



Microsoft  
**Research**  
微软亚洲研究院



# Jet tagging algorithm respecting Lorentz group symmetry

based on: S.Gong et al. *JHEP* 07 (2022) 030 ([arXiv:2201.08187](https://arxiv.org/abs/2201.08187));  
C.Li et al. [arXiv:2208.07814](https://arxiv.org/abs/2208.07814)

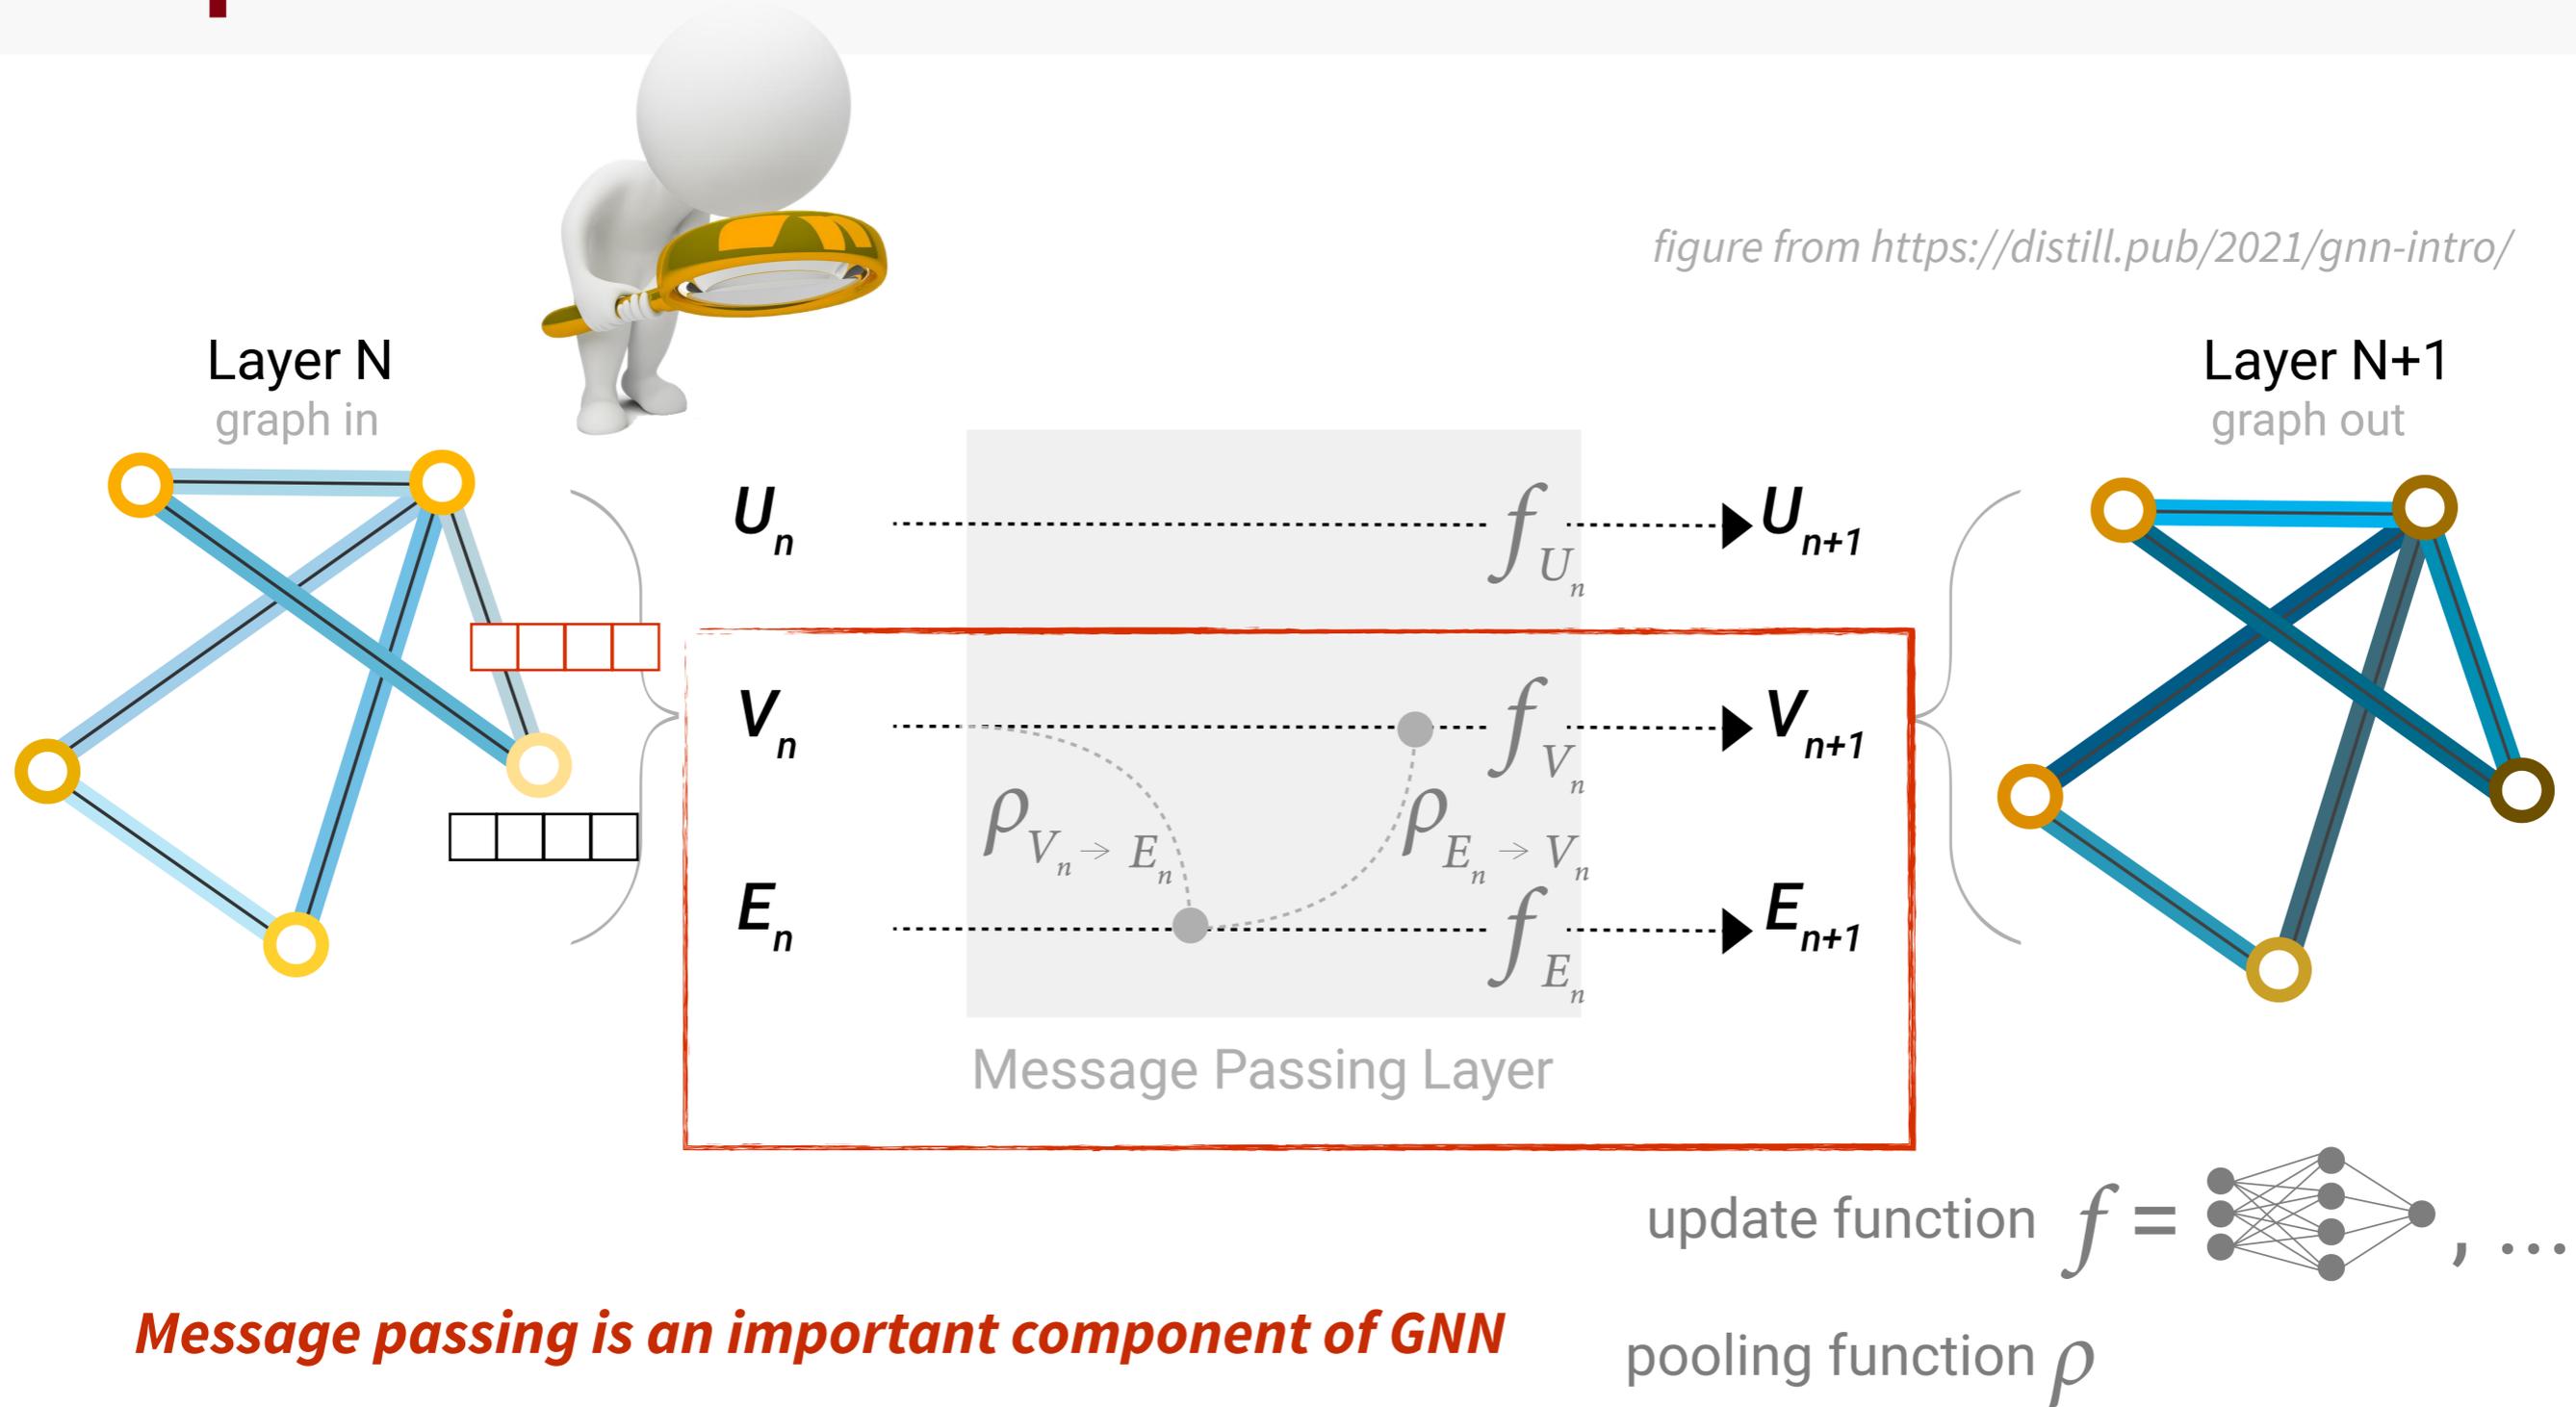
**Congqiao Li** (*Peking University*)

BOOST 2022 · Hamburg  
17 August, 2022

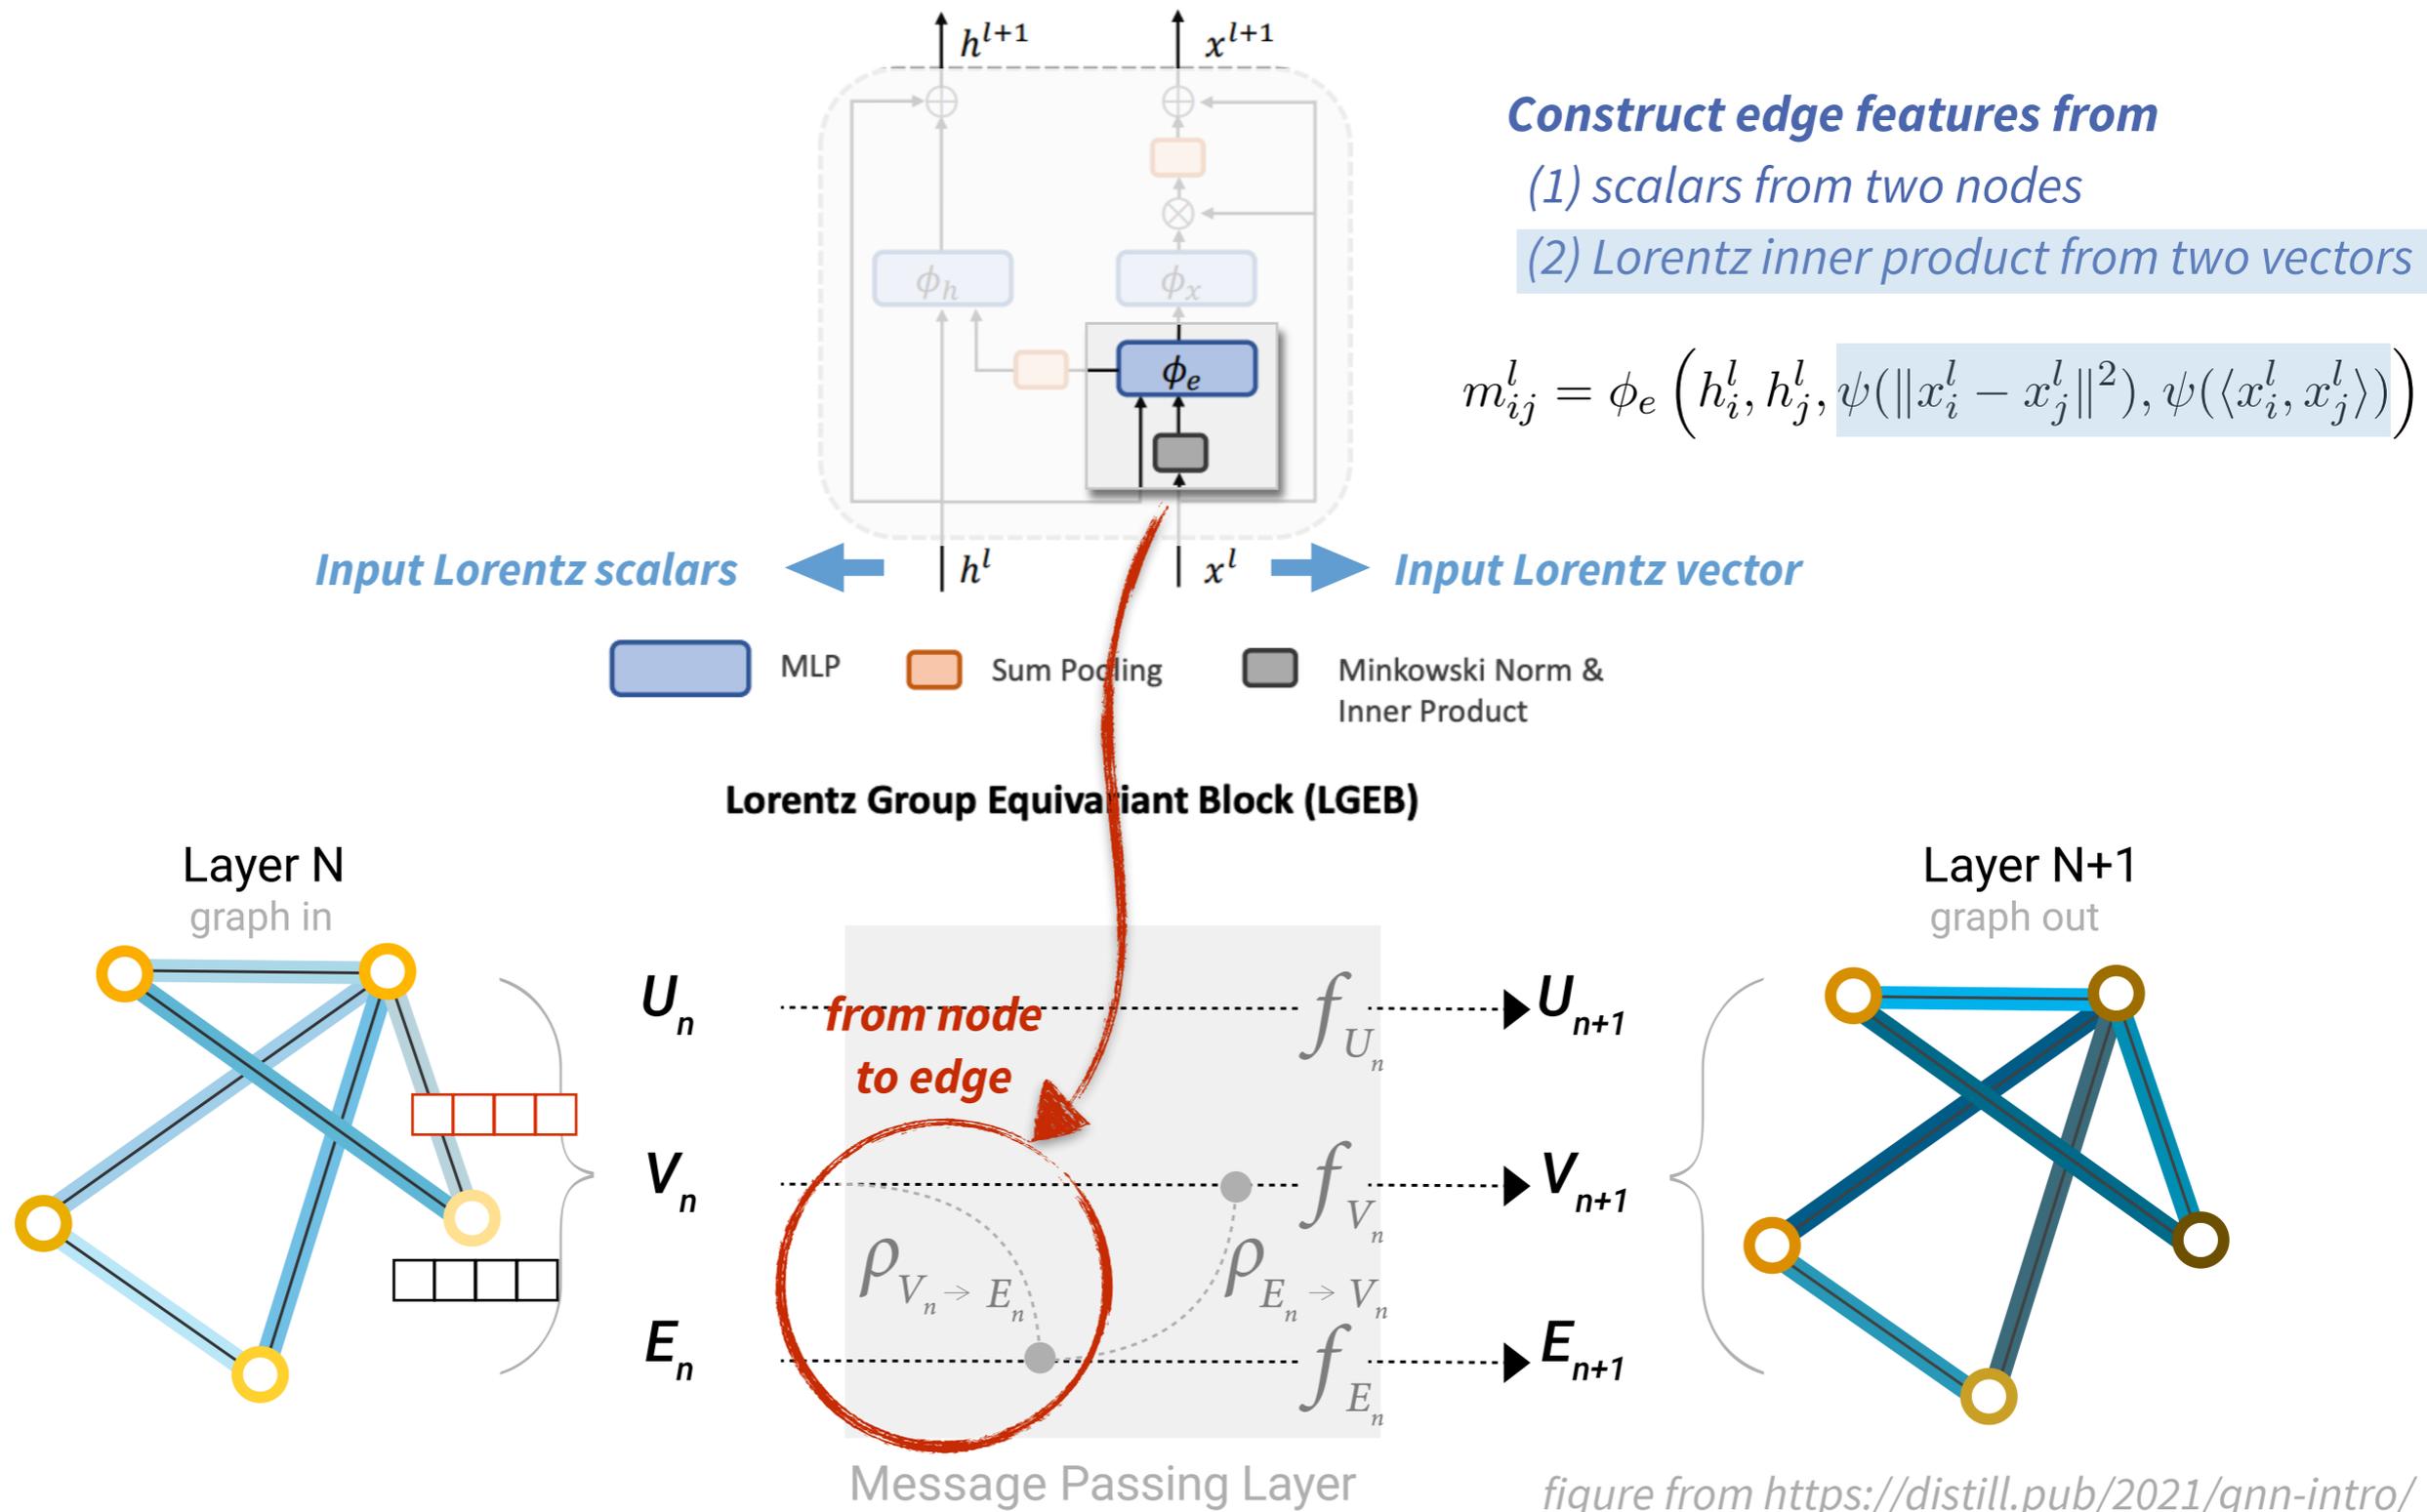
# Introduction

- ***HEP community is eager to study novel network design***
  - ❖ may directly lead to more accurate physics measurements at the LHC
- ***Full Lorentz symmetry has been studied for quite some time***
  - ❖ efforts including LoLa, Lorentz Boost Network, Lorentz Group Network, etc.
  - ❖ still cannot demonstrate the potential advantages of such a design
    - there lacks a fair comparison with the network without symmetric design
  - ❖ can we really see improvements with the use of Lorentz symmetry?
- ***LorentzNet makes a successful attempt***
  - ❖ made up of GNN + Lorentz symmetry preservation
  - ❖ then, are we able to understand where stands the underlying reason for benefits from LorentzNet?
- Let's start our journey in LorentzNet, then try to answer a broader question:  
which role does Lorentz-symmetry play in jet tagging

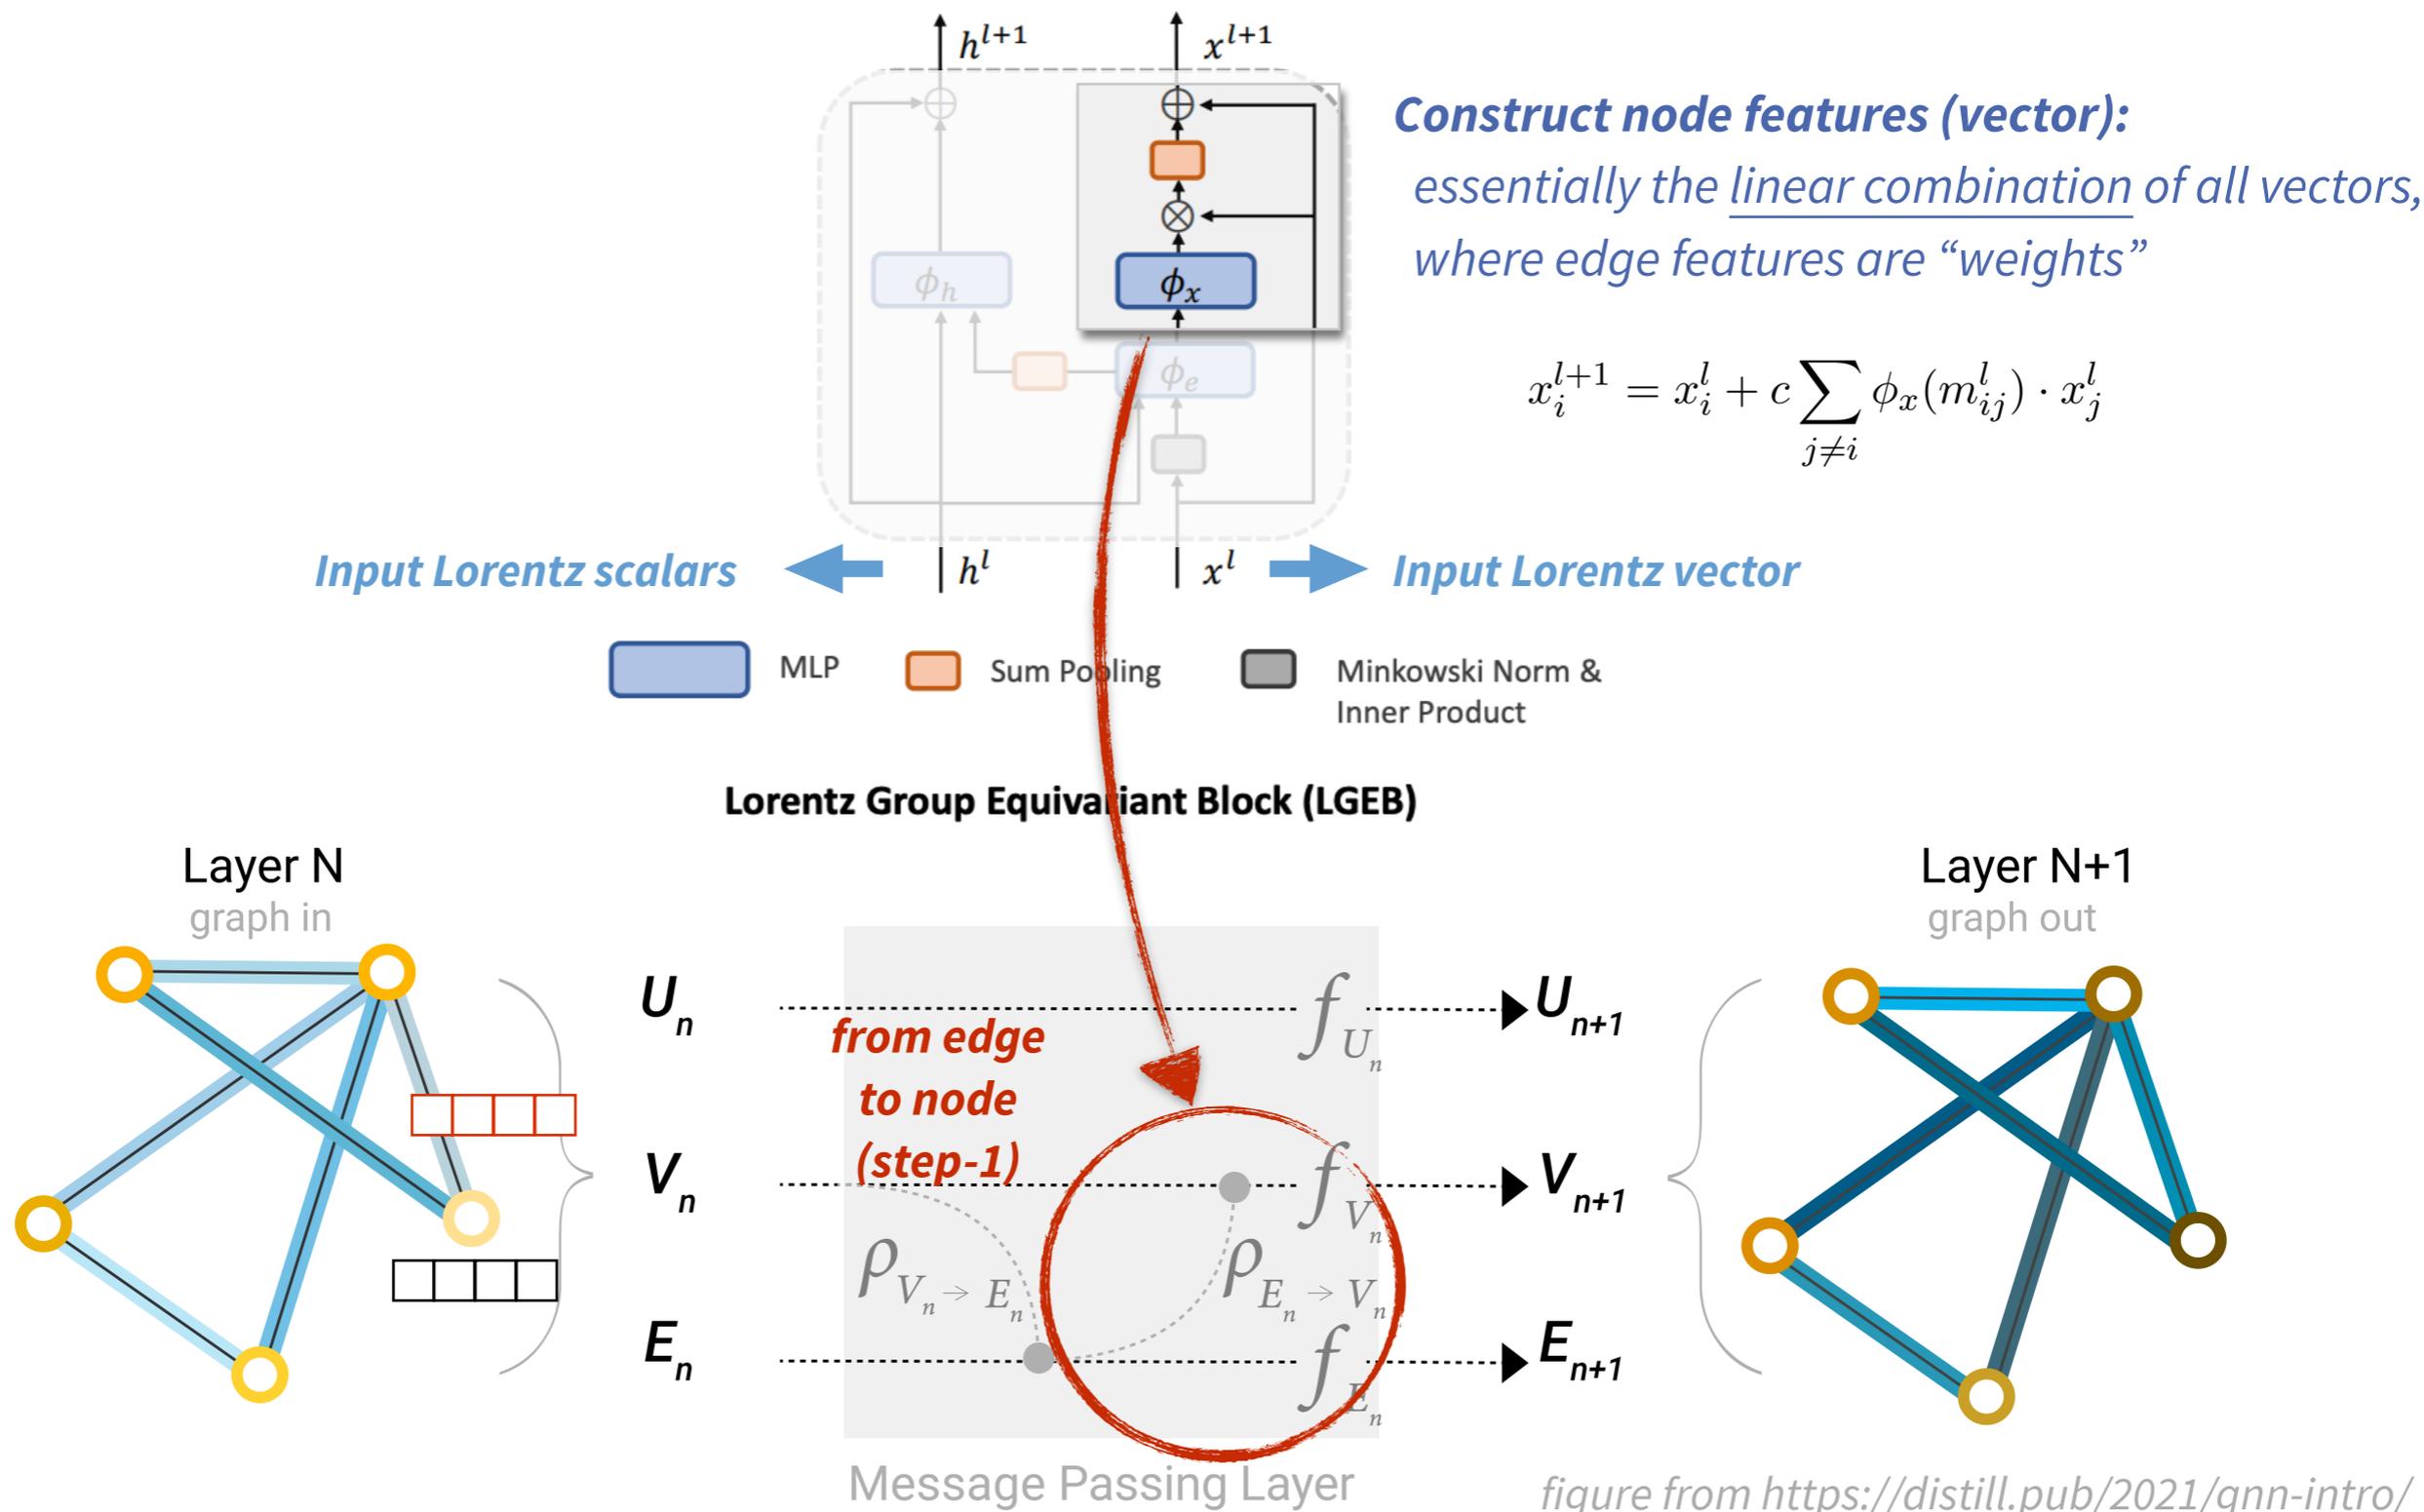
# Graph neural networks



# LorentzNet architecture



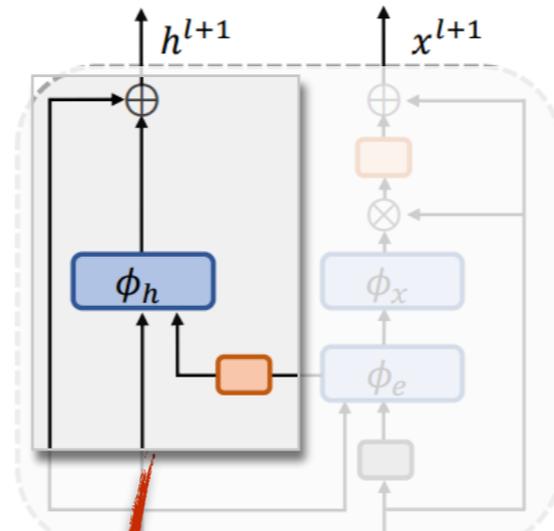
# LorentzNet architecture



# LorentzNet architecture

**Construct node features (scalar):**  
attentive pooling on all connecting edges

$$h_i^{l+1} = h_i^l + \phi_h(h_i^l, \sum_{j \neq i} w_{ij} m_{ij}^l)$$



Input Lorentz scalars

$h^l$

$x^l$

Input Lorentz vector



MLP

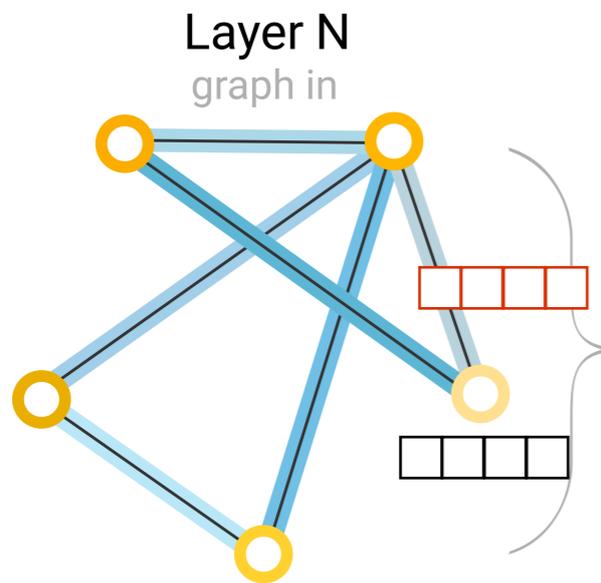


Sum Pooling



Minkowski Norm & Inner Product

## Lorentz Group Equivariant Block (LGEB)

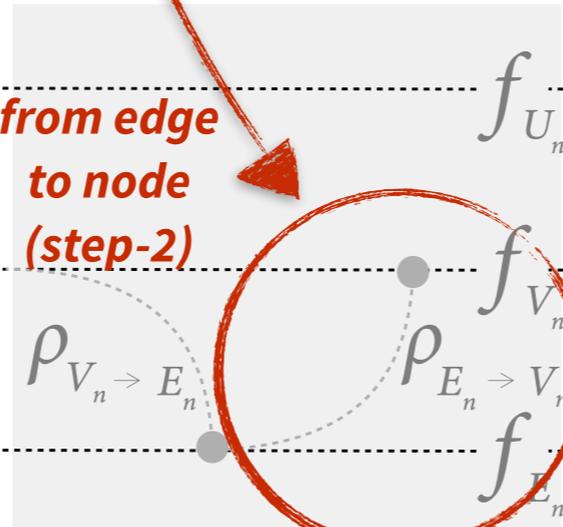


$U_n$

$V_n$

$E_n$

from edge  
to node  
(step-2)

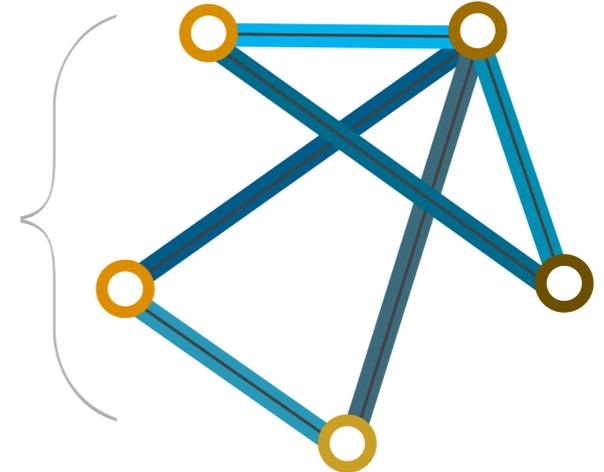


$U_{n+1}$

$V_{n+1}$

$E_{n+1}$

Layer N+1  
graph out



Message Passing Layer

figure from <https://distill.pub/2021/gnn-intro/>

# Summary of architecture

→ Now let's summarize the main architecture of LorentzNet

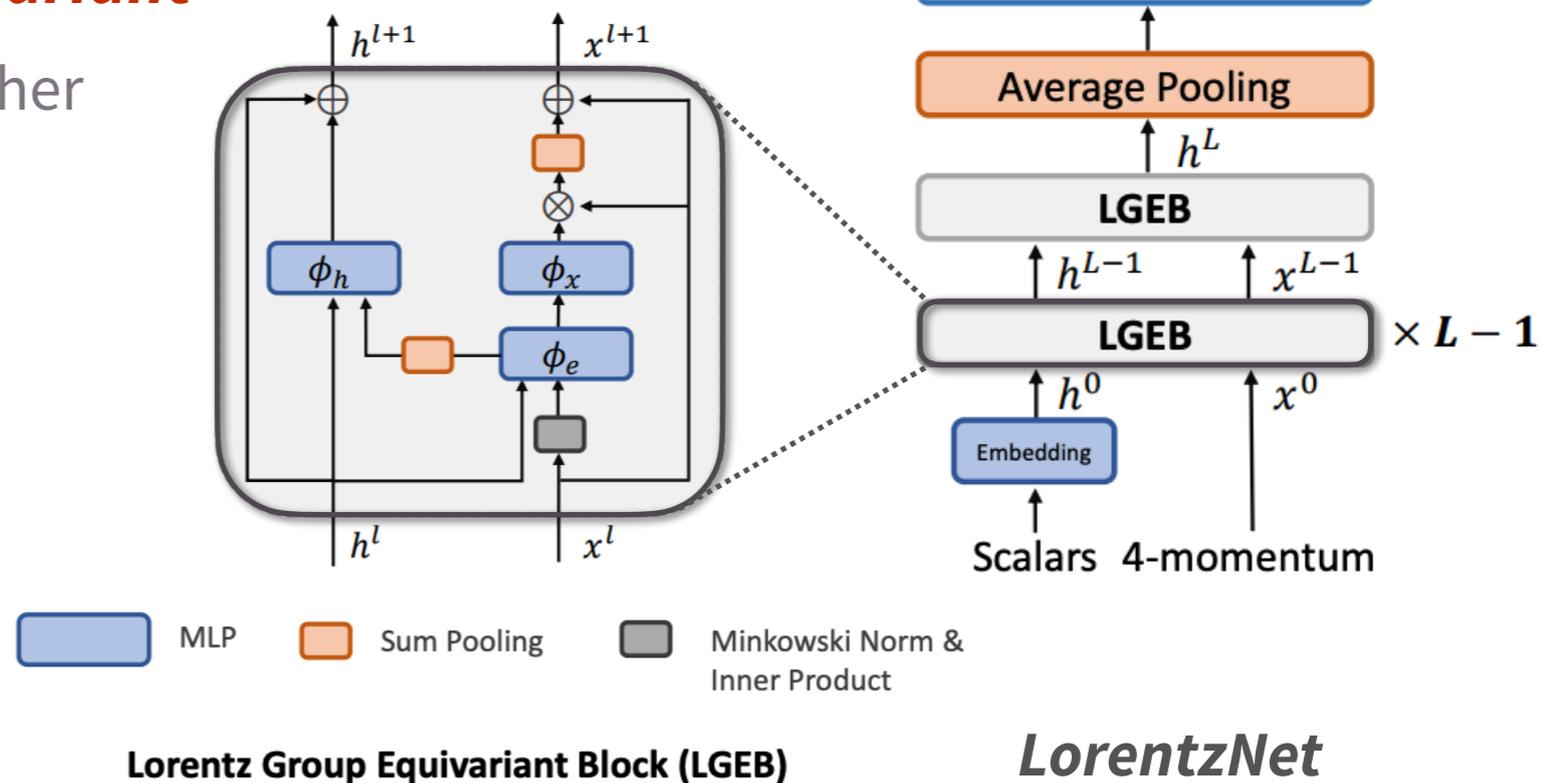
- ❖ **Graph neural network** as backbone

- ❖ **Fully connected**

- ▶ i.e., all  $N(N - 1)/2$  edges are computed
- ▶ ParticleNet use dynamic  $k$ -nearest neighbours to define edges (DGCNN), so it is not using the full pairs

- ❖ **Fully Lorentz invariant/equivariant**

- ▶ nodes can be grouped by either Lorentz scalars or vectors



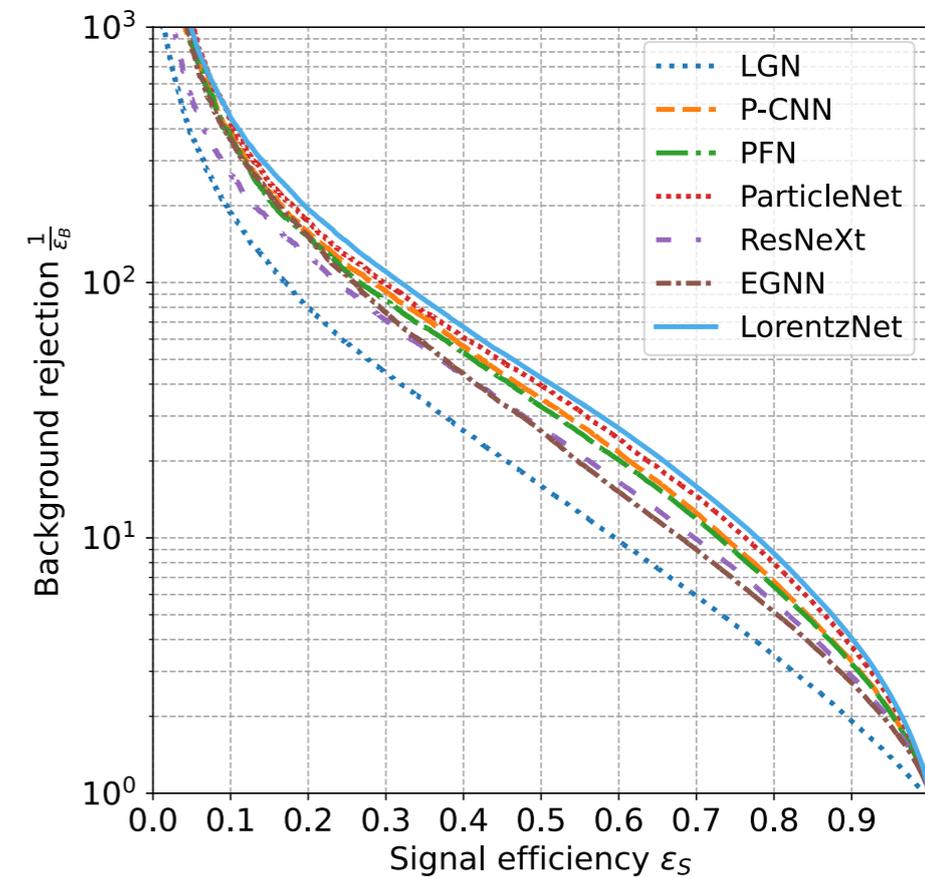
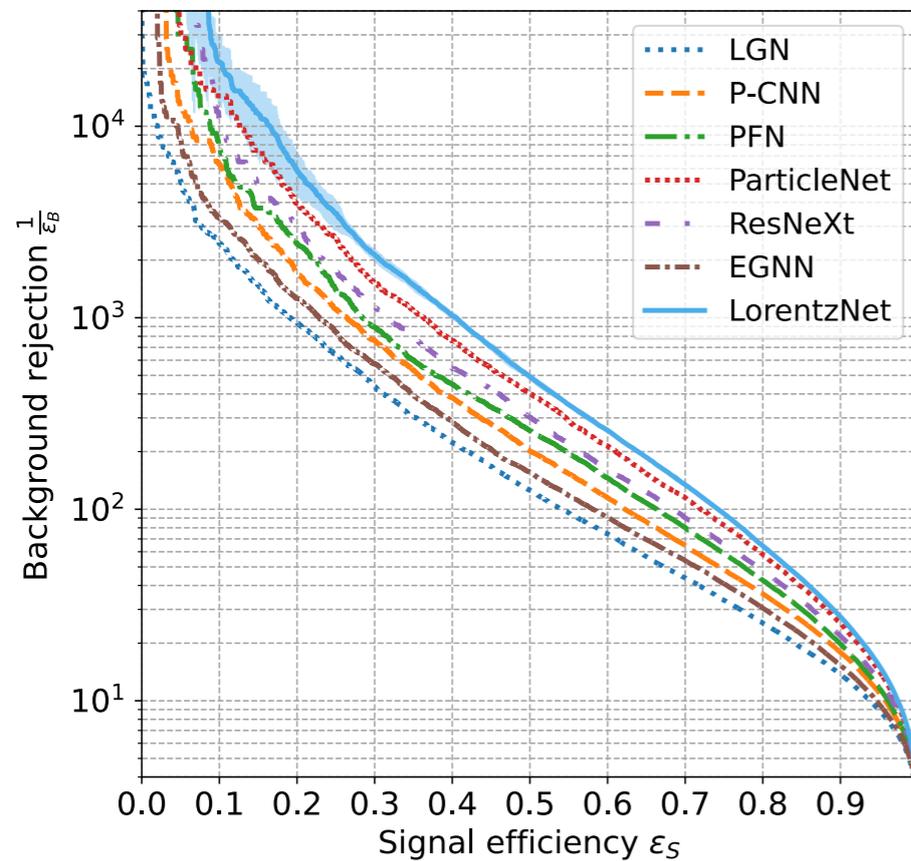
# Performance

*Top tagging benchmark [SciPost Phys. 7 (2019) 014]*

Model	Accuracy	AUC	$1/\varepsilon_B$ ( $\varepsilon_S = 0.5$ )	$1/\varepsilon_B$ ( $\varepsilon_S = 0.3$ )
ResNeXt	0.936	0.9837	$302 \pm 5$	$1147 \pm 58$
P-CNN	0.930	0.9803	$201 \pm 4$	$759 \pm 24$
PFN	0.932	0.9819	$247 \pm 3$	$888 \pm 17$
ParticleNet	0.940	0.9858	$397 \pm 7$	$1615 \pm 93$
EGNN	0.922	0.9760	$148 \pm 8$	$540 \pm 49$
LGN	0.929	0.9640	$124 \pm 20$	$435 \pm 95$
LorentzNet	<b>0.942</b>	<b>0.9868</b>	<b><math>498 \pm 18</math></b>	<b><math>2195 \pm 173</math></b>

*Quark-gluon tagging benchmark [JHEP 01 (2019) 121]*

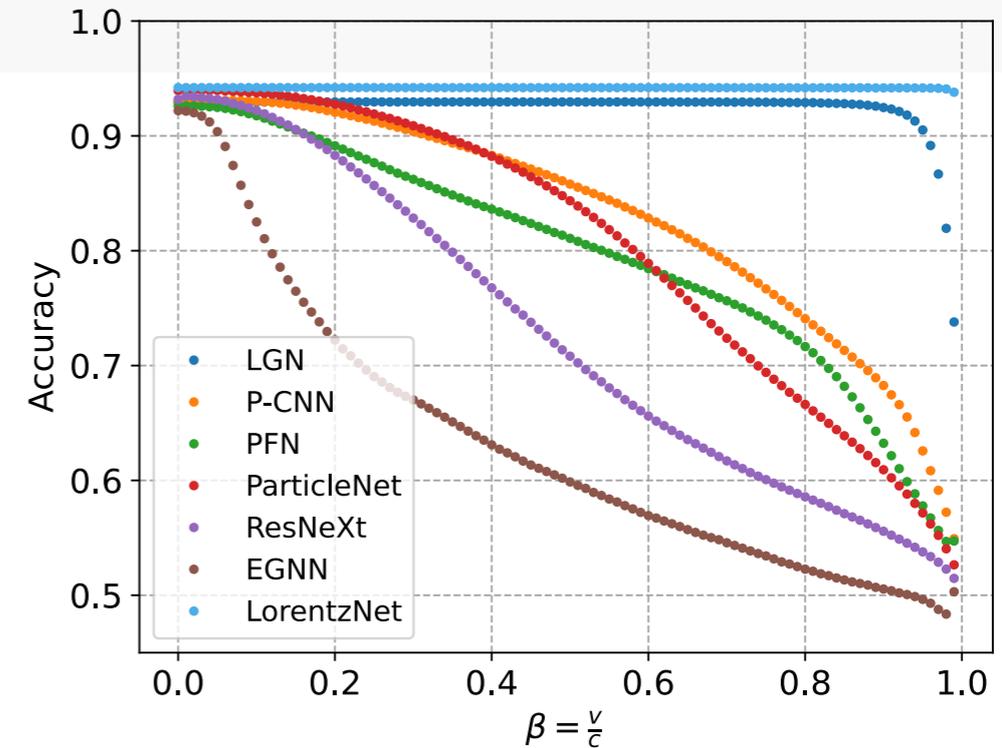
Model	Accuracy	AUC	$1/\varepsilon_B$ ( $\varepsilon_S = 0.5$ )	$1/\varepsilon_B$ ( $\varepsilon_S = 0.3$ )
ResNeXt	0.821	0.8960	30.9	80.8
P-CNN	0.827	0.9002	34.7	91.0
PFN	—	0.9005	$34.7 \pm 0.4$	—
ParticleNet	0.840	0.9116	$39.8 \pm 0.2$	$98.6 \pm 1.3$
EGNN	0.803	0.8806	$26.3 \pm 0.3$	$76.6 \pm 0.5$
LGN	0.803	0.8141	8.30	15.2
LorentzNet	<b>0.844</b>	<b>0.9156</b>	<b><math>42.4 \pm 0.4</math></b>	<b><math>110.2 \pm 1.3</math></b>



# Additional tests

## → *Equivariance test:*

- ❖ LorentzNet is more robust when the input jet undergoes a Lorentz transformation (consider Lorentz boosts on  $x$ -axis)



## → *Small training sample size:*

- ❖ LorentzNet is able to perform much better when trained on a smaller size of sample

Training Fraction	Model	Accuracy	AUC	$1/\varepsilon_B$ ( $\varepsilon_S = 0.5$ )	$1/\varepsilon_B$ ( $\varepsilon_S = 0.3$ )
0.5%	ParticleNet	0.913	0.9687	$77 \pm 4$	$199 \pm 14$
	LorentzNet	<b>0.929</b>	<b>0.9793</b>	<b><math>176 \pm 14</math></b>	<b><math>562 \pm 72</math></b>
1%	ParticleNet	0.919	0.9734	$103 \pm 5$	$287 \pm 19$
	LorentzNet	<b>0.932</b>	<b>0.9812</b>	<b><math>209 \pm 5</math></b>	<b><math>697 \pm 58</math></b>
5%	ParticleNet	0.931	0.9807	$195 \pm 4$	$609 \pm 35$
	LorentzNet	<b>0.937</b>	<b>0.9839</b>	<b><math>293 \pm 12</math></b>	<b><math>1108 \pm 84</math></b>

## → *Ablation study on Lorentz equivariant preserving structure*

- ❖ replacing the pairwise scalar (mass) has a negative effect on the network

Model	Equivariance	Accuracy	AUC	$1/\varepsilon_B$ ( $\varepsilon_S = 0.5$ )	$1/\varepsilon_B$ ( $\varepsilon_S = 0.3$ )
LorentzNet (w/o)	✗	0.934	0.9832	$290 \pm 30$	$1105 \pm 59$
LorentzNet	✓	<b>0.942</b>	<b>0.9868</b>	<b><math>498 \pm 18</math></b>	<b><math>2195 \pm 173</math></b>

# Conclusion

S.Gong et al. *JHEP* 07 (2022) 030

- We present LorentzNet, a Lorentz group equivariant GNN
  - ❖ the network has now reached state-of-the-art performance, when trained and evaluated on two mainstream benchmarks
  - ❖ its equivariance property confirmed on Lorentz-transformed test dataset
  - ❖ ablation study shows Lorentz-symmetry-preserving mechanism does help the network

- END OF SLIDES -



# Conclusion (?)

S.Gong et al. *JHEP* 07 (2022) 030

- We present LorentzNet, a Lorentz group equivariant GNN
  - ❖ the network has now reached state-of-the-art performance, when trained and evaluated on two mainstream benchmarks
  - ❖ its equivariance property confirmed on Lorentz-transformed test dataset
  - ❖ ablation study shows Lorentz-symmetry-preserving mechanism does help the network

- END OF SLIDES -

## → Wait.. **there are still questions to ask/solve**

- ❖ many novel network designs have been proposed and evaluated on these two benchmarks, but see no significant improvement - why?
  - why LorentzNet behave surprising well on them?
- ❖ can we dig deeper to extract the key component in LorentzNet?  
can it be applied to other networks as well?

→ Let's continue our game 🎮



# LorentzNet performance on JetClass



*JetClass* [[arXiv:2202.03772](https://arxiv.org/abs/2202.03772), proceedings of 39th ICML, Vol.162]

	All classes		$H \rightarrow b\bar{b}$	$H \rightarrow c\bar{c}$	$H \rightarrow gg$	$H \rightarrow 4q$	$H \rightarrow \ell\nu qq'$	$t \rightarrow bqq'$	$t \rightarrow bl\nu$	$W \rightarrow qq'$	$Z \rightarrow q\bar{q}$
	Accuracy	AUC	Rej <sub>50%</sub>	Rej <sub>50%</sub>	Rej <sub>50%</sub>	Rej <sub>50%</sub>	Rej <sub>99%</sub>	Rej <sub>50%</sub>	Rej <sub>99.5%</sub>	Rej <sub>50%</sub>	Rej <sub>50%</sub>
PFN	0.772	0.9714	2924	841	75	198	265	797	721	189	159
P-CNN	0.809	0.9789	4890	1276	88	474	947	2907	2304	241	204
ParticleNet	0.844	0.9849	7634	2475	104	954	3339	10526	11173	347	283
 LorentzNet	0.855	0.9869	9217	3425	117	1550	4425	19802	12500	480	353
<b>ParT</b>	<b>0.861</b>	<b>0.9877</b>	<b>10638</b>	<b>4149</b>	<b>123</b>	<b>1864</b>	<b>5479</b>	<b>32787</b>	<b>15873</b>	<b>543</b>	<b>402</b>
ParT (plain)	0.849	0.9859	9569	2911	112	1185	3868	17699	12987	384	311

→ LorentzNet performs much better than ParticleNet, slightly worse than ParT

- ❖ note that for #params LorentzNet < ParticleNet, and << ParT
- ❖ can prove that **LorentzNet is still more performant**

→ We may want to understand why there is performance difference between *top tagging dataset* and *JetClass*

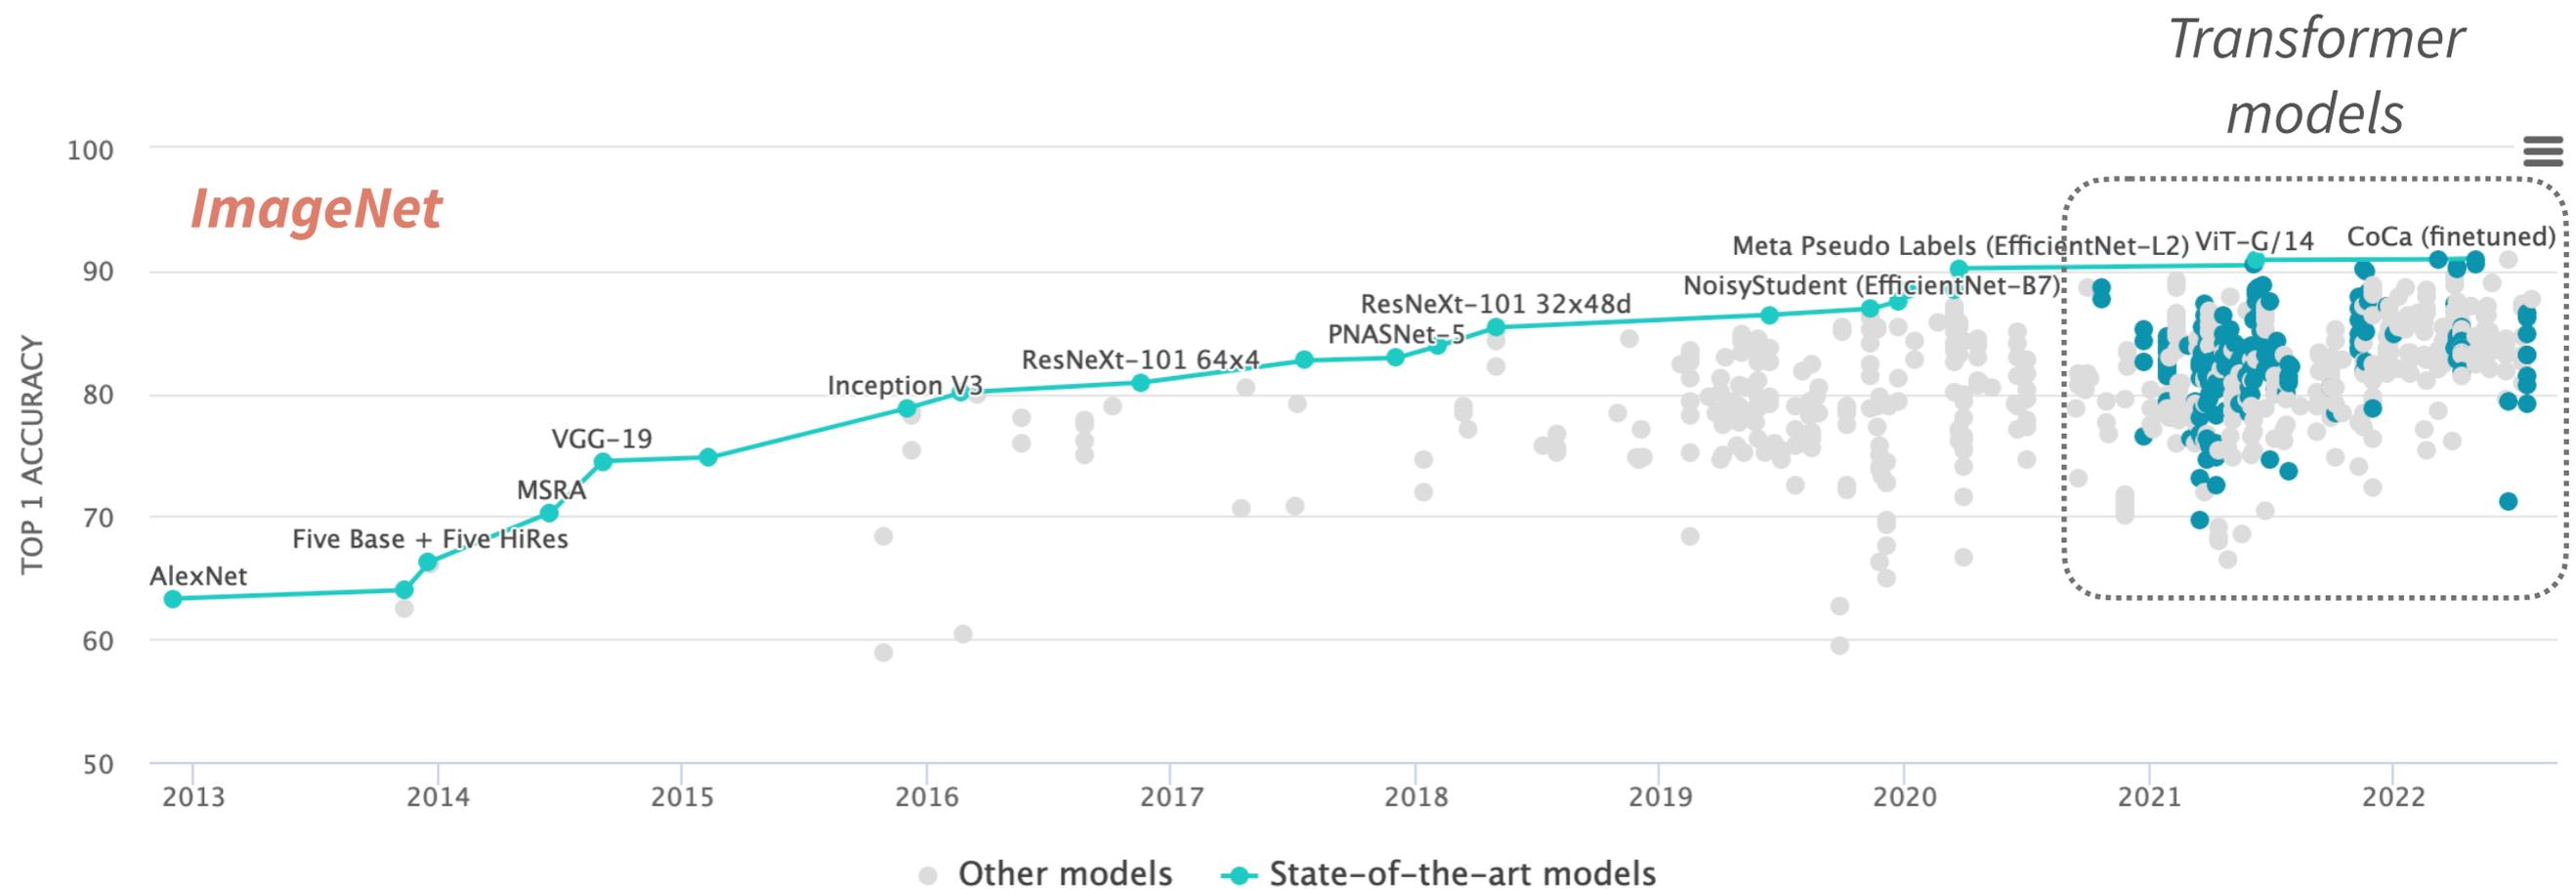


	Accuracy	# params	FLOPs
PFN	0.772	86.1 k	4.62 M
P-CNN	0.809	354 k	15.5 M
ParticleNet	0.844	370 k	540 M
LorentzNet	0.855	233 k	2.01 G
<b>ParT</b>	<b>0.861</b>	2.14 M	340 M
ParT (plain)	0.849	2.13 M	260 M

# Model, data size and “inductive bias”

## → Lessons from image classification from Computer Vision

- ❖ Training on ImageNet and its extension (224x224 pixel image classification)
  - Transformer models have led the performance, since the first application in 2020

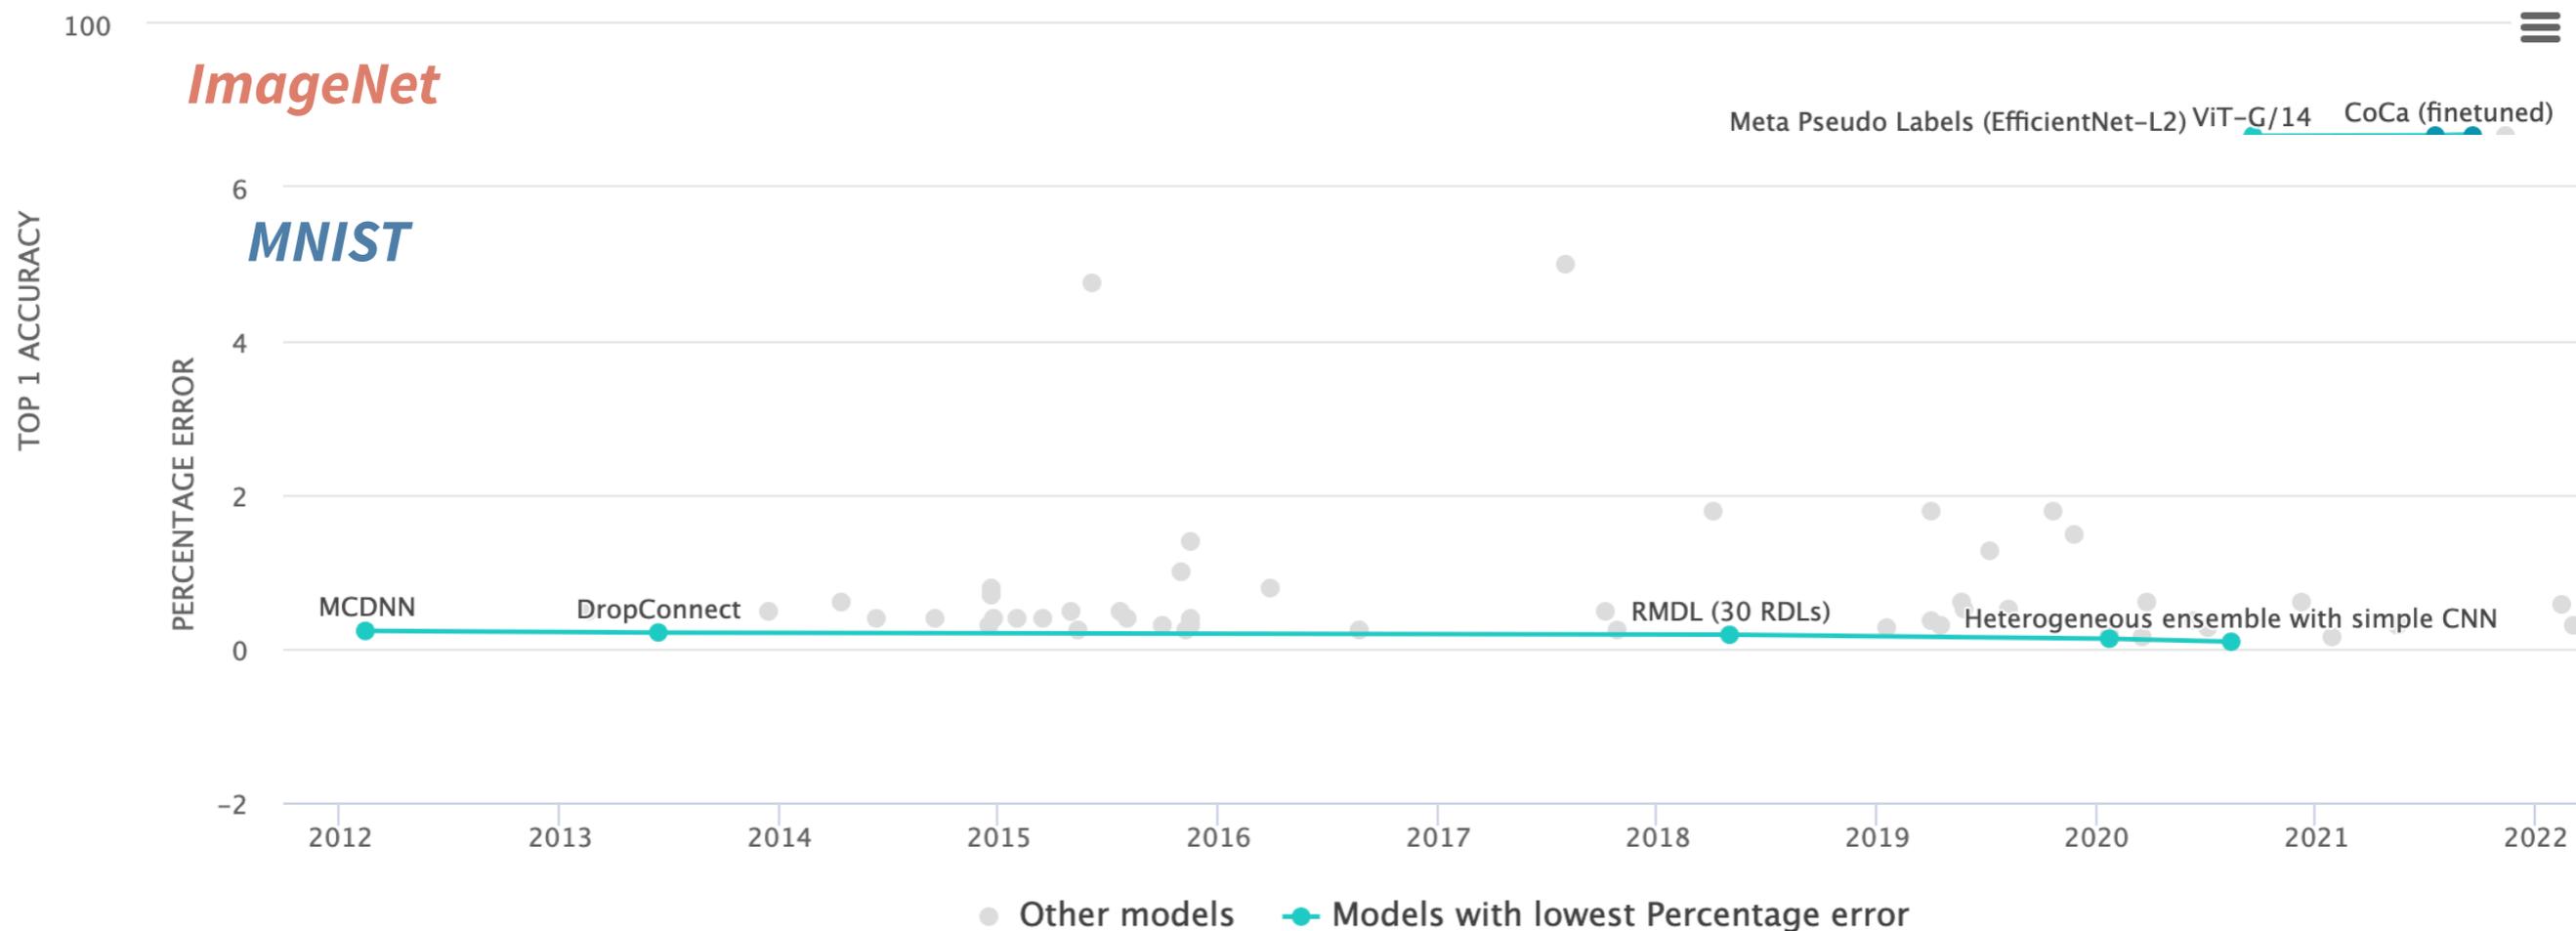


<https://paperswithcode.com/sota/image-classification-on-imagenet>

# Model, data size and “inductive bias”

## → Lessons from image classification from Computer Vision

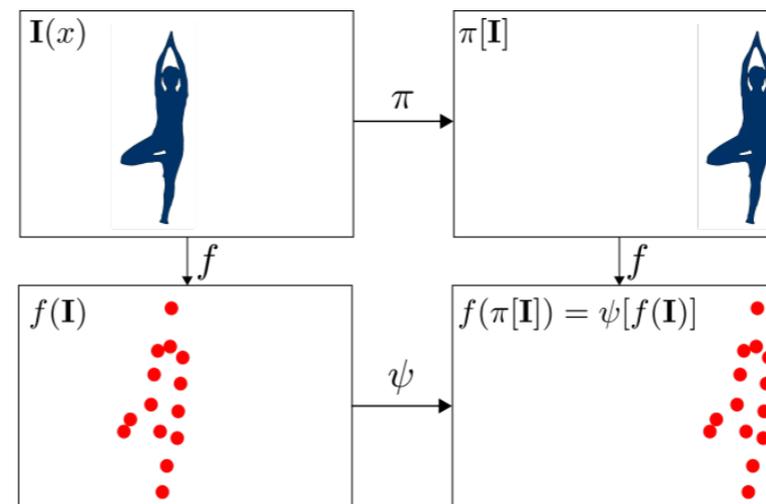
- ❖ Training on ImageNet and its extension (224x224 pixel image classification)
  - Transformer models have led the performance, since the first application in 2020
- ❖ But if we look back to MNIST dataset (hand-written digit classification)
  - still CNN-based networks rank higher



# Model, data size and “inductive bias”

## → Lessons from image classification from Computer Vision

- ❖ Training on ImageNet and its extension (224x224 pixel image classification)
  - ▶ Transformer models have led the performance, since the first application in 2020
- ❖ But if we look back to MNIST dataset (hand-written digit classification)
  - ▶ still CNN-based networks rank higher
- ❖ **Possible explanations:**
  - ▶ for MNIST dataset, we want more “efficient” model when training on small dataset
  - ▶ to be more efficient, cooperating with “**inductive bias**” in the network design is crucial
  - ▶ CNN respects the local translational symmetry, which is an inductive bias when processing real-world images



# Model, data size and “inductive bias”

## → Lessons from image classification from Computer Vision

- ❖ Training on ImageNet and its extension (224x224 pixel image classification)
  - Transformer models have led the performance, since the first application in 2020
- ❖ But if we look back to MNIST dataset (hand-written digit classification)
  - still CNN-based networks rank higher
- ❖ **Possible explanations:**
  - for MNIST dataset, we want more “efficient” model when training on small dataset
  - to be more efficient, cooperating with “**inductive bias**” in the network design is
  - We think it is the same situation for the top tagging dataset
    - the discrimination between top vs. QCD jets is an easier task, and performance are more saturated
    - it is the “**more efficient model**” that tends to win the top tagging competition
  - Now we need to verify if LorentzNet is using some inductive bias to be efficient



# Interpret Lorentz-symmetry as an inductive bias

## → Goal:

- ❖ we want to verify that Lorentz-symmetry preservation is a key component
- ❖ even better if we isolate “a patch” from LorentzNet, which can be applied to wider range of networks

## → Our experiments

- ❖ devise multiple choices of pairwise features, which are invariant to some or all Lorentz transformations
- ❖ want to see if this affects network performance as we expect

C.Li *et al.* [arXiv:2208.07814](https://arxiv.org/abs/2208.07814)

(following slides present a part of this work)

# Interpret Lorentz-symmetry as an inductive bias

## → Goal:

- ❖ we want to verify that Lorentz-symmetry preservation is a key component
- ❖ even better if we isolate “a patch” from LorentzNet, which can be applied to wider range of networks

## → Our experiments

- ❖ devise multiple choices of pairwise features, which are **invariant to some or all Lorentz transformations**
- ❖ want to see if this affects network performance as we expect

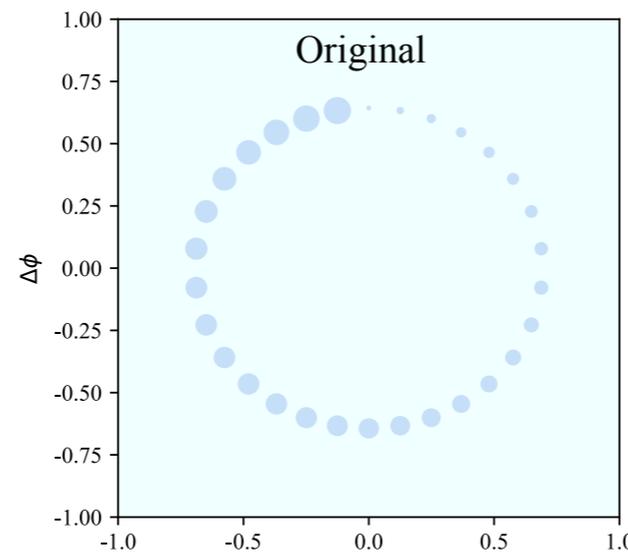
C.Li et al. [arXiv:2208.07814](https://arxiv.org/abs/2208.07814)

(following slides present a part of this work)

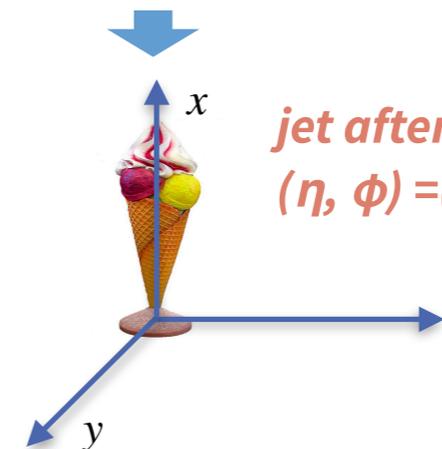
*First of all, we need to find a good way to categorize the possible Lorentz transformations acted on the jet*

# Lorentz transformations

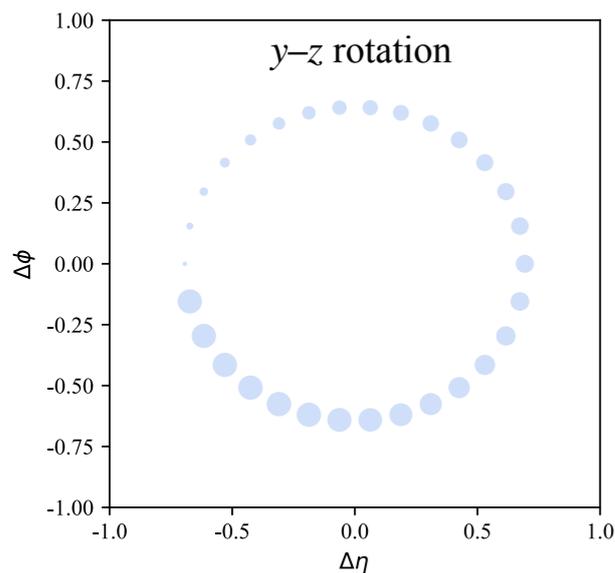
- By HEP convention, jets are represented on  $\eta$ - $\phi$  plane w.r.t. its mean axis
- ❖ equivalent as: boosts on  $z$ -axis  $\rightarrow$  rotation on  $x$ - $y$  plane (transverse plane)  $\rightarrow$  jet points to the  $x$ -axis
- ❖ then we have **four additional DoFs** for Lorentz transformation!



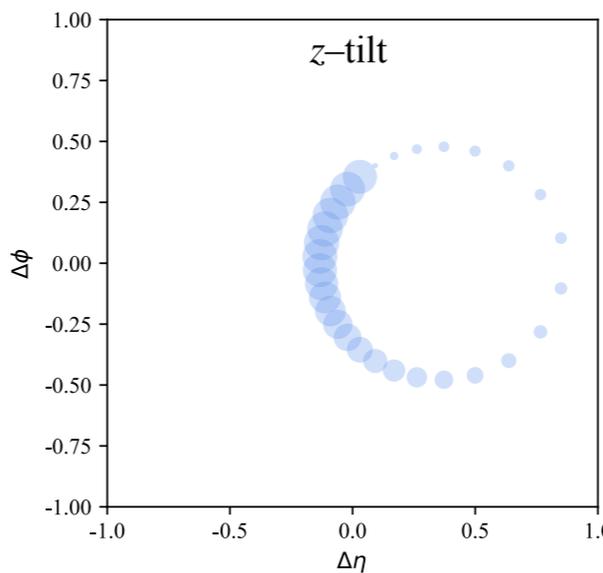
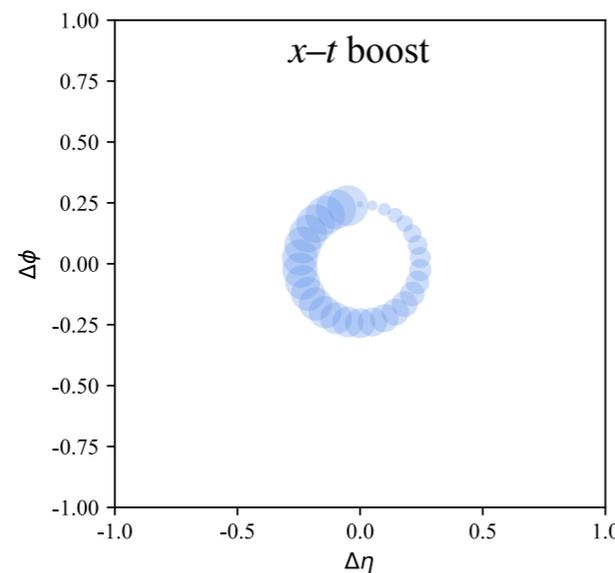
$\approx$  Viewing the jet from  $x$ -direction



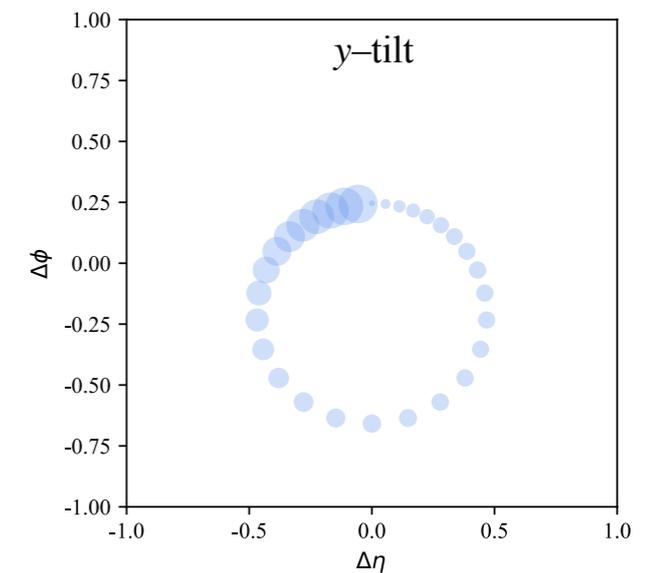
jet after preprocessing:  
 $(\eta, \phi) = (0, 0)$



$\approx \eta$ - $\phi$  rotation



$z$ -boost +  $x$ - $z$  rotation



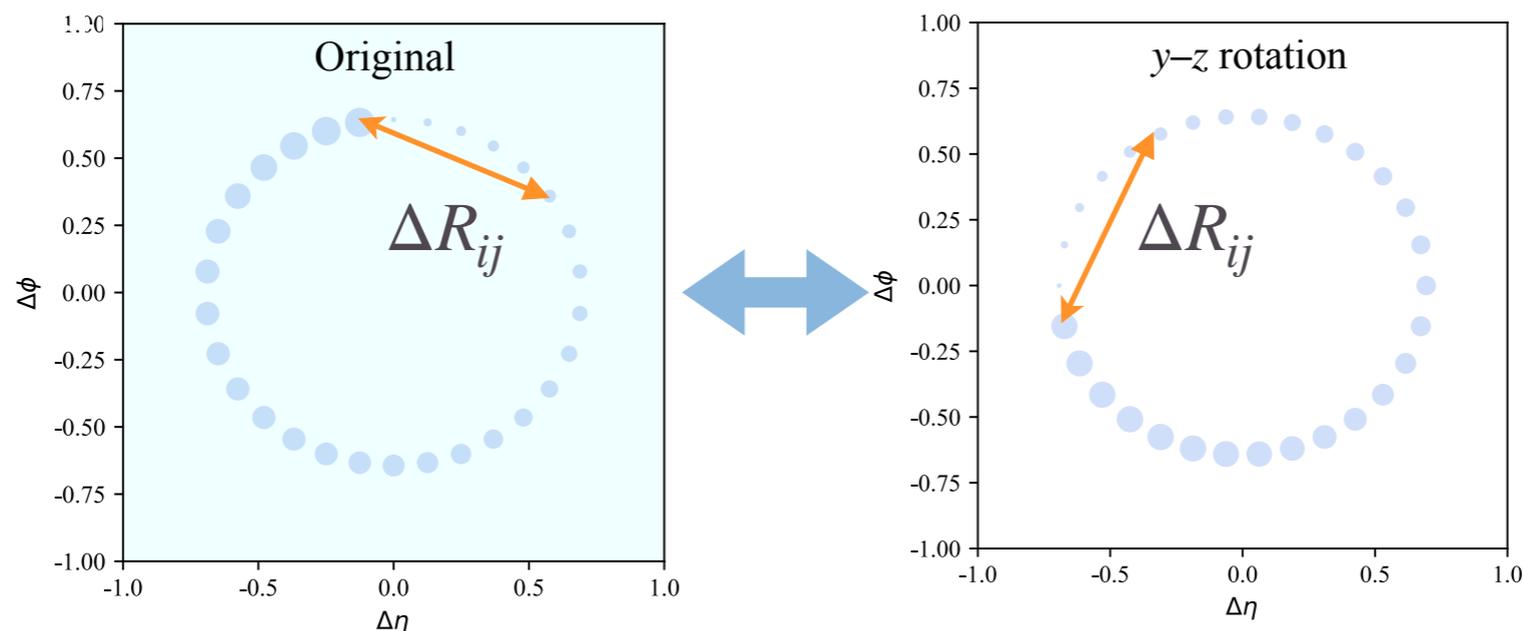
$y$ -boost +  $x$ - $y$  rotation

# Variable design

Pairwise variable	$z-t$ boost	$x-y$ rotation	$y-z$ rotation ( $y_{y,z} \sim o(1)$ )	$x-t$ boost ( $y_{y,z} \sim o(1)$ )	$z$ -tilt ( $y_{y,z} \sim o(1)$ )	$y$ -tilt ( $y_{y,z} \sim o(1)$ )
$m_{ij}^2$	✓	✓	✓	✓	✓	✓
$\Delta R_{ij}$	✓	✓	✓			
$\Delta R_{ij}(p_{T,i} + p_{T,j})$	✓	✓	✓	✓		
$E_{ij}$ (ablation study)		✓	✓			

→ Devise variables which are invariant under **some or all Lorentz (sub)symmetries**

- ❖ pairwise mass: invariant under all transformations
- ❖ pairwise  $\Delta R_{ij}$ : approx. invariant under  $y-z$  rotation ( $\approx \eta-\phi$  rotation)
- ❖

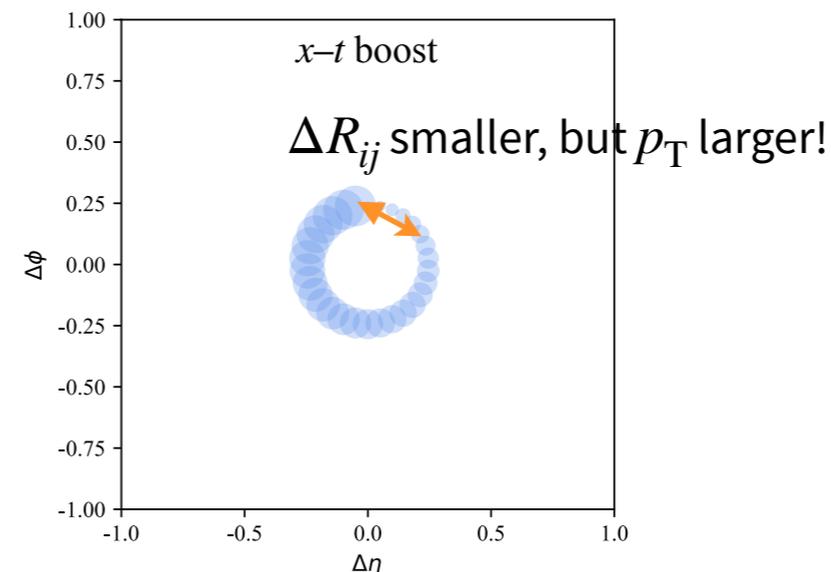
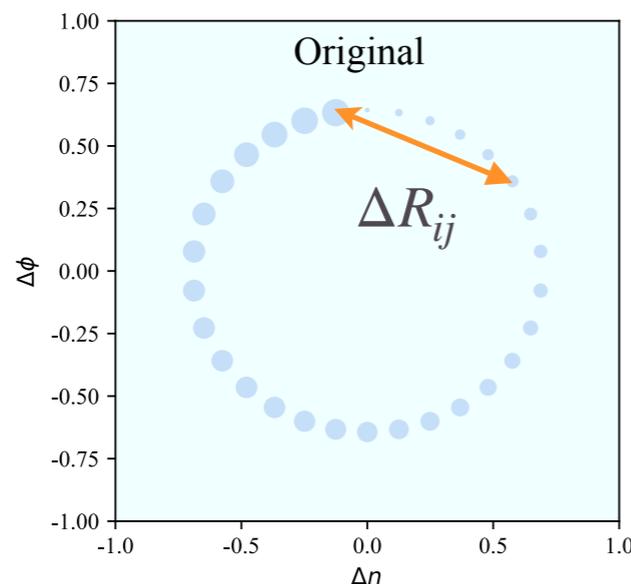


# Variable design

Pairwise variable	$z-t$ boost	$x-y$ rotation	$y-z$ rotation ( $y_{y,z} \sim o(1)$ )	$x-t$ boost ( $y_{y,z} \sim o(1)$ )	$z$ -tilt ( $y_{y,z} \sim o(1)$ )	$y$ -tilt ( $y_{y,z} \sim o(1)$ )
$m_{ij}^2$	✓	✓	✓	✓	✓	✓
$\Delta R_{ij}$	✓	✓	✓			
$\Delta R_{ij}(p_{T,i} + p_{T,j})$	✓	✓	✓	✓		
$E_{ij}$ (ablation study)		✓	✓			

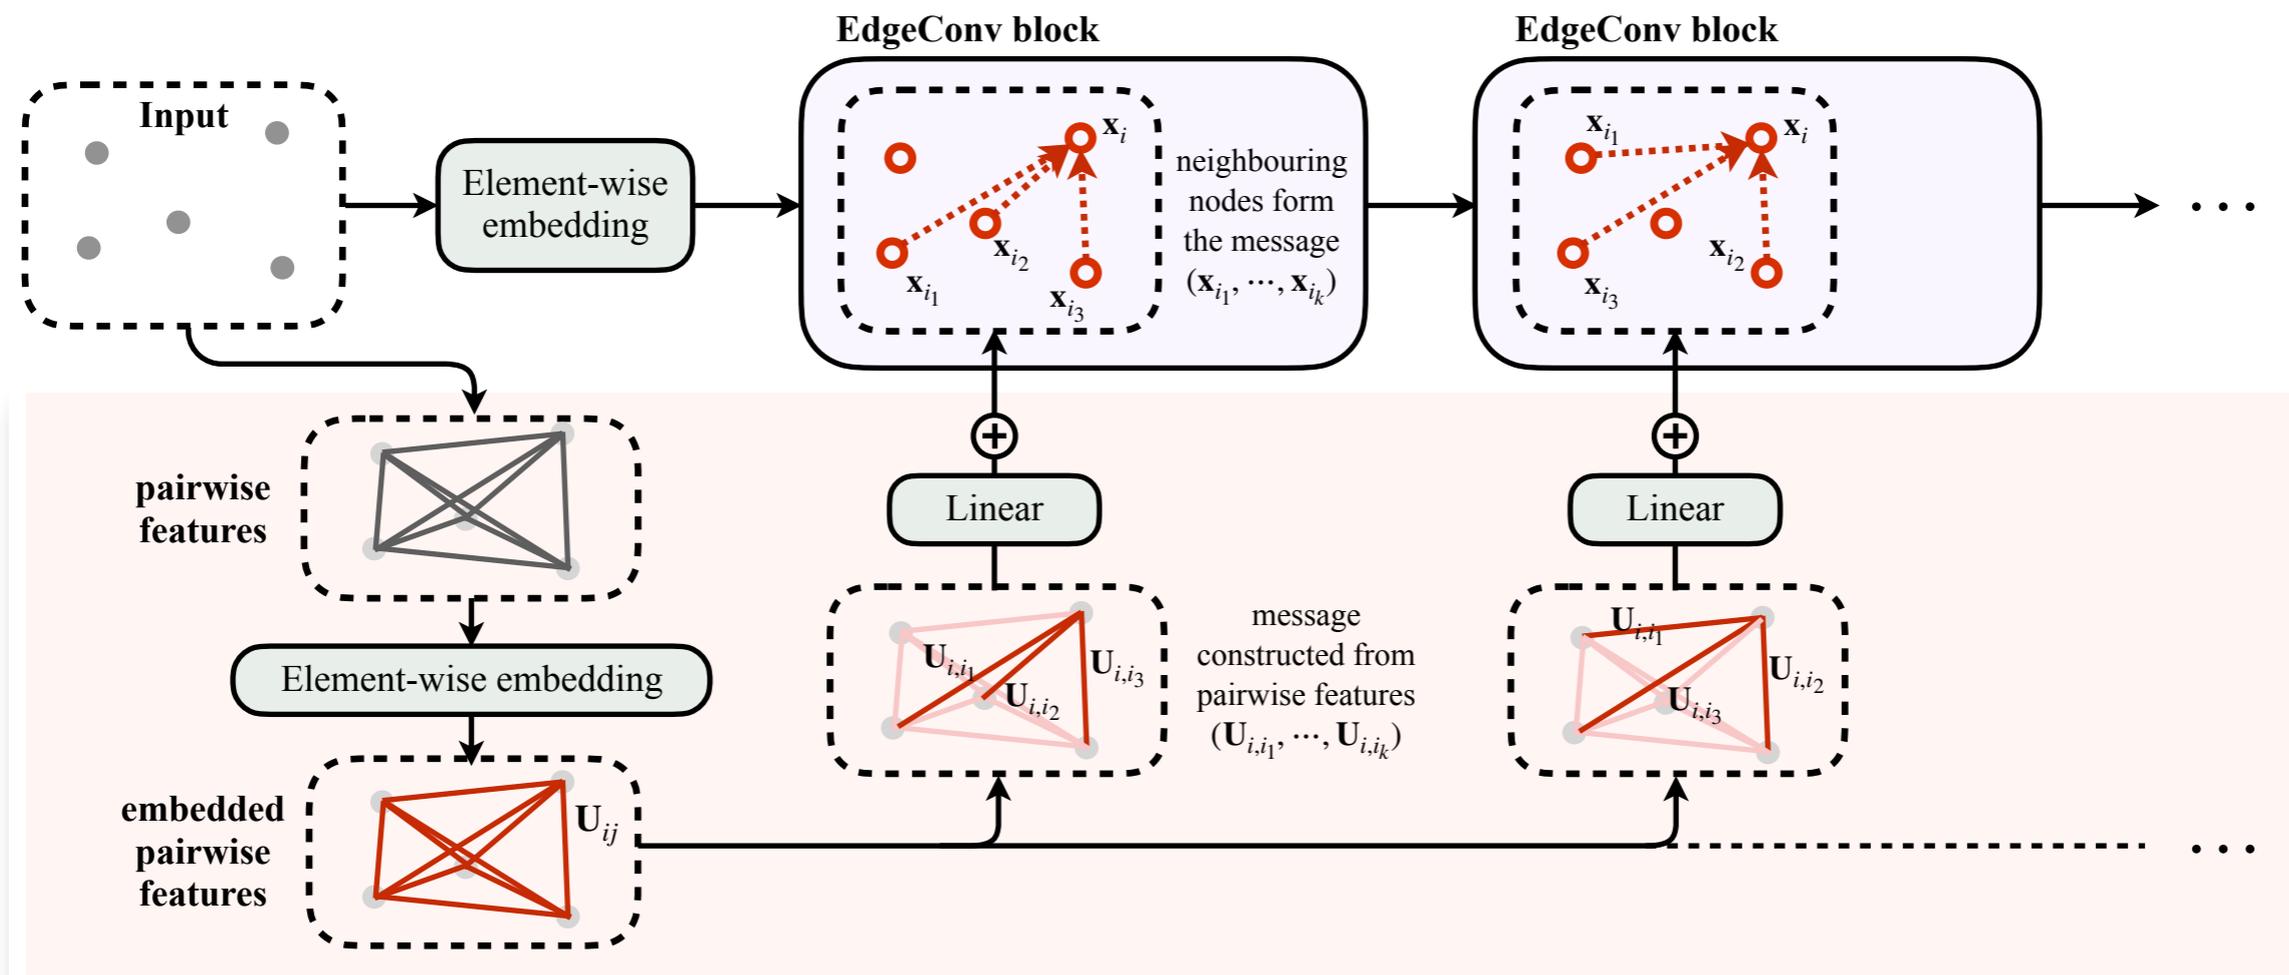
→ Devise variables which are invariant under **some or all Lorentz (sub)symmetries**

- ❖ pairwise mass: invariant under all transformations
- ❖ pairwise  $\Delta R_{ij}$ : approx. invariant under  $y-z$  rotation ( $\approx \eta-\phi$  rotation)
- ❖ manually construct variable  $\Delta R_{ij}(p_{T,i} + p_{T,j})$ : can prove that it is also approx. invariant under  $x$ -boost!



# Experiments on ParticleNet and LorentzNet

from paper [arXiv:2208.07814](https://arxiv.org/abs/2208.07814)



→ Two baseline networks to study pairwise feature effect: ParticleNet & LorentzNet<sub>base</sub>

- ❖ **ParticleNet**: now add an additional patch (in red colour) to incorporate pairwise features, based on ParticleNet's intrinsic kNN pairs
- ❖ **LorentzNet<sub>base</sub>**: LorentzNet has already included “pairwise mass”: remove it to create our baseline (but complete all node features as the case of ParticleNet)

# Performance for adding pairwise features

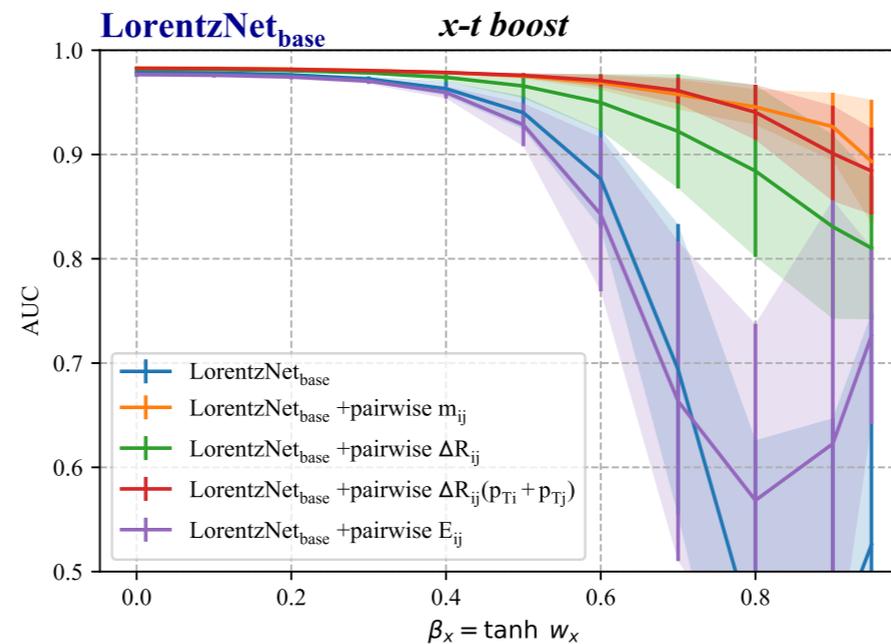
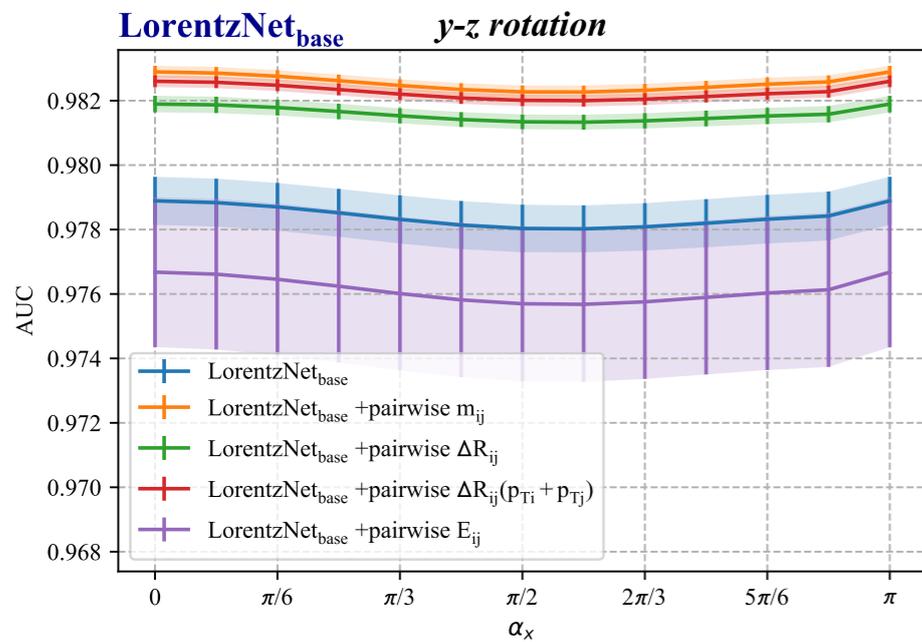
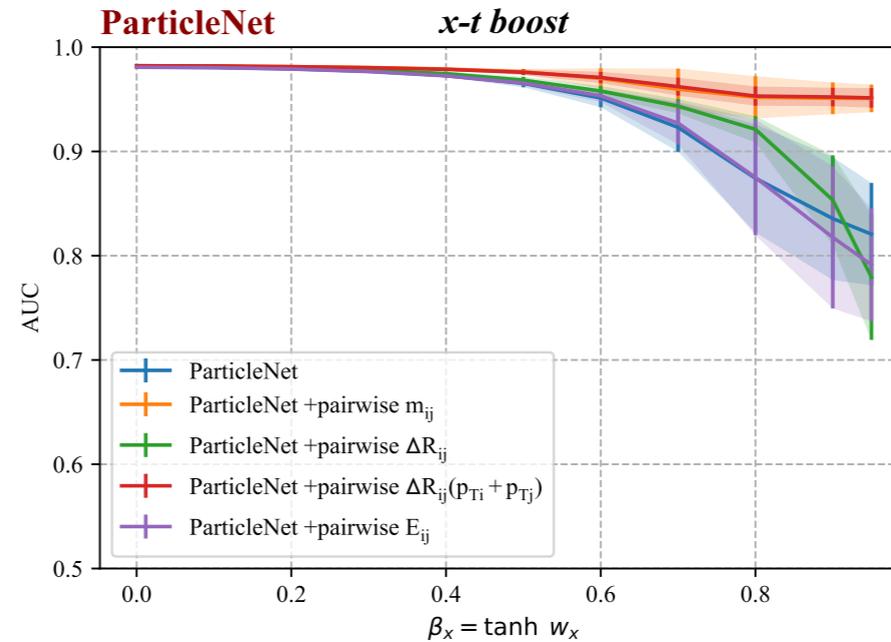
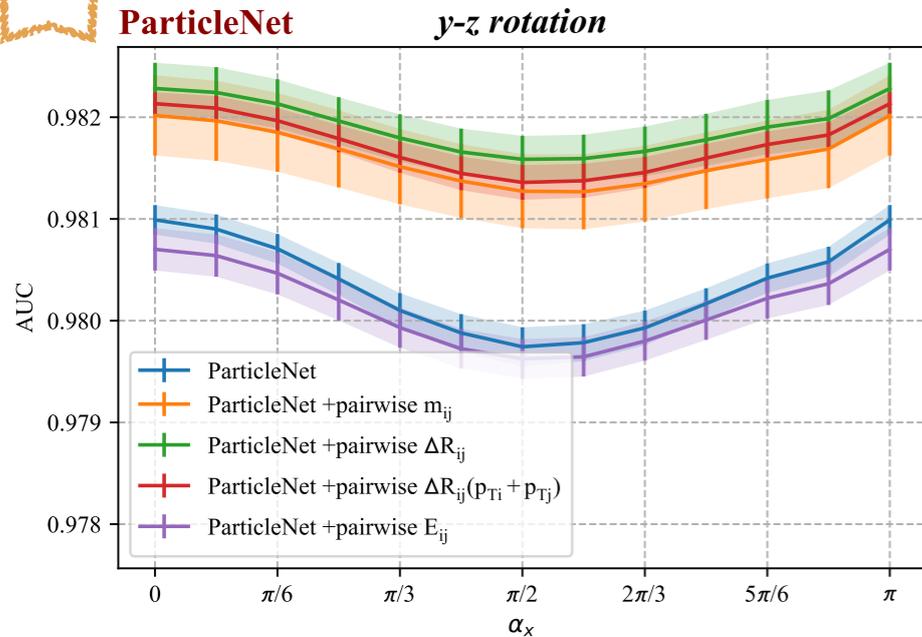
**Training on 60k top tagging dataset** (smaller dataset manifest the power of inductive bias)



Base model	Variation	Accuracy	AUC	$1/\epsilon_B$ ( $\epsilon_S = 50\%$ )	$1/\epsilon_B$ ( $\epsilon_S = 30\%$ )
ParticleNet	—	0.9310(3)	0.9810(2)	$198 \pm 7$	$640 \pm 29$
	+pairwise: $m_{ij}$	0.9334(8)	0.9820(4)	$222 \pm 13$	$722 \pm 52$
	+pairwise: $\Delta R_{ij}$	0.9334(6)	<b>0.9823(3)</b>	<b><math>231 \pm 10</math></b>	<b><math>752 \pm 43</math></b>
	+pairwise: $\Delta R_{ij}(p_{T,i} + p_{T,j})$	<b>0.9337(3)</b>	0.9821(1)	$223 \pm 6$	$741 \pm 36$
	+pairwise: $E_{ij}$	0.9303(5)	0.9807(2)	$200 \pm 6$	$651 \pm 23$
LorentzNet <sub>base</sub>	—	0.9276(12)	0.9789(7)	$172 \pm 13$	$581 \pm 53$
	+pairwise: $m_{ij}$	<b>0.9347(4)</b>	<b>0.9829(2)</b>	<b><math>260 \pm 6</math></b>	<b><math>931 \pm 50</math></b>
	+pairwise: $\Delta R_{ij}$	0.9328(4)	0.9819(3)	$232 \pm 10$	$807 \pm 35$
	+pairwise: $\Delta R_{ij}(p_{T,i} + p_{T,j})$	0.9342(4)	0.9826(2)	$251 \pm 6$	$919 \pm 34$
	+pairwise: $E_{ij}$	0.9243(37)	0.9767(23)	$144 \pm 29$	$485 \pm 108$

***better  
compared  
to baselines***

# Performance for adding pairwise features



▶ *Injecting  $\Delta R$  to the network  $\rightarrow$  more robust to y-z rotation*

▶ *Injecting  $\Delta R(p_{Ti}+p_{Tj})$  or mass  $\rightarrow$  more robust to y-z rotation and now also the x-boost*

# Performance for adding pairwise features

→ What does it mean?

- ❖ the full network tends to be more robust, even when we **introduce a very small patch structure** invariant under a certain symmetry (the original network is unaffected)
- ❖ **it means the invariance property of the small sub-network has a big impact on the learning, and can be reflected in the entire network**

Base model	Variation	# parameters	FLOPs
PFN	—	83.84 k	4.46 M
	+node-wise	+26.19 k	+3.41 M
ParticleNet	—	366.16 k	535.73 M
	+pairwise	+34.91 k	+285.29 M
	+node-wise	+21.97 k	+2.83 M
LorentzNet <sub>base</sub>	—	226.23 k	1997.69 M
	+pairwise	+0.43 k	+7.02 M
	+node-wise	+37.35 k	+4.8 M

# Performance for adding pairwise features

→ What does it mean?

- ❖ the full network tends to be more robust, even when we **introduce a very small patch structure** invariant under a certain symmetry (the original network is unaffected)
- ❖ **it means the invariance property of the small sub-network has a big impact on the learning, and can be reflected in the entire network**

Base model	Variation	# parameters	FLOPs
PFN	—	83.84 k	4.46 M
	+node-wise	+26.19 k	+3.41 M
ParticleNet	—	366.16 k	535.73 M
	+pairwise	+34.91 k	+285.29 M
	+node-wise	+21.97 k	+2.83 M
LorentzNet <sub>base</sub>	—	226.23 k	1997.69 M
	+pairwise	+0.43 k	+7.02 M
	+node-wise	+37.35 k	+4.8 M

- This experiment shows that “**pairwise mass**” is the key component
- **We reveal that the underlying logic lies in the Lorentz symmetry preservation!**
- Besides, we device other ways to add variables which respect Lorentz symmetries to boost the network performance → see more experiments in [arXiv:2208.07814](https://arxiv.org/abs/2208.07814)

# Conclusion

- We present LorentzNet, a Lorentz group equivariant GNN
- ❖ the network has now reached state-of-the-art performance, when trained and evaluated on two mainstream benchmarks

## *additional take-home message*

- ▶ **Key component for success of LorentzNet?**
  - ▶ fully-connected GNN as the backbone
  - ▶ Lorentz-symmetry-preserving design to respect the inductive bias of jet data
- ▶ **Possible hints to other work?**
  - ▶ the “pairwise mass” is one crucial component of the network—reason lies in the symmetry preservation
  - ▶ it can be used as a patch for many networks - in fact, [ParticleNeXt](#), [LorentzNet](#), [ParT](#) have implemented it in design
- ▶ ***Let's enjoy our new understanding of jet tagging in the DL era, and embrace the new improvement ahead***

# Backup

---