Graph neural networks

Nikita Kazeev
Yandex School of Data Analysis
HSE University
Schaffhausen Institute of Technology

05.10.2021

Machine learning in high energy physics: a conversation over ice cream

About me

- Graduated from Yandex School of Data Analysis in 2015
- Graduated from Moscow Institute of Physics and Technology (MIPT) in 2016
- Defended PhD in 2020 on machine learning for particle identification in LHCb
- Teach at Machine Learning for High Energy Physics (MLHEP) summer schools since their inception in 2015

Plan for the next 20 minutes

- The conceptual problems solved by graph neural networks
- General algorithm for graph neural networks
- A (very brief) look at HEP applications

Toy (not really) problem: predict potential energy of a molecule

Naïve attempt 1

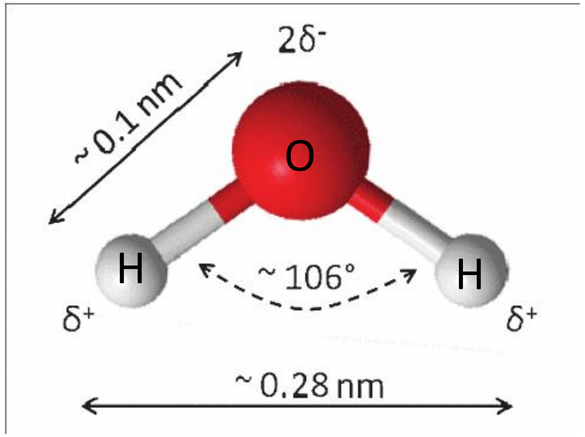

Source

Charge	x	y	z
1	1.1	-1.4	0
1	1.1	1.4	0
16	0	0	0

[1 1.4 -1.4 0 1 1.1 1.4 0 16 0 0 0]

CatBoost

Profit?

No profit

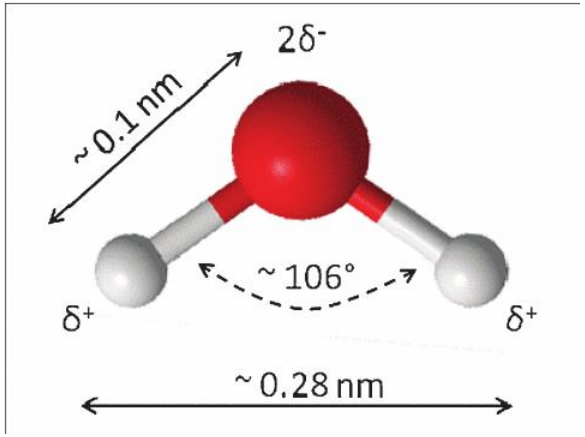

Charge	x	y	z
1	1.1	-1.4	0
1	1.1	1.4	0
16	0	0	0

= [1 1.4 -1.4 **0 1 1.1 1.4 0 16 0 0 0**]

Why not this?

Charge	x	y	z
1	1.1	-1.4	0
16	0	0	0
1	1.1	1.4	0

= [1 1.4 -1.4 **0 16 0 0 0 1 1.1 1.4 0**]

Toy (not really) problem: predict potential energy of a molecule

Naïve attempt 2

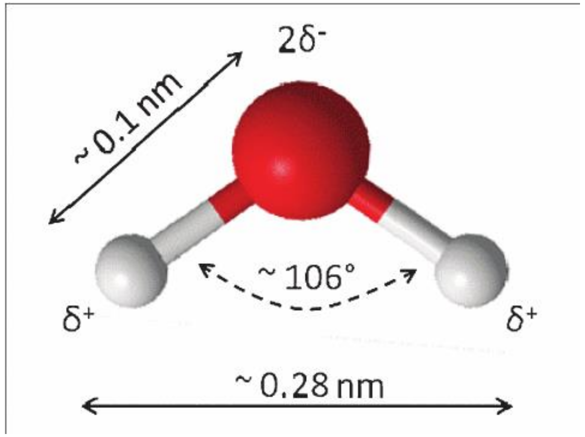

Source

Charge	x	y	z
1	1.1	-1.4	0
1	1.1	1.4	0
16	0	0	0

Augmented training with random permutation

[1 1.4 -1.4 0 1 1.1 1.4 0 16 0 0 0]

….

CatBoost

Profit?

No profit

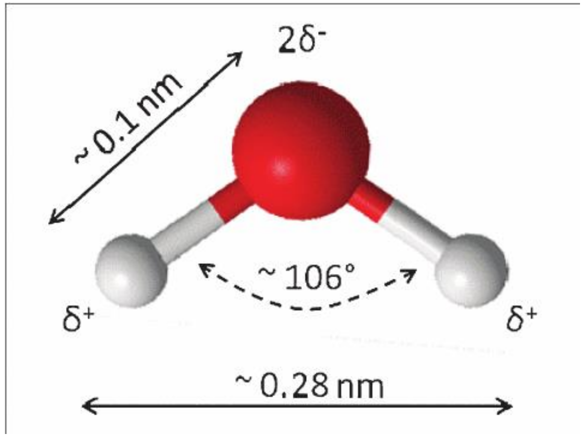

Charge	x	y	z
1	1.1	-1.4	0
1	1.1	1.4	0
16	0	0	0

12 features

= [1 1.4 -1.4 0 1 1.1 1.4 0 16 0 0 0]

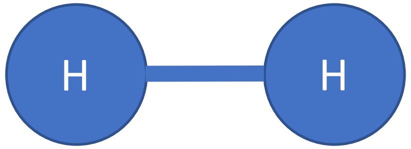

Charge	x	y	z
1	0	0	0
1	0.74	0	0

8 features

= [1 0 0 0 1 0.74 0 0 0]

Toy (not really) problem: predict potential energy of a molecule

Naïve attempt 3

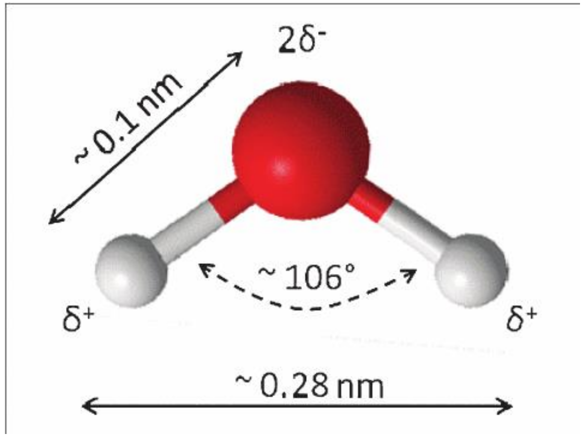

Source

Charge	x	y	z
1	1.1	-1.4	0
1	1.1	1.4	0
16	0	0	0

Augmented training with random permutation

[1 1.4 -1.4 0 1 1.1 1.4 0 16 0 0 0]

….

Zero-padding up to maximum molecule size

[1 1.4 -1.4 0 1 1.1 1.4 0 16 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0]

CatBoost

Profit?

Given enough data, this might work, but incredibly data-inefficient.

Meet graphs

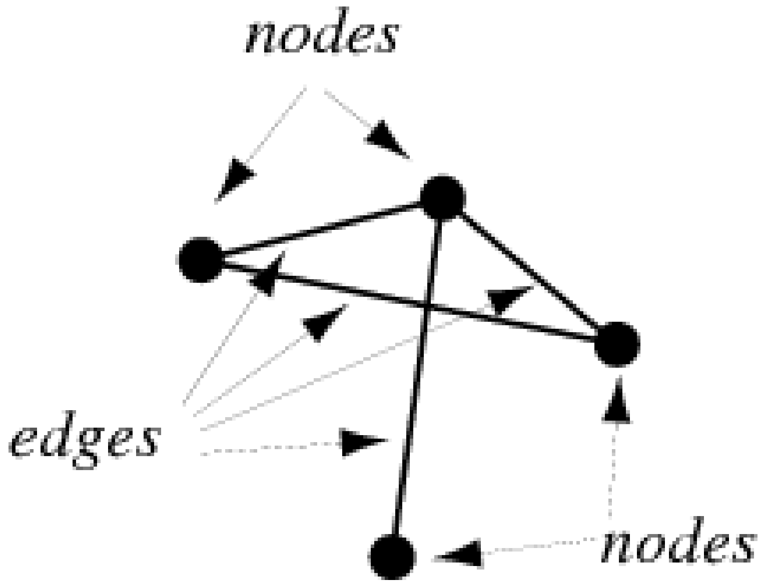

Source

Graphs, graphs everywhere!

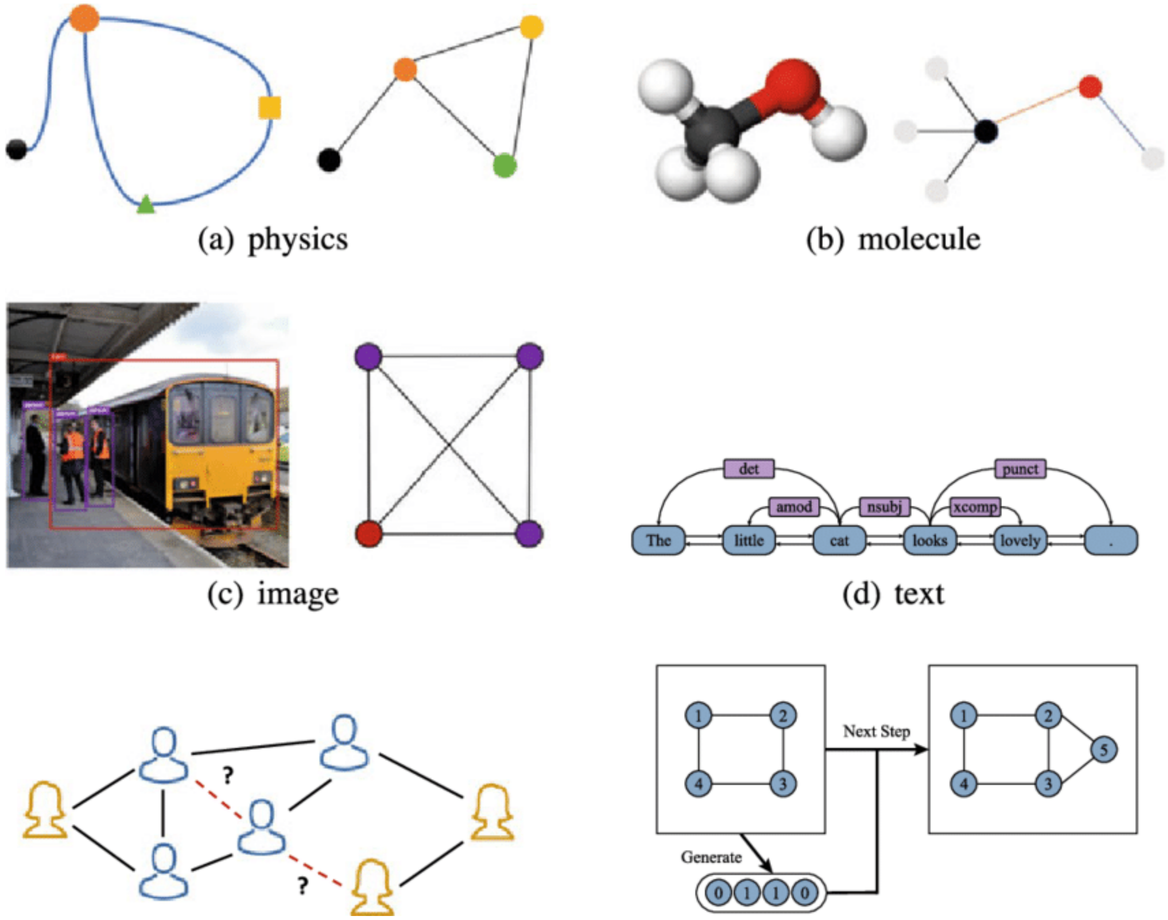

[Source](#)

Graph as a representation

Easy to construct, but hard to machine learn

In machine learning inductive bias is everything

Graphs provide:

- Permutation invariance
- Different system size
- Locality of interactions

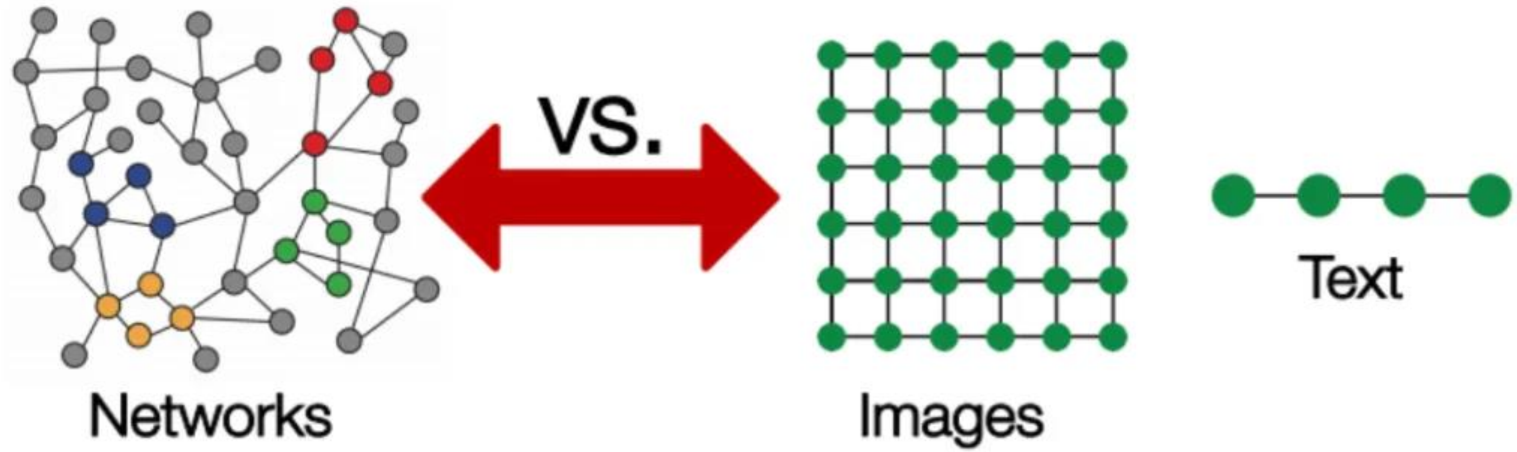

Source

Outline of a GNN

GraphNN layer at a glance

- Input: graph, each edge and node has a state vector; global state vector
- Output: graph; global state vector
- Doesn't change connectivity
- Steps:
  - Compute new edge states
  - Compute new node states
  - Compute new global state

Definitions

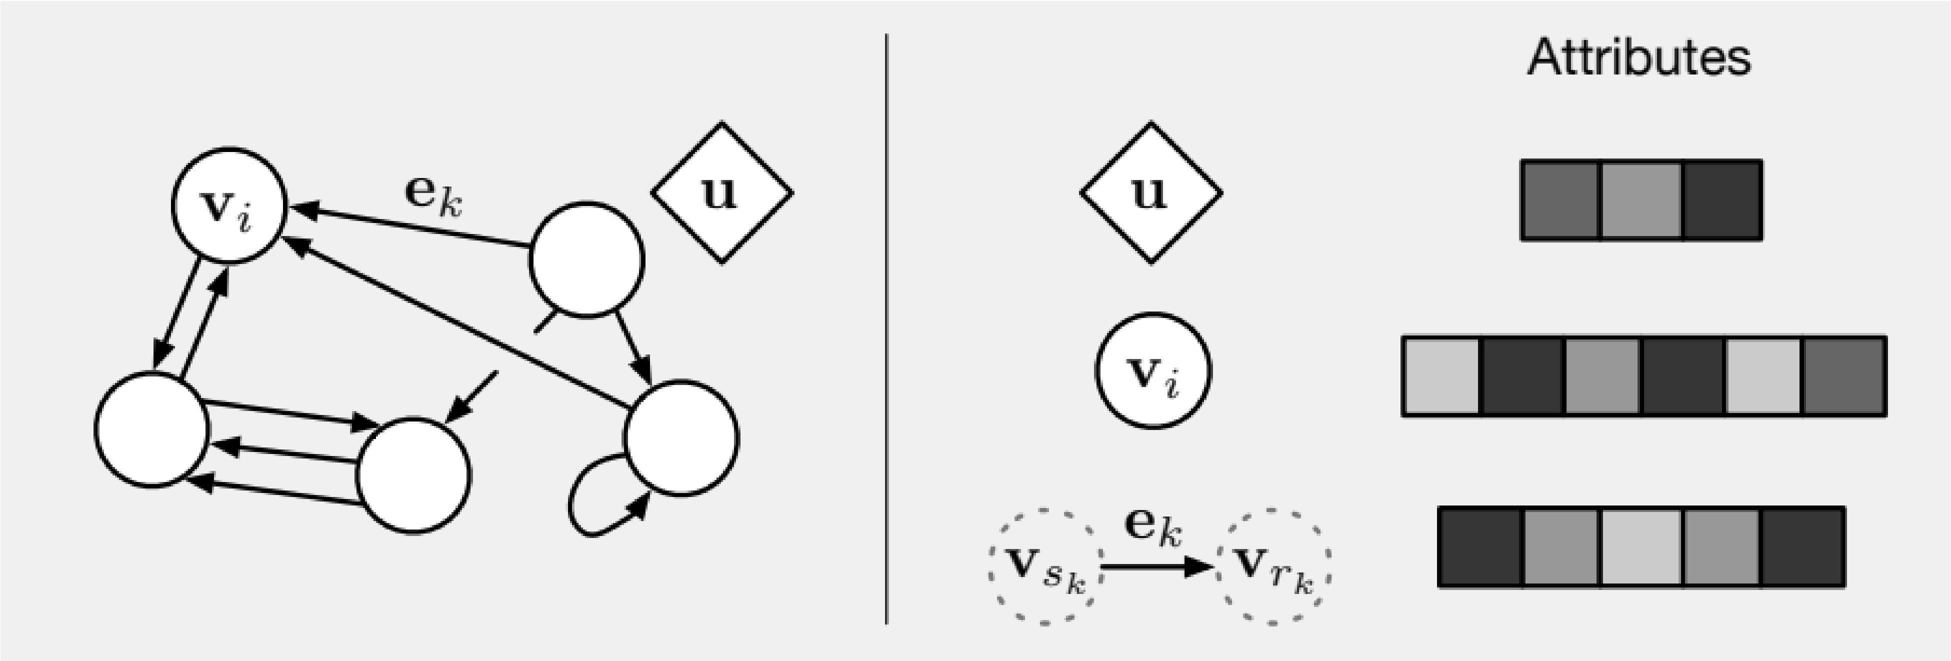

$\boldsymbol{v}_i$ – state of node i, vector
$\boldsymbol{e}_k$ – state of edge k, vector
$\boldsymbol{u}$ – global state, vector

[Source]

Edge update

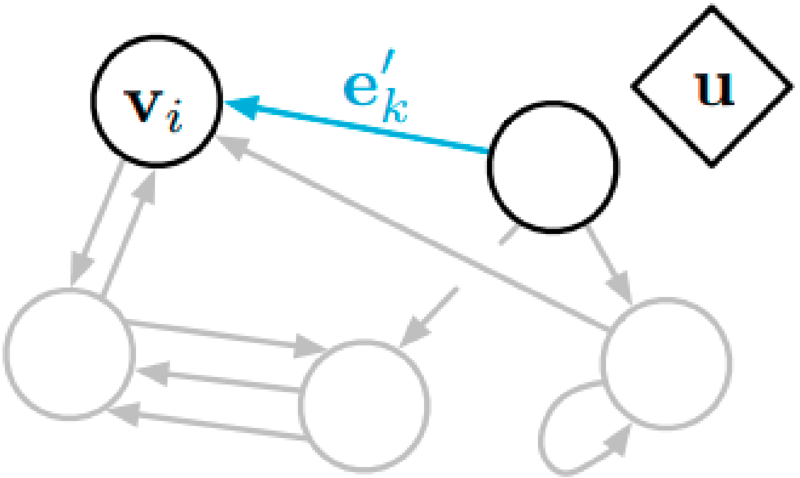

new edge state = **NN**(old edge state, incident vertices’ states, global state)

Node update

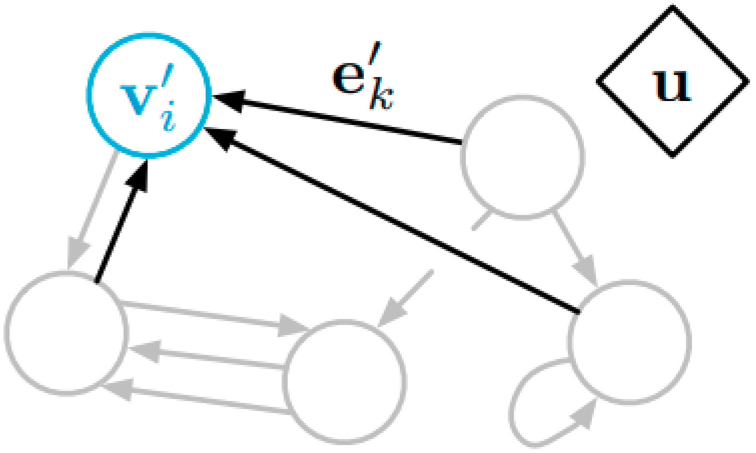

new node state = **NN**(old node state, **sum**(incoming edges' states), global state)

[Source]

Global state update

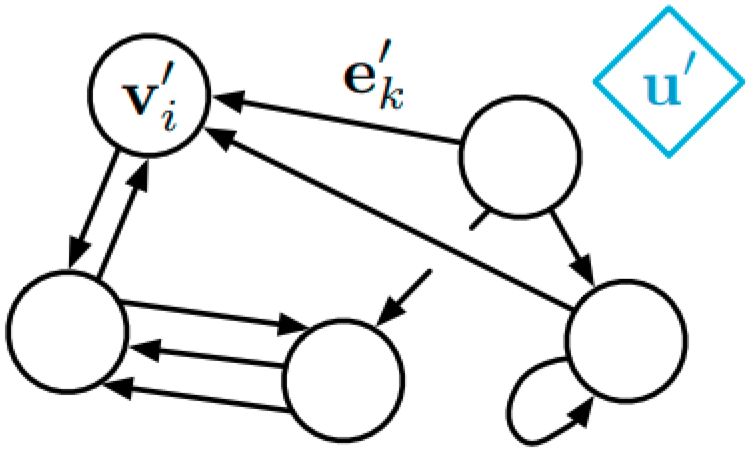

new global state = **NN**(**sum**(vertices' states), **sum**(edges' states), old global state)

[Source](#)

Deep GraphNN

- Input has the same structure as output, stack layers to make it deep
- Can mark any of the states as the output
- Trainable via back-propagation
- Suitable for predicting global, node and edge targets
- By no means the only one possible, read the more in the references

Graph neural networks in physics

HEP

- Graph-level target
  - Jet classification (LHC)
  - Whole event classification (IceCube)
- Node-level target
  - Pileup (noise) identification
  - Calorimeter reconstruction
- Edge-level target
  - Charged-particle tracking
  - Secondary vertex reconstruction

Learn more

- A very good textbook-like description
- Battaglia, Peter W., et al. "Relational inductive biases, deep learning, and graph networks." *arXiv preprint arXiv:1806.01261* (2018). - a through and math-heavy derivation and description
- Zhou, Jie, et al. "Graph neural networks: A review of methods and applications." *AI Open* 1 (2020): 57-81.
- Shlomi, Jonathan, Peter Battaglia, and Jean-Roch Vlimant. "Graph neural networks in particle physics." *Machine Learning: Science and Technology* 2.2 (2020): 021001.

Shameless advertisement

- Laboratory of methods for Big Data Analysis at HSE University
    - Applications of Machine Learning to **natural science challenges** at CERN and beyond
    - HSE has joined LHCb in 2018!
- Co-organizer of Flavours of Physics @Kaggle (2015), TrackML challenge (2018)
- Education activities (ML at ICL, ClermonFerrand, URL Barcelona, Coursera)
    - Summer school on Machine Learning in Hamburg, 2019, Oxford 2018, Reading 2017, Lund 2016, …
- **We have a lot of opportunities to do ML for physics, join us!**

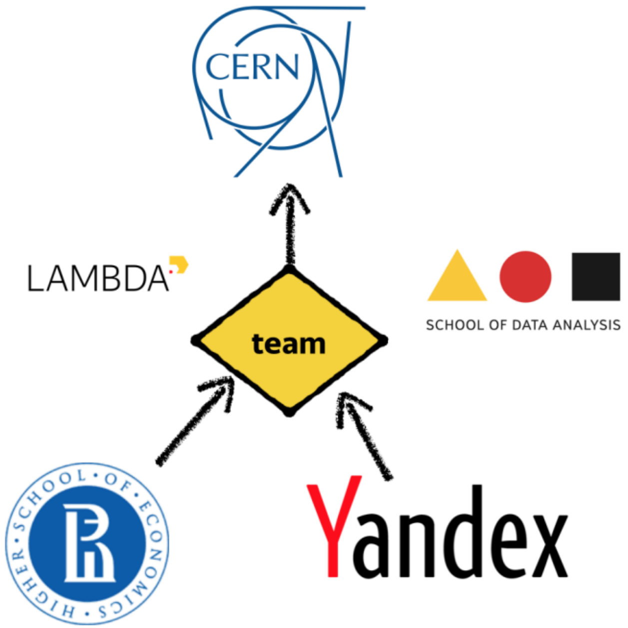

Thanks, and looking forward to hearing from you!

nikita.kazeev@cern.ch