

Bringing light to AI

Iacopo Poli – Lead Machine Learning Engineer

iacopo@lighton.io

Microsoft boss: World needs more computing power

By Joe Miller
BBC News, Davos



Thom Quinn
@tpq_



Is deep learning right for you? Now in 1 easy step:

(Q) Do you have more than 10,000 samples?

> If no -- sorry, you don't have enough samples

> If yes -- sorry, you don't have enough compute


INCREASING DEMAND OF COMPUTE



Eliot Andres

@EliotAndres

Follow

We just received the new iPhone 11! Wondering how it improved regarding machine learning? We put together a small benchmark. A thread 



ECOLOGICAL IMPACT OF AI



Dr Chloé Azencott
@cazencott

Follow

In a single day, I heard both Marc Schoenauer and @katecrawford discuss the ecological impact of AI and we need much more of this conversation.



Andrej Karpathy ✓
@karpathy

Follow

"Hybrid Optical-Electronic Convolutional Neural Networks" computationalimaging.org/publications/h ... incredibly interesting work - develops a hybrid optoelectronic CNN with an optical CONV1 layer that operates at zero power consumption (with rest of the forward pass in electronics (for now))

Green AI

Roy Schwartz, Jesse Dodge, Noah A. Smith, Oren Etzioni

(Submitted on 22 Jul 2019 (v1), last revised 13 Aug 2019 (this version, v3))

The role of artificial intelligence in achieving the Sustainable Development Goals

Ricardo Vinuesa, Hossein Azizpour, Iolanda Leite, Madeline Balaam, Virginia Dignum, Sami Domisch, Anna Felländer, Simone Langhans, Max Tegmark, Francesco Fuso Nerini

(Submitted on 30 Apr 2019)

OPTICAL PROCESSING UNIT



$$\mathbf{y} = |\mathbf{R}\mathbf{x}|^2 \quad R_{ij} \in \mathbb{C}$$

$$\text{Re}\{R_{ij}\} \sim \mathcal{N}(0, \sigma^2)$$

$$\text{Im}\{R_{ij}\} \sim \mathcal{N}(0, \sigma^2)$$

1M input – 1M output
Speed: 2 kHz
Power: 30W

Random Features for Large-Scale Kernel Machines


Ali Rahimi
Intel Research Seattle
Seattle, WA 98105
`ali.rahimi@intel.com`

Benjamin Recht
Caltech IST
Pasadena, CA 91125
`brecht@ist.caltech.edu`

Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions

Nathan Halko, Per-Gunnar Martinsson, Joel A. Tropp

LIGHT SCATTERING




Credit: Emmanuel Bossy- Simsonic Software



Georges de la Tour – Saint Joseph charpentier

LIGHT SCATTERING




Credit: Emmanuel Bossy- Simsonic Software



Georges de la Tour – Saint Joseph charpentier

LIGHT SCATTERING



Credit: Emmanuel Bossy- Simsonic Software



Georges de la Tour – Saint Joseph charpentier

Model-Free Episodic Control

Charles Blundell
Google DeepMind
cblundell@google.com

Benigno Uria
Google DeepMind
buria@google.com

Alexander Pritzel
Google DeepMind
apritzel@google.com

Yazhe Li
Google DeepMind
yazhe@google.com

Avraham Ruderman
Google DeepMind
aruderman@google.com

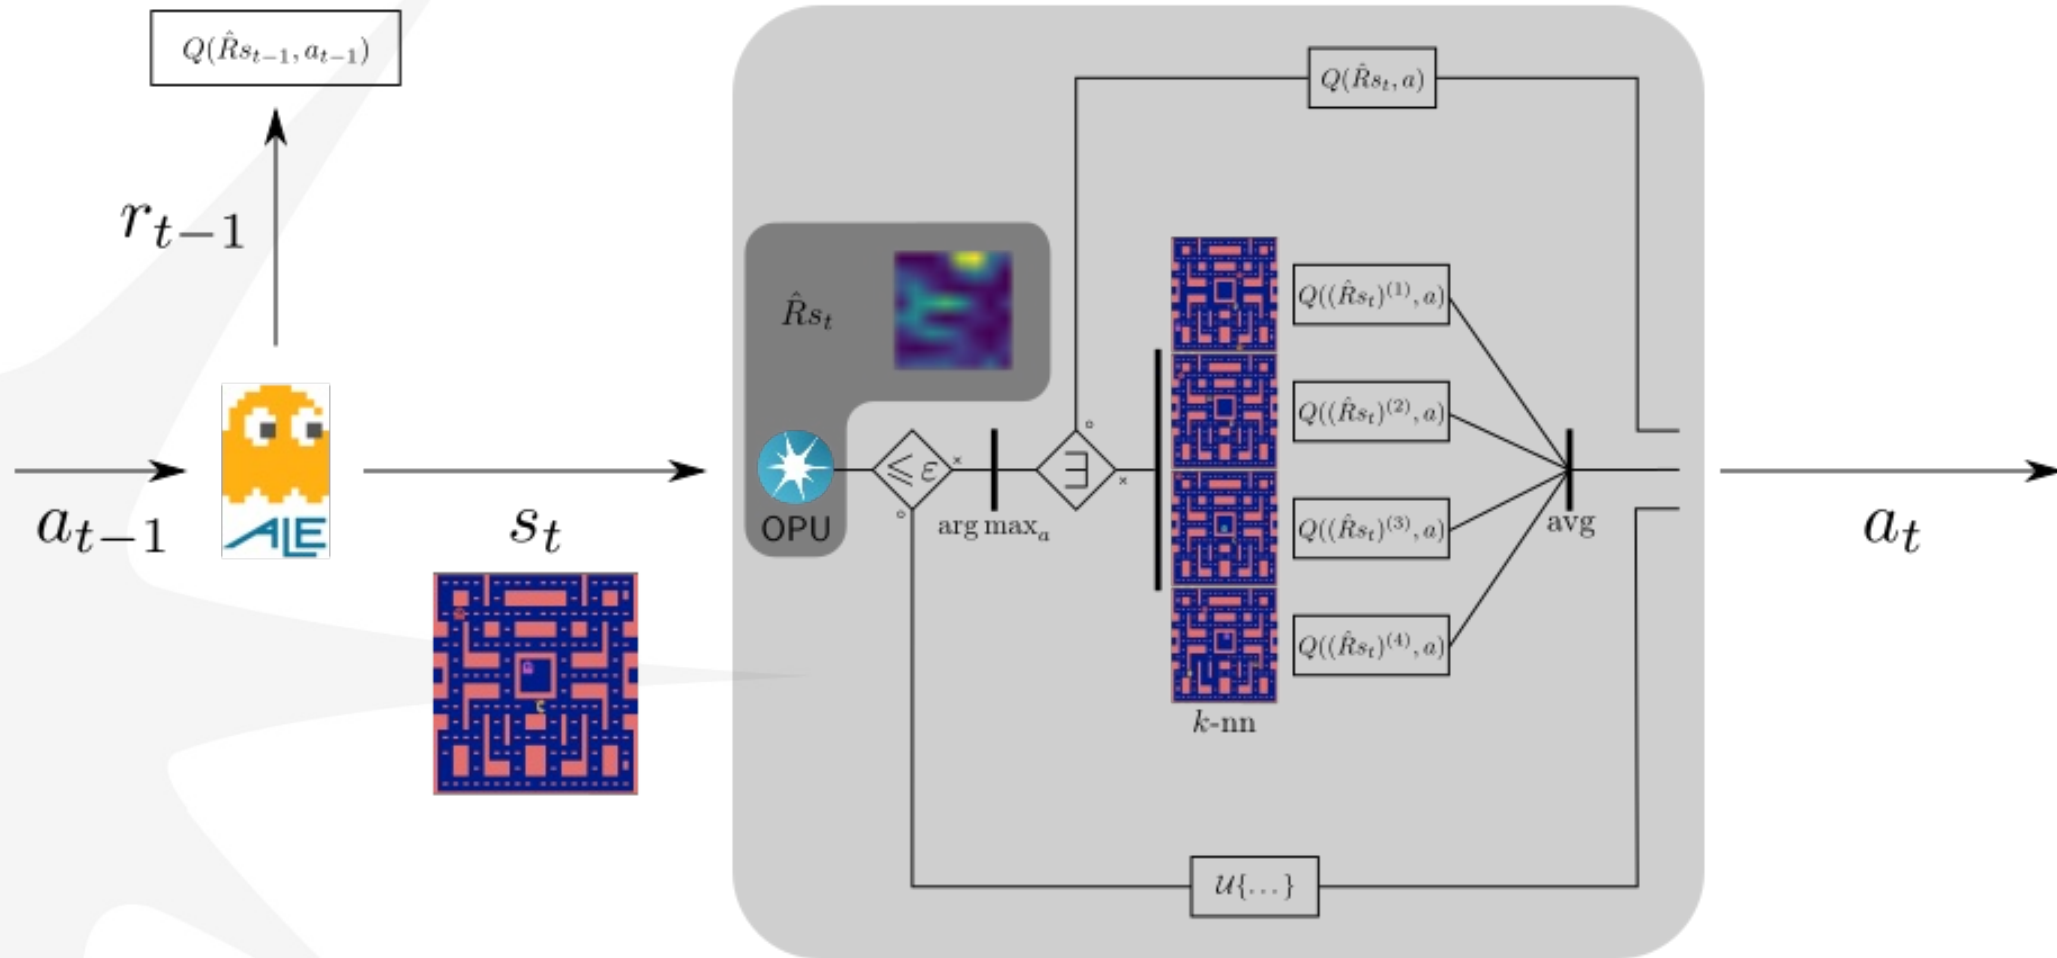
Joel Z Leibo
Google DeepMind
jzl@google.com

Jack Rae
Google DeepMind
jwrae@google.com

Daan Wierstra
Google DeepMind
wierstra@google.com


Demis Hassabis
Google DeepMind
demishassabis@google.com

REINFORCEMENT LEARNING




Credit: Martin Graive - Lighton

REINFORCEMENT LEARNING




Credit: Martin Graive - Lighton

REINFORCEMENT LEARNING



Credit: Martin Graive - Lighton

REINFORCEMENT LEARNING



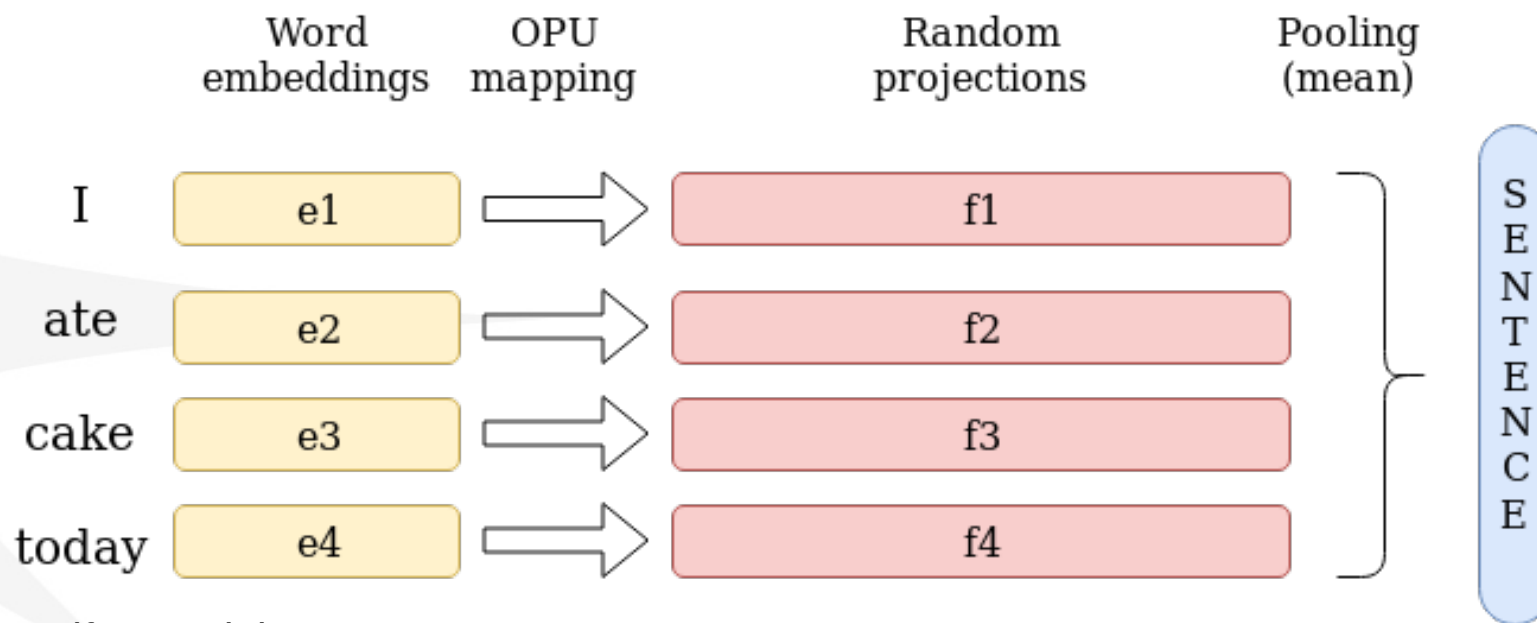
Credit: Martin Graive - Lighton

NOT ONLY IMAGES...

NO TRAINING REQUIRED: EXPLORING RANDOM ENCODERS FOR SENTENCE CLASSIFICATION

John Wieting*
Carnegie Mellon University
jwieting@cs.cmu.edu

Douwe Kiela
Facebook AI Research
dkiela@fb.com



Credit: François Boniface - Lighton

Shedding Light on the “Grand Débat”



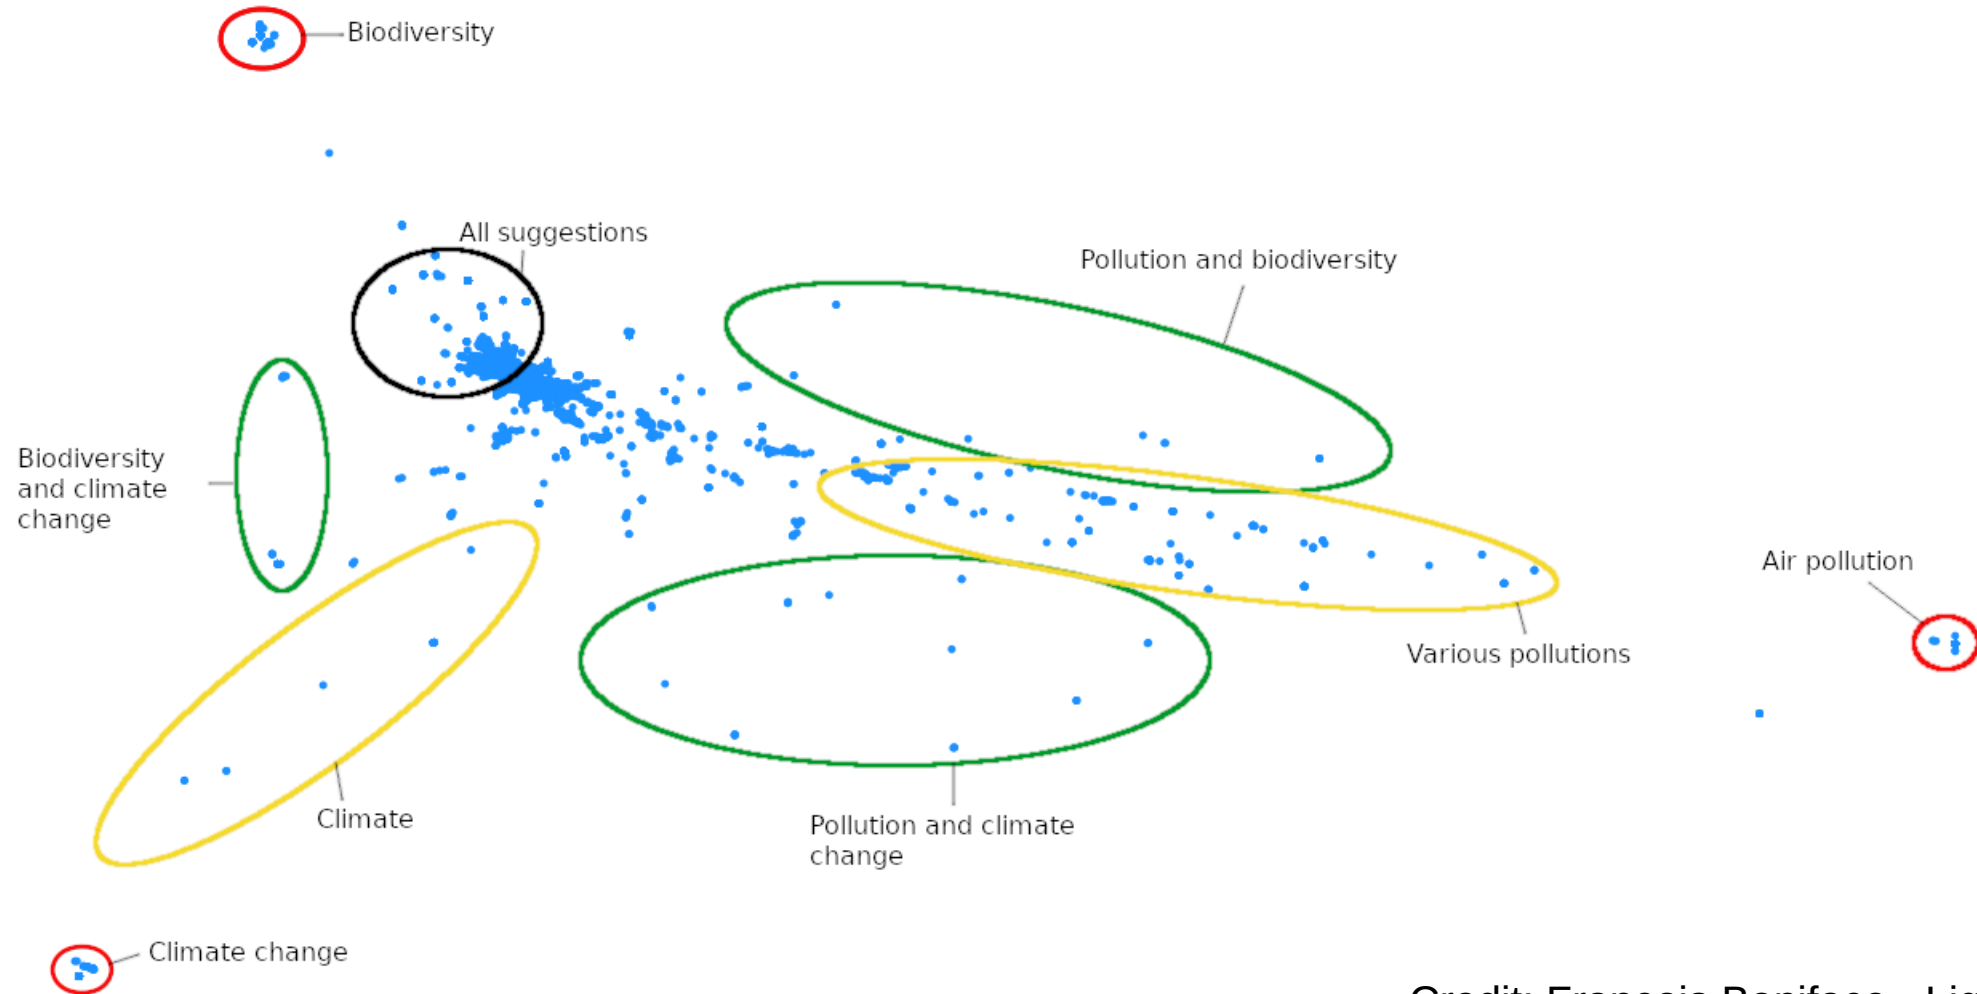
LightOn

Apr 11 · 11 min read



Credit: François Boniface - Lighton

NOT ONLY IMAGES...



Credit: François Boniface - Lighton

TRY IT OUT !

LightOn

CLOUD

Thank you !