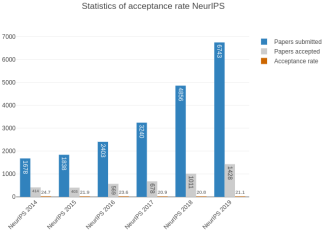

NeurIPSConf@medium.com D. Charrez@medium.com

Synced@medium.com

Conference sponsors. . .

DIAMOND SPONSORS			$80,000
Lambda	Microsoft	Graphcore	Sony
Amazon	Anthem.ai	IBM Research	Deep Mind
Apple	Facebook	Qualcomm	Waymo
Intel AI	Google	Citadel & Citadel Securities	Salesforce
Lyft	OpenAI	Doc.ai	

PLATINUM SPONSORS			$40,000
Invenia Labs	Habana Labs	Voleon Group	QuantumBlack
Naver Line	Uber	Borealis AI	PDT Partners
XPRIZE	D. E. Shaw Group	Twitter	Bloomberg
SigOpt	Cerebras Systems	Yokogawa Electric Corp.	Bosch
J.P. Morgan	Netflix	Inspur	Alibaba Group
Baidu	EY	DiDi	Two Sigma
Jump Trading	Jane Street	Unity Technologies	Disney Research
BioMind	Sportlogiq	Hudson River Trading	ByteDance
Cruise	National Security Agency	Yandex	Google X (The moonshot Factory)

GOLD SPONSORS			$20,000
MathWorks	Weights & Biases	XTX Markets	Squarepoint
G-Research	VinAI Research	Horizon Robotics	Neural Magic
BenevolentAI	Preferred Networks	Prowler.io	Petuum Inc.
Scale AI	Pryon	Splunk Inc.	Sberbank
Mipsology	Snap Inc	Zillow	Novartis
	Point 72/Cubist Systematic Strategies	Cisco Systems	

SILVER SPONSORS			$10,000
Avira	Khosla Ventures	Edgestream Partners	GHELIA Inc.
Tencent AI Lab	Arm Inc	Element AI	Accenture
Expedia Groups	Johnson & Johnson	Blackberry Cylance	Capital Group
Vectra AI	TerraQuanta	Moqi Technology	Alegion
Walmart Labs	Booz Allen Hamilton	Happy Elements	Centurion Capital
	Siemens Medical Solutions		

BRONZE SPONSORS				$5,000
NextAI (NEXT CANADA)	Simon Fraser U.	Grameen Research	Lab41	Wadhwani AI

Conference venue

- *“. . . And all I got for this price was a mug!”*
- Beautiful setting, huge rooms
- Impressive coffee breaks
- 9’ in 15”: heading to keynote ▸ © A. Kurenkov

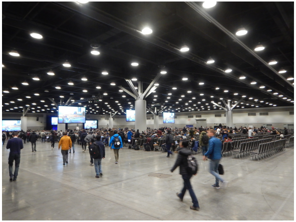

Vancouver

Vancouver

- More typical, 2 hours later. . .

Vancouver


LANDFILL
RECYCLE
ORGANICS
STOP
Is that really garbage?
STYROFOAM
BROKEN GLASS
POTATO CHIP BAGS
CANDY WRAPPERS
FOOD SOILED PLASTICS
UNSOILED PAPER &
CARDBOARD
CANS & BOTTLES
PLASTIC BAGS
COFFEE CUP LIDS
FOOD SCRAPS
WOODEN STIR STICKS &
CHOPSTICKS
COFFEE CUPS &
SLEEVES
COMPOSTABLE WARES
& UTENSILS
FOOD SOILED PAPER

Vancouver

Diversity and inclusion / Conf name

Diversity and inclusion / Conf name

Diversity and inclusion / Conf name

- Women in Machine Learning
- Black in AI (BAI)
- LatinX in AI
- Queer In AI
- {Dis}Ability in AI
- Jews in Machine Learning

Diversity and inclusion / Conf name

- Women in Machine Learning
- Black in AI (BAI)
- LatinX in AI
- Queer In AI
- {Dis}Ability in AI
- Jews in Machine Learning

- Daycare (kids from 75 delegates)
- Whova app for agenda, meetups, social networking, comments, . . .

Conference schedule

- Obviously for ML experts, more on the theoretical side than, e.g., ICML
- Extremely dense ⇒ very tiring exercise for me, but very interesting
- From Sunday morning to Saturday evening
  - Sunday: Expo 2019. Workshops, demonstrations, talks & panels by sponsors and private companies Brochure
  - Monday: 9 tutorials (8:30–18:35)
  - Tuesday–Thursday: Invited talks at 8:30 and 14:15, parallel tracks with one talk + 5(am)/10(pm) spotlights, poster sessions 10:45–12:45/17:30–19:30
  - Friday–Saturday: workshops (8:00–18:40)
- Socials: 15 social events (19:00-22:00, Tue–Thu). Talks, food, interactions, etc. E.g. *ML 4 Space*, *AI for Good*, *British Parliamentary style debate*
- Many videos (live+slides, very browsable) available at https://slideslive.com/neurips/

Authors

- 15,920 authors of submitted papers
- Often multiple authors per paper
- Often mix of academia and industry
- Sergey Levine (UC Berkeley) most contributing author with 12 accepted papers
  - followed by Francis Bach (INRIA–Ecole Normale Supérieure), 10 papers

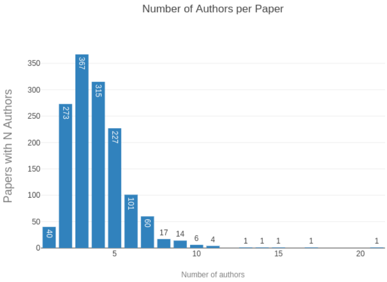

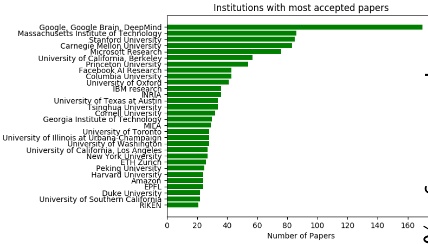

12% of accepted papers

D. Charrez@medium.com

Invited talks

- Breiman lecture: Bin Yu (UC Berkeley), *Veridical Data Science*. Predictability, computability, and stability (PCS)
- Dana Pe'er (Sloan Kettering Institute), *Machine Learning Meets Single-Cell Biology: Insights and Challenges*. Biology becoming a data science
- Blaise Aguera y Arcas (Google), *Social Intelligence*. What is the loss function? Data privacy, energy consumption, ...
- Posner lecture: Yoshua Bengio (MILA, U. of Montreal), *From System 1 Deep Learning to System 2 Deep Learning*. Unconscious/current DL → conscious/future DL, from IID data to OOD generalisation and transfer
- Kafui Dzirasa (Duke U.), *Mapping Emotions: Discovering Structure in Mesoscale Electrical Brain Recordings*. Modelling major depressive disorder with ANN, to understand in vivo signals. Causality, vulnerability
- Jeff Heer (UW), *Agency + Automation: Designing Artificial Intelligence into Interactive Systems*. Complementarity between human and machine, e.g., in data visualisation/cleaning, translation

Awards

- Outstanding Paper Award:
  - I. Diakonikolas, T. Gouleakis, Ch. Tzamos, *Distribution-Independent PAC Learning of Halfspaces with Massart Noise*
  - Honorable Mention:
    - A. Uppal, S. Singh, B. Poczos, *Nonparametric Density Estimation & Convergence Rates for GANS under Besov IPM Losses*
    - A. Maalouf, I. Jubran, D. Feldman, *Fast and Accurate Least-Mean-Squares Solvers*
- Outstanding New Directions Paper Award:
  - V. Nagarajan and Zico Kolter, *Uniform Convergence May Be Unable to Explain Generalization In Deep Learning*
  - Honorable Mention:
    - S. Löwe, Peter O'Connor, Bastiaan Veeling, *Putting An End to End-to-End: Gradient-Isolated Learning of Representations*
    - V. Sitzmann, Michael Zollhoefer, Gordon Wetzstein, *Scene Representation Networks: Continuous 3D-Structure-Aware Neural Scene Representations*
- Test of Time Award:
  - Lin Xiao, *Dual Averaging Method for Regularized Stochastic Learning and Online Optimization*

Datasets

- Too often MNIST or CIFAR, or other “simple” datasets
  - MNIST: handwritten digits, 60k training set and 10k test set
  - CIFAR-10 = 60k 32x32 color images in 10 different classes (6k/class)
  - CIFAR-100 = 100 classes (600/class)
- Bias from “theoretical” conference?
- Risk: many studies not scaling to real life applications

Take away: generalisation in deep learning

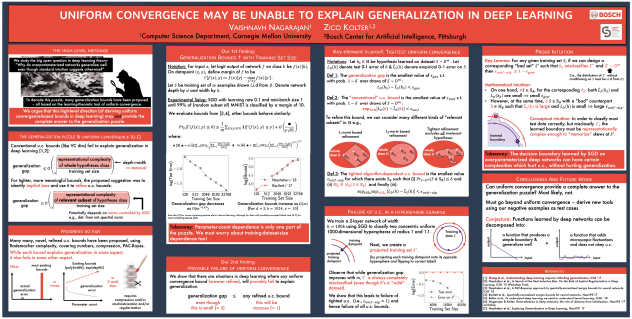

- Actually works surprisingly well
- Over-parameterised DNN should overfit but don't: why?
- Neural tangent kernel (NTK): helps thinking in infinite-width limit. But can do better in reality
- Robustness to adversarial attacks
- Start with large learning rate to learn easy features, then decrease to learn low noise, hard-to-fit patterns

Take away trends

- More people active with neurosciences: ML to understand NS, and NS to understand ML
- Meta learning (learning to learn)
- Reinforcement learning is gaining ground. Other keywords: bandit, feedback, regret, control

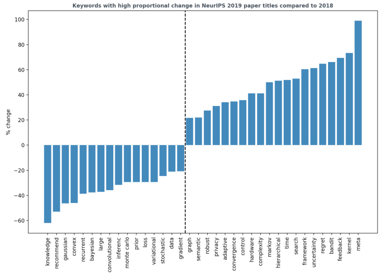

- Attributing uncertainty to ML algorithms (often with Bayesian methods in deep learning)
- Generative models still popular
- Hardware keyword on the rise, signaling more hardware-aware algorithms: hardware = bottleneck?
- "Recurrent and convolutional neural networks are literally so last year"
- Growing consciousness of potential impact on society

AI and society

- ML achieves super-human performance for well-designed problems, or games with score ⇒ where one can define a proper loss function or reward
- Scale to "real" problems?
  - explainability
  - causality
  - "moral" stand
  - culture, art
- Many advances in medical imaging, modelling of various phenomena, supernova analysis or LHC physics, but issues with:
  - out-of-distribution generalisation
  - scalability of computing resources, carbon footprint
  - reliability
  - decision bias (gender, race, etc.)
- Workshops/socials on Fairness & ethics, AI for Good, Tackling Climate Change with ML, AI for Humanitarian Assistance and Disaster Response, Safety and Robustness in Decision-making, ...
- Importance of personal decisions

Machine Learning and the Physical Sciences

- 91 short papers accepted for poster presentation (6 selected for talks)
- 70 "digital acceptance" papers (above rejection threshold but beyond capacity)
- 228 referees   web site (incl. videos)
- 5 invited speakers:
  - Alan Aspuru-Guzik: *Recent progress in ML for chemistry: SELFIES, inverse design of drug candidates and materials, and Bayesian algorithms for self-driving laboratories*
  - Yasaman Bahri: *Towards an understanding of wide, deep neural networks*
  - Katie Bouman: Cannot find title, about Event Horizon Telescope imaging technique
  - Bernhard Schölkopf: *Causality and Exoplanets*
  - Maria Schuld: *Innovating machine learning with near-term quantum computing*
  - Lenka Zdeborova: *Understanding machine learning via exactly solvable statistical physics models*

Suggested areas

- Application of machine and deep learning to physical sciences
- Generative models
- Likelihood-free inference
- Variational inference
- Simulation-based models
- Implicit models
- Probabilistic models
- Model interpretability
- Approximate Bayesian computation
- Strategies for incorporating prior scientific knowledge into machine learning algorithms
- Experimental design
- Any other area related to the subject of the workshop

Hidden information

- Predicting the Politics of an Image Using Webly Supervised Data

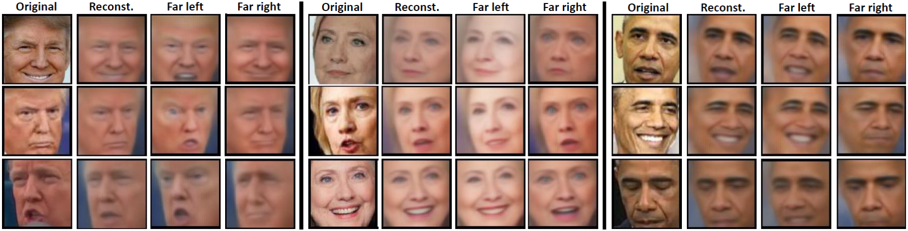

- Computational Mirrors: Blind Inverse Light Transport by Deep Matrix Factorization

(edited video)

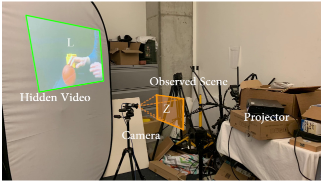

Decision trees are not dead!

- PIDForest: Anomaly Detection via Partial Identification ▸ NeurIPS
- A Debiased MDI (Mean Decrease of Impurity) Feature Importance Measure for Random Forests ▸ NeurIPS
- MonoForest framework for tree ensemble analysis ▸ NeurIPS
- Faster Boosting with Smaller Memory (Yoav S Freund) ▸ NeurIPS
- Minimal Variance Sampling in Stochastic Gradient Boosting ▸ NeurIPS
- Regularized Gradient Boosting ▸ NeurIPS
- Partitioning Structure Learning for Segmented Linear Regression Trees ▸ NeurIPS
- Random Tessellation Forests ▸ NeurIPS
- Optimal Sparse Decision Trees ▸ NeurIPS
- Provably robust boosted decision stumps and trees against adversarial attacks ▸ NeurIPS
- Robustness Verification of Tree-based Models ▸ NeurIPS

Conclusion

- Great location, great venue, great organisation
- A lot to process: "An individual reviewing that much content would need to read 43 pages every single day for a year" (NeurIPS 2019 General Chair Hanna Wallach)
- 75% of articles provide direct link to code implementation
  $\Rightarrow$ helps dissemination, reproducibility
- Some of it can/will be interesting for us (faster algos, hardware implementation, generative networks, uncertainty, ...)