DeepLearn School - Guimaraes

Matthieu François - Simon Lebeaud - Nicolas Tirel

GreenAI U.P.P.A.

May 2, 2022

1. DeepLearn 2022 Spring
   - Figures
   - Structure
2. Graph Lecture
3. Biometrics Trustworthiness
4. Medicine, high physics and Weather
5. Ecology survey

1. DeepLearn 2022 Spring
   - Figures
   - Structure
2. Graph Lecture
3. Biometrics Trustworthiness
4. Medicine, high physics and Weather
5. Ecology survey

Figures

- 3 keynotes

Figures

- 3 keynotes
- 21 lectures

Figures

- 3 keynotes
- 21 lectures
  3*1h30 lecture each
  - Interpretability
  - Generative Models
  - Vision
  - Audio
  - ...

Figures

- 3 keynotes
- 21 lectures
  3*1h30 lecture each
  - Interpretability
  - Generative Models
  - Vision
  - Audio
  - ...
- 160 participants

1. DeepLearn 2022 Spring
   - Figures
   - Structure
2. Graph Lecture
3. Biometrics Trustworthiness
4. Medicine, high physics and Weather
5. Ecology survey

Structure

Institute for Research Development, Training and Advice

AlCoB

International Conference on Algorithms for Computational Biology

BigDat

International School on Big Data

DeepLearn

International School on Deep Learning

LATA

International Conference on Language and Automata Theory

SLSP

International Conference on Statistical Language and Speech

TPNC

International Conference on the Theory and Practice of Natural

Presentation

Lecturer: Michalis Vazirgiannis
Theme: Graph Mining - generators & community detection

Applications fields

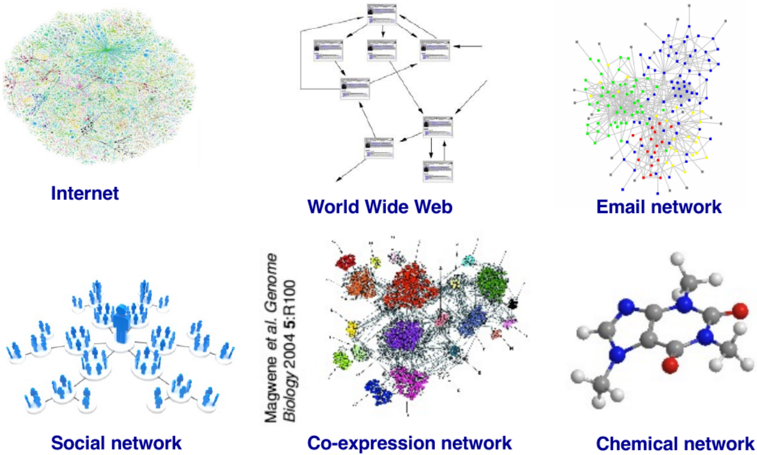

1. DeepLearn 2022 Spring
2. Graph Lecture
   - Presentation
   - Applications fields
   - Graph types
   - Some Methods
   - Graph Embedding
   - DeepLearning based embedding
3. Biometrics Trustworthiness
4. Medicine, high physics and Weather
5. Ecology survey

Graph types

Overview of graph types : directed, unidirected, complete, tree, bipartite graphs...

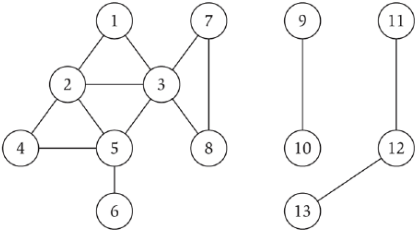

Isomorphism = best similarity between 2 graphs

Graph Generation

Aim: simulate graph data with same distribution / patern

- Regarding dregree law

Graph Generation

Aim: simulate graph data with same distribution / patern

- Regarding dregree law

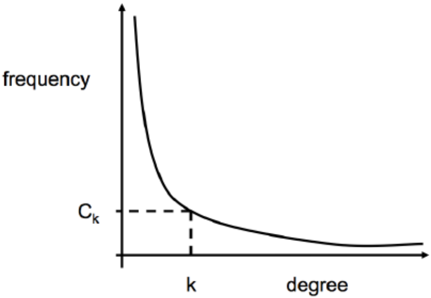

We visualise number of neighbors for each node in graph

Graph Generation

Aim: simulate graph data with same distribution / patern

- Regarding dregree law

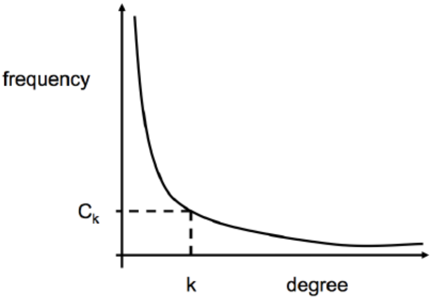

We visualise number of neighbors for each node in graph

- Regarding subpatern (Kronecker model)

Graph Generation

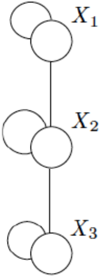

Graph G_1

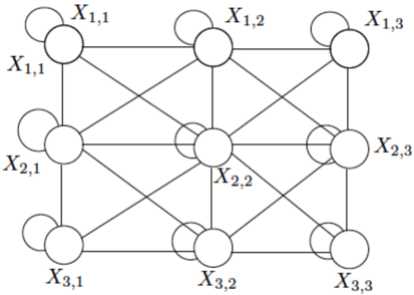

Graph $G_2 = G_1 \boxtimes G_1$

Graph Generation

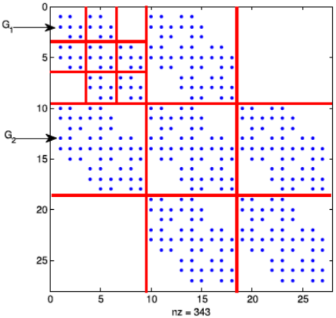

(α) $A(G_3) = A(G_2) \otimes A(G_1)$

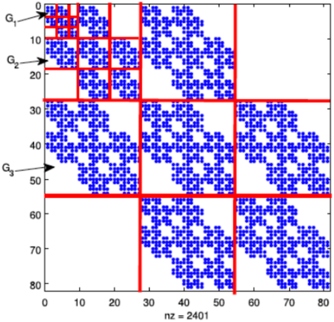

(β) $A(G_4) = A(G_3) \otimes A(G_1)$

Clustering

- based on number of triangles
- based on number of edges

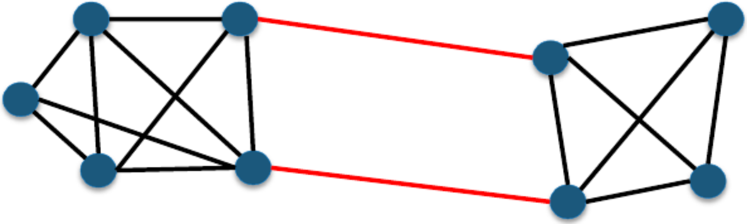

Clustering

- based on modularity (Louvain, Newman-Girvan)
- based on deeplearning (auto-encoder & node embedding)

Embedding & Similarity study

Aim

Create Embedding vectors for nodes, keeping proximity and similarity edges between them.

Embedding & Similarity study

Aim

Create Embedding vectors for nodes, keeping proximity and similarity edges between them.

How to?

We could take the adjacency matrix as embedding but inner product between them would produce a lot of 0.

$$\begin{pmatrix} 0 & 1 & \dots & 0 \\ 1 & 0 & \dots & 1 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 1 & \dots & 0 \end{pmatrix}$$

Embedding & Similarity study

Many methods:

- based on kernel

Embedding & Similarity study

Many methods:

- based on kernel
  Better than DeepLearning Methods but not scalable on big data.

Embedding & Similarity study

Many methods:

- based on kernel
  Better than DeepLearning Methods but not scalable on big data.
- based on deep learning

Laplacian eigenmaps

Early Method: matrix-factorization using **Laplacian eigenmaps**

Laplacian eigenmaps

Early Method: matrix-factorization using **Laplacian eigenmaps**

- Laplacian eigenmaps projects two nodes i and j close to each other when the weight of the edge between the two nodes A_{ij} is high

Laplacian eigenmaps

Early Method: matrix-factorization using **Laplacian eigenmaps**

- Laplacian eigenmaps projects two nodes i and j close to each other when the weight of the edge between the two nodes A_{ij} is high

$$y^* = argmin \sum (y_i - y_j)^2 A_{ij} \tag{1}$$

With A_ij the edge weight between i and j.

DeepWalk

Recent Method: inspired by Language Modeling **DeepWalk**

DeepWalk

Recent Method: inspired by Language Modeling **DeepWalk**

- Employ random walks to capture structural relationships between nodes. Each sample can be considered as a "sentence" of a document in NLP.

DeepWalk

Recent Method: inspired by Language Modeling **DeepWalk**

- Employ random walks to capture structural relationships between nodes. Each sample can be considered as a "sentence" of a document in NLP.

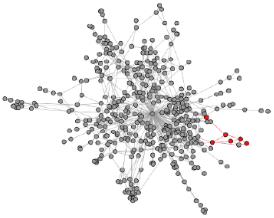

$\rightarrow$

$$v_5 \rightarrow v_8 \rightarrow v_{32} \rightarrow v_{28} \rightarrow v_6 \rightarrow v_{10} \rightarrow v_9$$
$$v_3 \rightarrow v_5 \rightarrow v_{28} \rightarrow v_8 \rightarrow v_9 \rightarrow v_{10} \rightarrow v_{25}$$
$$v_{20} \rightarrow v_{10} \rightarrow v_{12} \rightarrow v_6 \rightarrow v_8 \rightarrow v_4 \rightarrow v_5$$
$$v_{23} \rightarrow v_5 \rightarrow v_{32} \rightarrow v_{10} \rightarrow v_8 \rightarrow v_3 \rightarrow v_1$$
$$v_4 \rightarrow v_3 \rightarrow v_1 \rightarrow v_5 \rightarrow v_1 \rightarrow v_{12} \rightarrow v_{10}$$
$$\vdots$$

Why Random Walk?

Why Random Walk?

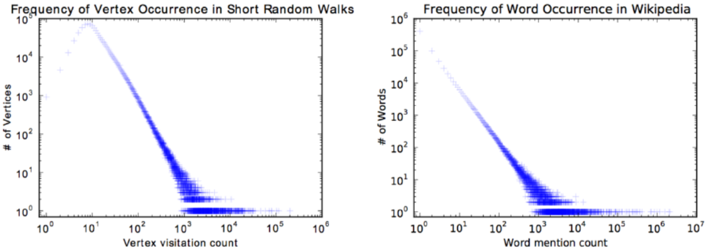

(a) YouTube Social Graph (b) Wikipedia Article Text

RandomWalk nodes distribution is the same as words in Wikipedia article.

DeepWalk

Next step: Use skipgram to generate embedding

- two sampling strategies, to capture:

DeepWalk

Next step: Use skipgram to generate embedding

- two sampling strategies, to capture:
    - structure: depth-first sampling (DFS)
    - local similarity: breadth-first sampling (BFS)

DeepWalk

Next step: Use skipgram to generate embedding

- two sampling strategies, to capture:
  - structure: depth-first sampling (DFS)
  - local similarity: breadth-first sampling (BFS)

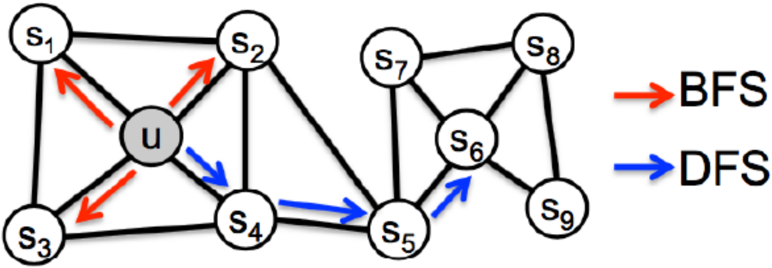

DeepWalk

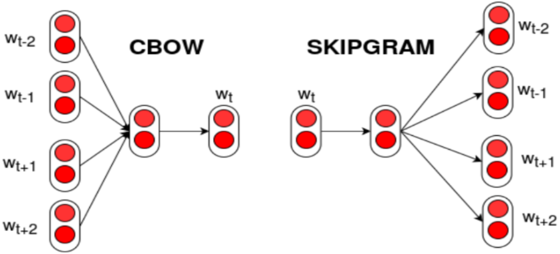

$$\frac{1}{T}\sum_{t=1}^{T}\sum_{-c\leq j\leq c, j\neq 0} logp(w_t|w_{t+j})$$

$$\frac{1}{T}\sum_{t=1}^{T}\sum_{-c\leq j\leq c, j\neq 0} logp(w_{t+j}|w_t)$$

Other Methods

It exists also

- GraRep
- Line
- SDNE

Other Methods

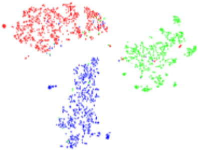

(a) *SDNE*

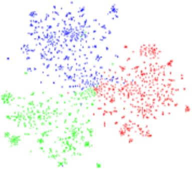

(b) *LINE*

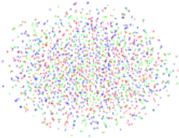

(c) *DeepWalk*

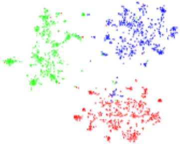

(d) *GraRep*

Challenges

- Proofing face recognition for security control (public security, border control, building access, face purchase...)

Challenges

- Proofing face recognition for security control (public security, border control, building access, face purchase...)
- Remove ubiquitous presence of deepfake and preventing rapid dissemination of "fake news"

Attack types

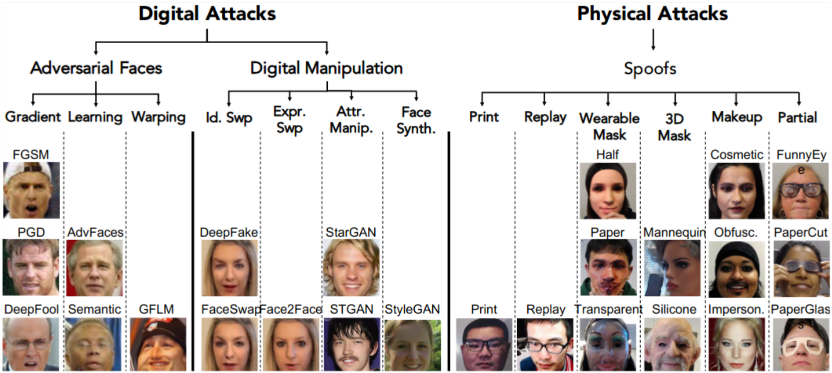

Anti Physical Spoofing Methods

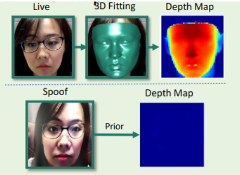

Figure 1: Depth Estimation

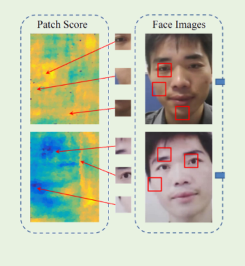

Figure 2: Patch-based CNNs

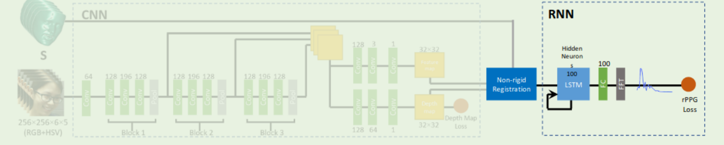

Figure 3: rPPG Estimation

Generative FAS

- CNN is trained to generate a spoof trace image
- Spoof is considered as an added layer to your original image
- If we can generate the spoof layer from a false image then we are capable of taking it of the image and then "recover" the original image

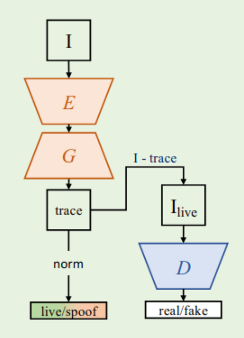

Attacks

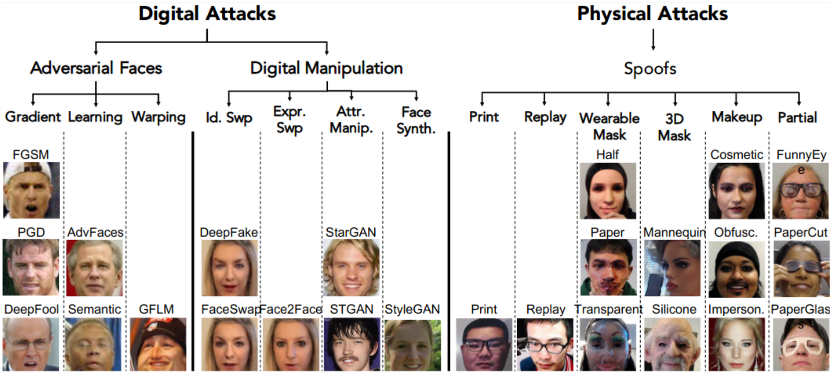
Digital Attacks
Physical Attacks
Adversarial Faces
Digital Manipulation
Spoofs
Gradient
Learning
Warping
Id. Swp
Expr. Swp
Attr. Manip.
Face Synth.
Print
Replay
Wearable Mask
3D Mask
Makeup
Partial
FGSM
Half
Cosmetic
FunnyEye
PGD
AdvFaces
DeepFake
StarGAN
Paper
Mannequin
Obfusc.
PaperCut
DeepFool
Semantic
GFLM
FaceSwap
Face2Face
STGAN
StyleGAN
Print
Replay
Transparent
Silicone
Imperson.
PaperGlass

Dynamic Methods -> Inconsistent Head Poses

Figure 4: Xin et al. Exposing deep fakes using inconsistent head poses

Static Methods -> Face X-Ray

Detect images altered by the blending of a new face on the original

Figure 5: Li et al. Face X-ray for more general face forgery detection

Multi manipulation detection

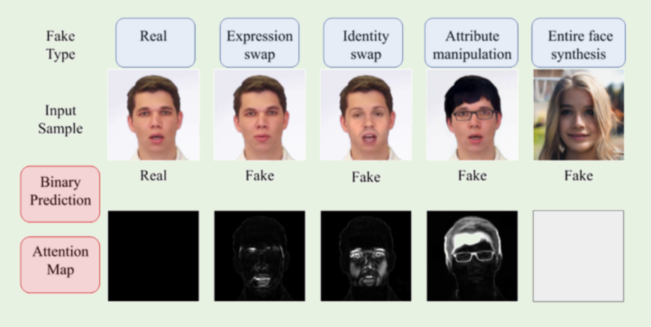

Figure 6: Dang et al. On the Detection of Digital Face Manipulation

Multi manipulation detection

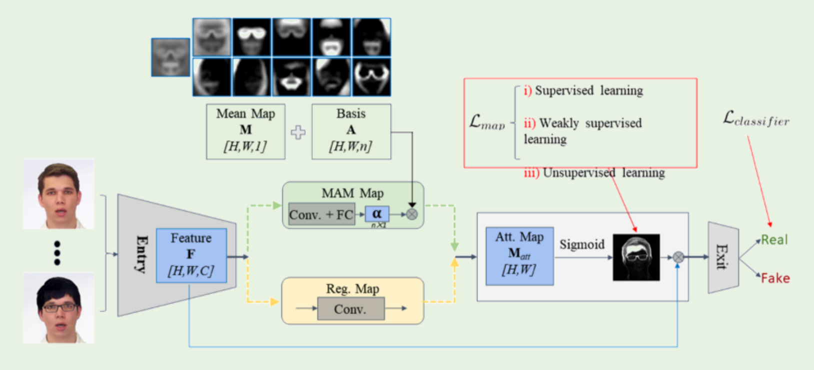

Figure 7: Dang et al. On the Detection of Digital Face Manipulation

Proactive schemes

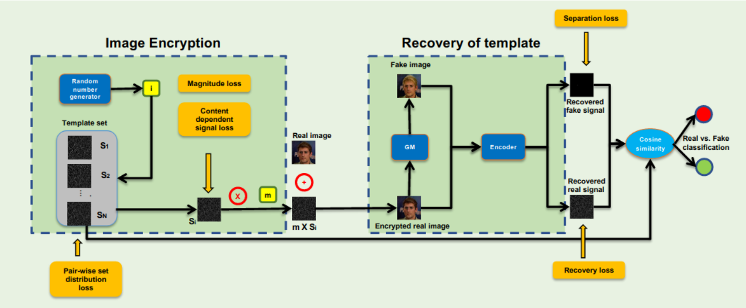

Figure 8: Framework to proactively counter digital spoofing

Adversarial Attacks

- Obfuscation attack: Falsely reject a genuine subject
- Impersonation attack: falsely match to an impostor subject

Solutions:

- Detection
- Purification
- Robustness

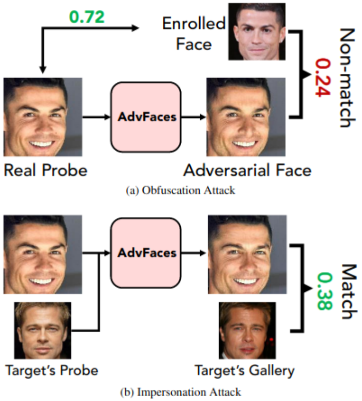

Figure 9: AdvFaces: Adversarial Face Synthesis

Unified Detection of Digital and Physical Face Attacks

- Cluster each types of attacks with high similarities to form groups
- Train a multi-branch model each branch trying to do classification for a cluster of attacks
- Fuse the classification for each cluster

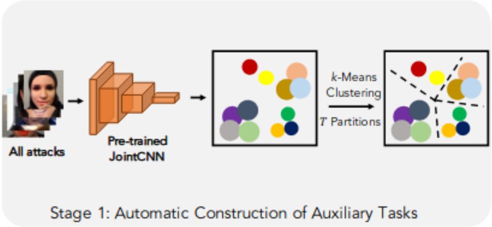

Figure 10: JointCNN: 53.19% TDR

Proposed UniFAD

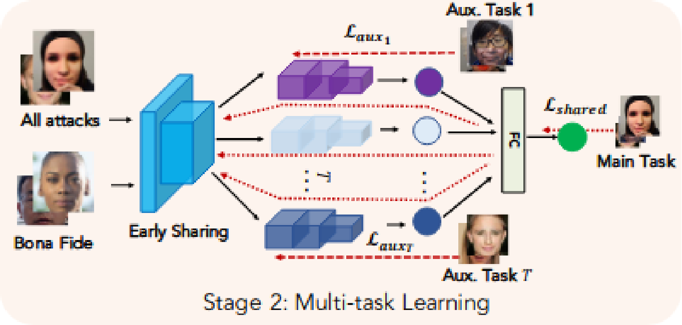

Figure 11: Debayan Deb, Xiaoming Liu, Anil Jain, Unified Detection of Digital and Physical Face Attacks

UniFAD Results

	TDR (%) @ 0.2% FDR	Year	Proposed For	Adv.	Dig. Man.	Phys.	Overall	Time (ms)
w/o Re-train	FaceGuard	2020	Adversarial	**99.91**	22.28	00.58	29.64	01.41
	FFD	2020	Digital Manipulation	09.49	**94.57**	01.25	34.55	11.57
	SSRFCN	2020	Spoofs	00.25	00.76	93.19	22.71	02.22
	MixNet	2020	Spoofs	00.36	09.83	78.21	21.12	12.47
Baselines	FaceGuard	2020	Adversarial	**99.86**	41.56	04.35	56.69	01.41
	FFD	2020	Digital Manipulation	76.06	91.32	87.43	68.25	11.57
	SSRFCN	2020	Spoofs	08.23	27.67	89.19	43.26	02.22
	One-class	2020	Spoofs	04.81	45.96	79.32	39.40	07.92
	MixNet-*UniFAD*	2021	All	82.33	91.59	**94.60**	**90.07**	12.47
Fusion Schemes	Cascade	–	–	88.39	81.98	69.19	77.46	05.16
	Min-score	–	–	03.65	11.08	00.43	07.22	16.14
	Median-score	–	–	10.87	42.33	47.19	39.48	16.12
	Mean-score	–	–	14.53	47.18	61.32	38.23	16.12
	Max-score	–	–	85.32	61.93	56.87	73.89	16.13
	Sum-score	–	–	74.93	58.01	50.34	69.21	16.11
	LightGBM	–	–	76.25	81.28	88.52	85.97	17.92
	Proposed UniFAD	2021	All	92.56	**97.21**	**98.76**	**94.73**	02.59

Figure 12: Debayan Deb, Xiaoming Liu, Anil Jain, Unified Detection of Digital and Physical Face Attacks

1. DeepLearn 2022 Spring
2. Graph Lecture
3. Biometrics Trustworthiness
4. Medicine, high physics and Weather
   - Sofia Vallecorsa- Generative Models in High Energy Physics: Examples from CERN
   - Lucila Ohno-machado - Use of Predictive Models in Medicine and Biomedical Research
   - Rylan Conway - Deep Learning for Digital Assistants
   - Martin Schultz - Deep Learning for Air Quality, Weather and Climate
5. Ecology survey

Generative models

Examples in science

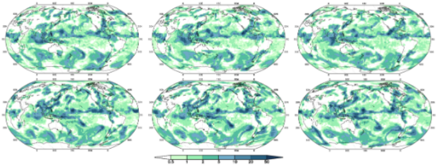

Ayala, Alexis, et al. "Loosely Conditioned Emulation of Global Climate Models With Generative Adversarial Networks." *arXiv:2105.06386* (2021).

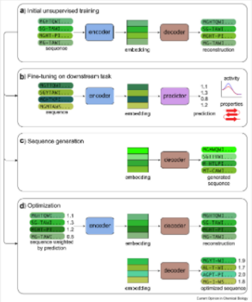

Wu, Zachary, et al. "Protein sequence design with deep generative models." *Current opinion in chemical biology* 65 (2021): 18-27

CERN openlab

Replace with DNN

Accelerating the training process

- Introducing techniques to **parallelise** training
- **Data parallelism**
  - Compute gradients on several batches independently
  - Update the model synchronously or asynchronously
- **Model Parallelism**, **Hybrid techniques**
- Use **reduced precision** representation (INT6, BF16, …)
- Extreme parallelism using **combined strategies** and SGD algorithm optimisation
  - DeepSpeed and ZeRO-2 on Microsoft Azure

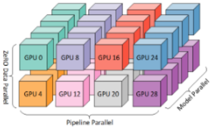

https://www.microsoft.com/en-us/research/blog/deepspeed-extreme-scale-model-training-for-everyone/

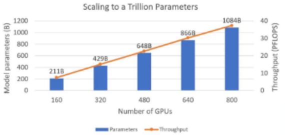

CERN openlab

Sustainable AI

Sustainable AI

- AI inference more **energy efficient** than classical algorithms
- Energy cost of **AI training** can be high
- The community is defining **best practices**[1]
  - **Efficient AI architectures** can reduce computation by 3x–10x.
  - **AI-optimized processors vs general-purpose** can improve energy efficiency by 2x–5x[2].
  - **Cloud computing vs on-prem** reduces energy usage by 1.4x–2x
- **Efficient training strategies**
  - Self-supervision, few-short learning, pre-training

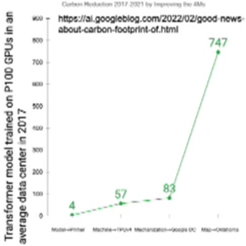

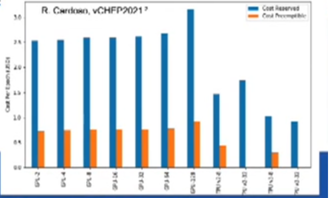

CERN openlab

1 [illegible] Carbon Footprint of Machine Learning Training Will Plateau, Then Shrink." (2022).
2 Cardoso, Renato, et al. "Accelerating GAN training using highly parallel hardware on public cloud." EPJ Web of Conferences. Vol. 251. EDP Sciences, 2021.

Detector Response

Monte Carlo simulation

Detector response as images

Monte Carlo simulation of detector response is extremely **demanding in terms of computing resources**

→ 50 % of LHC Computing Grid resources today

Interpret **detector output as images**

Sensors outputs become pixels in a image

Use **computer vision techniques** to interpret results

Replace Monte Carlo approach with **Generative Models**

Interfacing DL to standard software is not trivial!

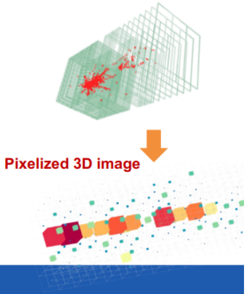

CERN openlab 04.04.22 12

CaloGAN

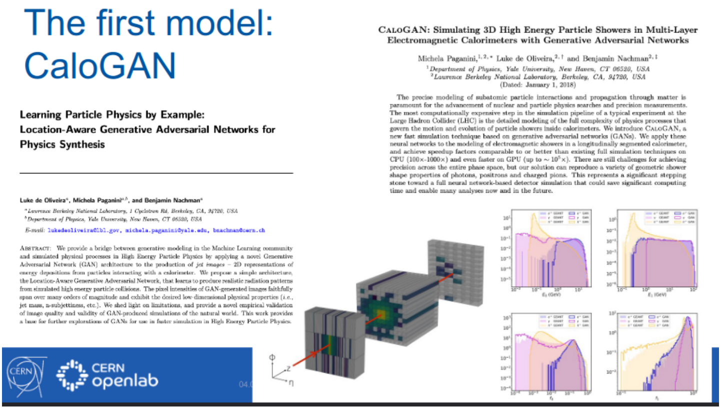
The first model: CaloGAN
Learning Particle Physics by Example: Location-Aware Generative Adversarial Networks for Physics Synthesis
Luke de Oliveira[a], Michela Paganini[a,b], and Benjamin Nachman[a]
[a] Lawrence Berkeley National Laboratory, 1 Cyclotron Rd, Berkeley, CA, 94720, USA
[b] Department of Physics, Yale University, New Haven, CT 06520, USA
E-mail: lukedeoliveira@lbl.gov, michela.paganini@yale.edu, bnachman@cern.ch
Abstract: We provide a bridge between generative modeling in the Machine Learning community and simulated physical processes in High Energy Particle Physics by applying a novel Generative Adversarial Network (GAN) architecture to the production of jet images – 2D representations of energy depositions from particles interacting with a calorimeter. We propose a simple architecture, the Location-Aware Generative Adversarial Network, that learns to produce realistic radiation patterns from simulated high energy particle collisions. The pixel intensities of GAN-generated images faithfully span over many orders of magnitude and exhibit the desired low-dimensional physical properties (i.e., jet mass, n-subjettiness, etc.). We shed light on limitations, and provide a novel empirical validation of image quality and validity of GAN-produced simulations of the natural world. This work provides a base for further explorations of GANs for use in faster simulation in High Energy Particle Physics.
CERN openlab
CaloGAN: Simulating 3D High Energy Particle Showers in Multi-Layer Electromagnetic Calorimeters with Generative Adversarial Networks
Michela Paganini,[1,2,*] Luke de Oliveira,[2,†] and Benjamin Nachman[2,‡]
[1] Department of Physics, Yale University, New Haven, CT 06520, USA
[2] Lawrence Berkeley National Laboratory, Berkeley, CA, 94720, USA
(Dated: January 1, 2018)
The precise modeling of subatomic particle interactions and propagation through matter is paramount for the advancement of nuclear and particle physics searches and precision measurements. The most computationally expensive step in the simulation pipeline of a typical experiment at the Large Hadron Collider (LHC) is the detailed modeling of the full complexity of physics processes that govern the motion and evolution of particle showers inside calorimeters. We introduce CaloGAN, a new fast simulation technique based on generative adversarial networks (GANs). We apply these neural networks to the modeling of electromagnetic showers in a longitudinally segmented calorimeter, and achieve speedup factors comparable to or better than existing full simulation techniques on CPU (100×-1000×) and even faster on GPU (up to ~ 10^5×). There are still challenges for achieving precision across the entire phase space, but our solution can reproduce a variety of geometric shower shape properties of photons, positrons and charged pions. This represents a significant stepping stone toward a full neural network-based detector simulation that could save significant computing time and enable many analyses now and in the future.
φ
z
η

1. DeepLearn 2022 Spring
2. Graph Lecture
3. Biometrics Trustworthiness
4. Medicine, high physics and Weather
   - Sofia Vallecorsa- Generative Models in High Energy Physics: Examples from CERN
   - Lucila Ohno-machado - Use of Predictive Models in Medicine and Biomedical Research
   - Rylan Conway - Deep Learning for Digital Assistants
   - Martin Schultz - Deep Learning for Air Quality, Weather and Climate
5. Ecology survey

Introduction

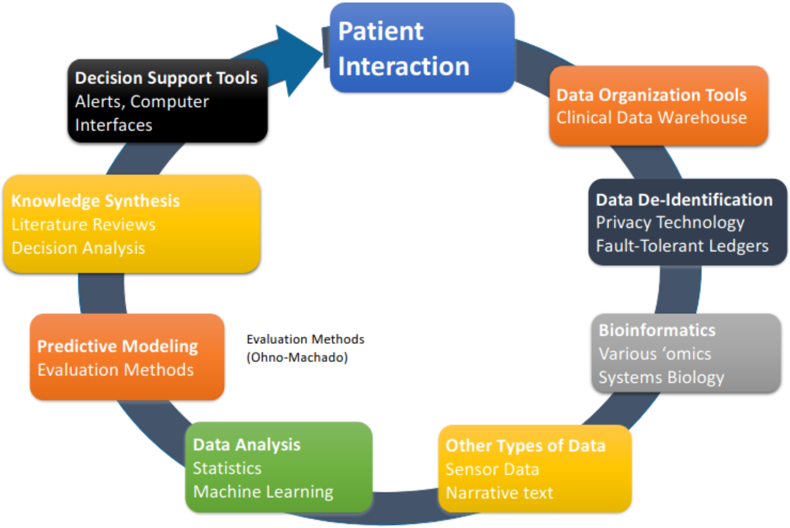

Data sharing

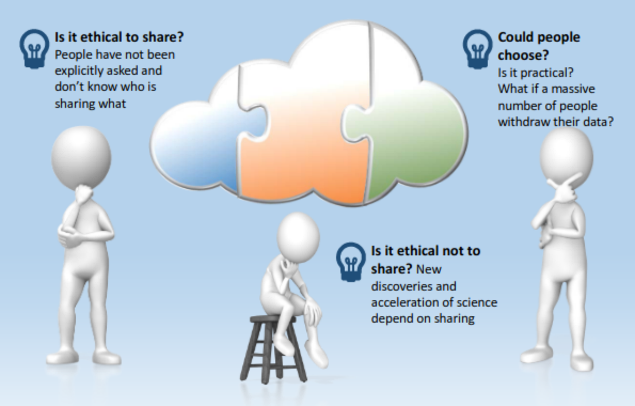

Data sharing

Privacy vs. Data Sharing

- Personalized medicine depends on big data
- Getting enough data to study rare diseases is already difficult without thinking of privacy risks
- It is difficult to quantify the privacy risk and potential benefits

- Streamlined access-controlled sharing: quickly determine who will use the data and why, and provide controlled access
  - Authenticate Users & Authorize Users
  - Monitor Use

Source: DOE

De-identification

People think *re-identification* means finding the names (of everyone) in the database

	biometric	diagnosis	income
Lisa	ABDSFHG	pregnancy	60k
Mike	BQEHGKK	rare disease 1	100k
Alice	WOEIMIV	depression	20k

	A	B	diagnosis	income
Lisa	10	20	pregnancy	60k
Mike	10	21	rare disease 1	100k
Alice	11	20	depression	20k

Re-identification

But they forget someone can target an individual

biometric	diagnosis	income
ABDSFHG	pregnancy	60k
BQEHGKK	rare disease 1	100k
WOEIMIV	depression	20k

I know something about the target person

A	B	diagnosis	income
10	20	pregnancy	60k
10	21	rare disease 1	100k
11	20	depression	20k

But I want to know more...

A	B	diagnosis	income
10	20	pregnancy	60k
10	21	rare disease 1	100k
11	20	depression	20k
10	20	pregnancy	20k

k anonymity, l-diversity, ...

1 DeepLearn 2022 Spring

2 Graph Lecture

3 Biometrics Trustworthiness

4 Medicine, high physics and Weather
   - Sofia Vallecorsa- Generative Models in High Energy Physics: Examples from CERN
   - Lucila Ohno-machado - Use of Predictive Models in Medicine and Biomedical Research
   - Rylan Conway - Deep Learning for Digital Assistants
   - Martin Schultz - Deep Learning for Air Quality, Weather and Climate

5 Ecology survey

Full ASR

ASR System Overview

This task is commonly approached by breaking it down into subproblems

- Wake Word Detection: identifying when the user is speaking to the assistant
- Acoustic modeling: transforming the raw audio signal into units of spoken language called "phonemes" (these usually map directly to sub-word tokens)
- Pronunciation modeling: mapping sequences of phonemes into words
- Language modeling: assign probabilities to word sequences
- Decoding: determine which word sequence is most likely to represent to the input audio signal

Key word selection

Wake Word Detection

- The first step in the ASR pipeline for digital assistants is Wake Word Detection
  - This is a special case of Keyword Spotting (KWS), a larger subfield of ASR, with some unique challenges
    - It has to be done in real time
      - A high latency response is perceived as a False Negative
    - It has to work with very limited computational resources
      - These models need to be able to run on the device itself (i.e., they are not run in the cloud)
    - It needs to have high Precision
      - False Positives lead to a lot of user friction

Goal

Wake Word Models: DeepKWS

- The DeepKWS model used for Wake Word detection was released in 2014: [Small-Footprint Keyword Spotting Using Deep Neural Networs]

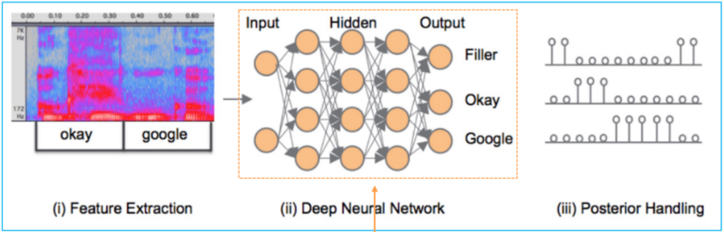

Fully-connected feedforward neural network with three hidden layers (128d) with ReLU activations and a softsax output layer

106

Reference: Small-Footprint Keyword Spotting Using Deep Neural Networs

CNN

Wake Word Detection as Image Classification

- The input LFBE features to the DeepKWS model can be thought of as an image!
- This means we can try to use image classification techniques to try and solve the Wake Word detection problem.
  - Image Classification → Convolutional Neural Networks!

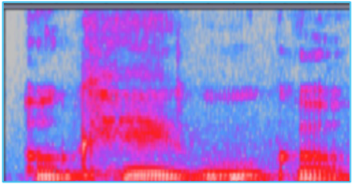

LBFE Features

CNN

Wake Word Detection Models using CNNs

- It was shown that using Convolutional Neural Networks (CNNs) can improve Wake Word detection even more (Google showed a ~40% improvement over DeepKWS).
    - These models are attractive in a limited resource setting as the computational complexity can be easily controlled by adjusting various hyperparameters (e.g., kernel size, pooling size, and stride)
    - The model below is the word-level keyword spotter (WL CNN)

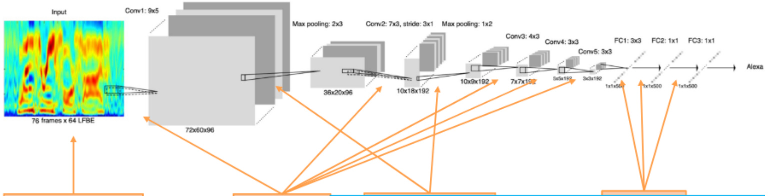

Input LFBE features

5 conv layers

References [Convolutional Neural Networks for Small-footprint Keyword Spotting]
[Accurate Detection of Wake Word Start and End Using a CNN]

Challenges

Wake Word Detection Challenges

- The models need to be able to run with very limited computational resources (in terms of compute power, available memory, and power constraints).
  - Model quantization (even down to 4 bits!), knowledge distillation, and pruning are all effective ways build effective models in this setting.
- The Wake Word needs to be detected in a far-field, noisy environment.
  - Robustness training is key! Augmenting training data with noise perturbations added to the input signals is a common way to improve performance.
  - Automatic Gain Control (AGC) can applied to the input where the signal is amplified whenever speech is detected.
  - Beam forming (using a microphone array to identifying the direction of the sound) can also be used to identify background noise.
    - When the direction of the signal jumps round → it's probably noise

A comprehensive reference on KWS[148]

1. DeepLearn 2022 Spring
2. Graph Lecture
3. Biometrics Trustworthiness
4. Medicine, high physics and Weather
   - Sofia Vallecorsa- Generative Models in High Energy Physics: Examples from CERN
   - Lucila Ohno-machado - Use of Predictive Models in Medicine and Biomedical Research
   - Rylan Conway - Deep Learning for Digital Assistants
   - Martin Schultz - Deep Learning for Air Quality, Weather and Climate
5. Ecology survey

Introduction

Atmospheric radiation budget

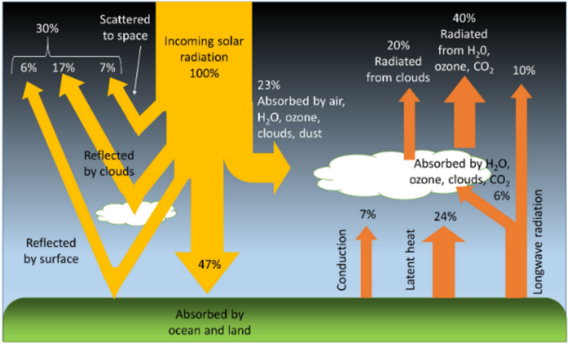

https://rwu.pressbooks.pub/app/uploads/sites/7/2017/01/figure8.1.1.png

Martin Schultz
Slide 9

JÜLICH Forschungszentrum

Data

Global weather observations: vertical profiles

According to WMO (~1,300 sites)

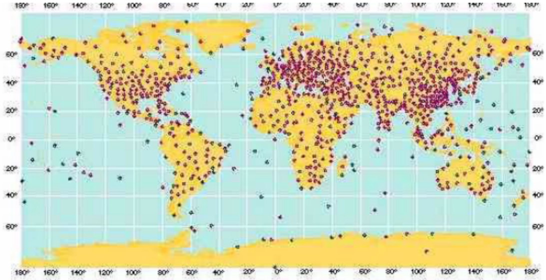

Map from https://public.wmo.int/en/programmes/global-observing-system

Martin Schultz
Slide 24

JÜLICH Forschungszentrum

One application, weather forecasting

Machine learning approaches to weather forecasting

Forecast problem

$$\mathbf{x}(t) \rightarrow \mathbf{x}(t + \Delta t)$$

$\mathbf{x}(t)$: State of the atmosphere at time t
$\mathbf{x}(t+\Delta t)$: Forecast of atmospheric state at lead time Δt

Martin Schultz
Slide 2

JÜLICH Forschungszentrum

One application, weather forecasting

NN Weather Forecasting

Precipitation nowcasting with radar data

Obs.

GNN

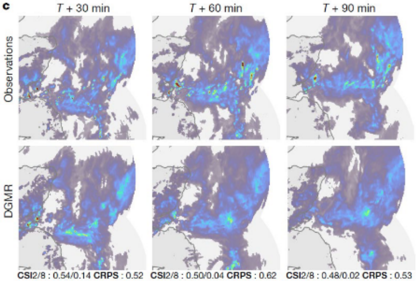

30, 60, and 90 minute forecasts of a complex precipitation event over Scotland (24 June 2019, 15:15 UTC)

In a case study, 90% of meteorologists chose output from generative model over other methods.

Ravuri et al. (2021), Nature https://doi.org/10.1038/s41586-021-03854-z

Martin Schultz
Slide 4

JÜLICH Forschungszentrum

Questions (21 answers)

- How much are you aware of climate change?
- Is climate change caused by humans?
- Do you know IPCC?
- Do you know Labos 1Point5/carbon footprint calculator?
- What was your main means of transport to go to Guimaraes?

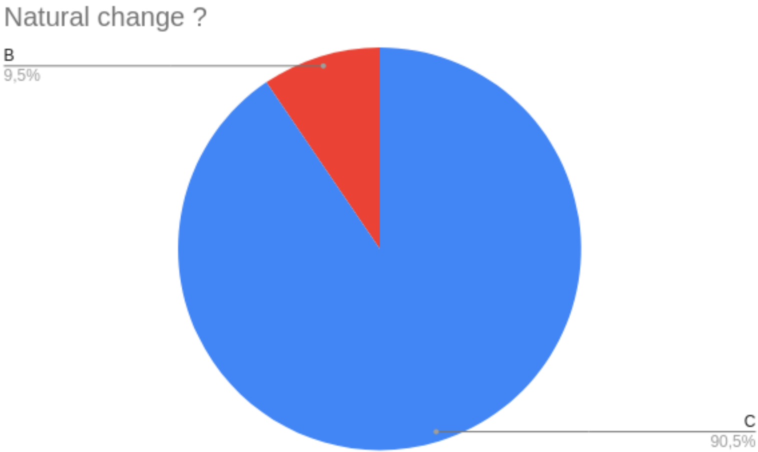

Figure 13: Is climate change caused by humans ?

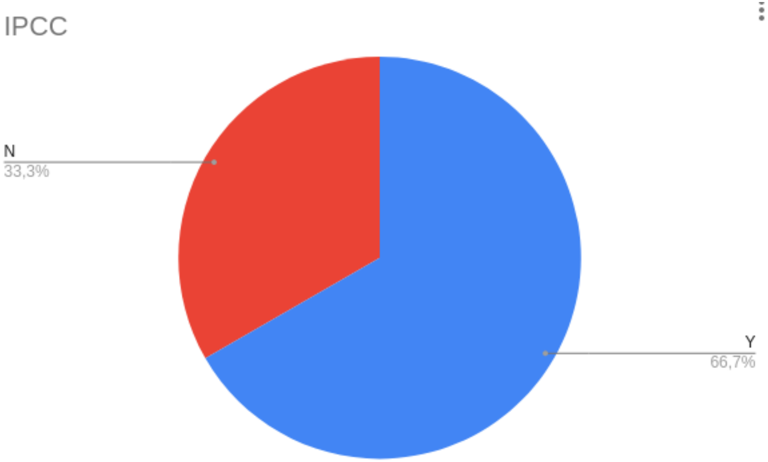

Figure 14: Do you know IPCC ?

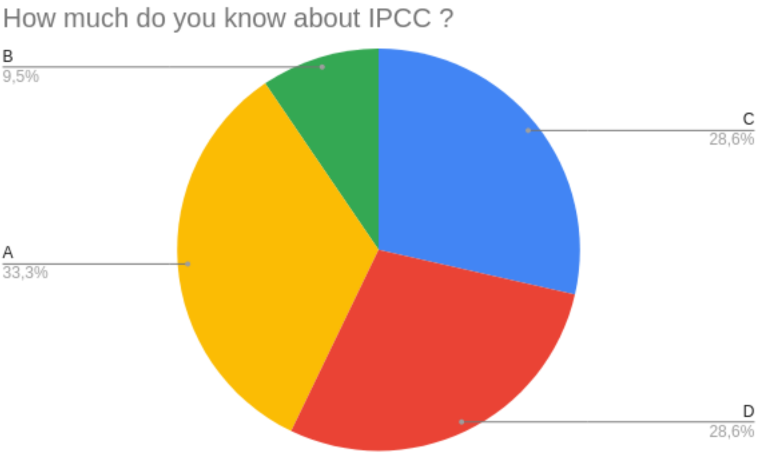

Figure 15: How much do you know about it ?

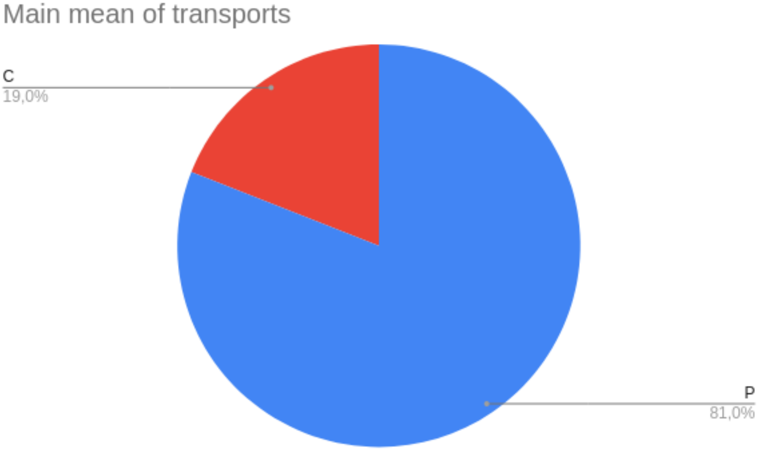

Figure 16: What was your main means of transport to go to Guimaraes ?

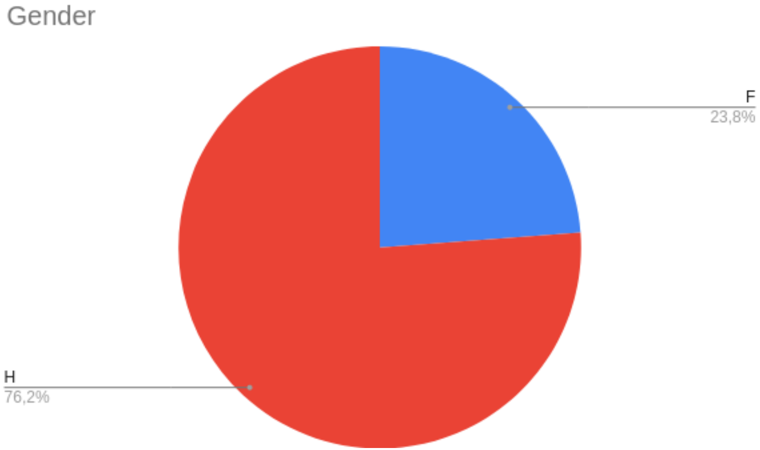

Figure 17: Gender distribution