Emergent Equivariance in Deep Ensembles

Jan E. Gerken

CHALMERS UNIVERSITY OF TECHNOLOGY

UNIVERSITY OF GOTHENBURG

WASP WALLENBERG AI, AUTONOMOUS SYSTEMS AND SOFTWARE PROGRAM

in collaboration with

Pan Kessel

Philipp Misof

Symmetries in deep learning

Symmetries in deep learning

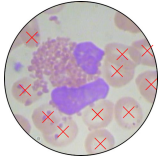

Symmetries in deep learning

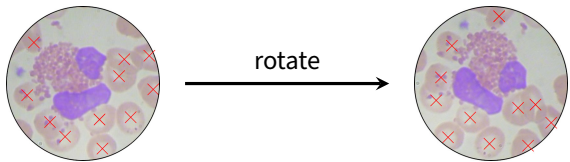

Symmetries in deep learning

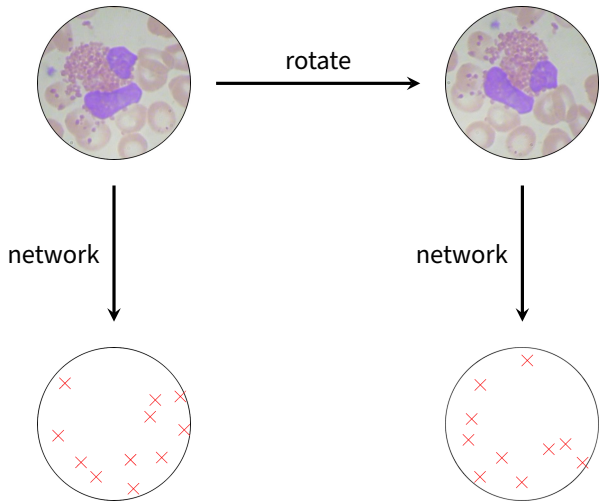

Symmetries in deep learning

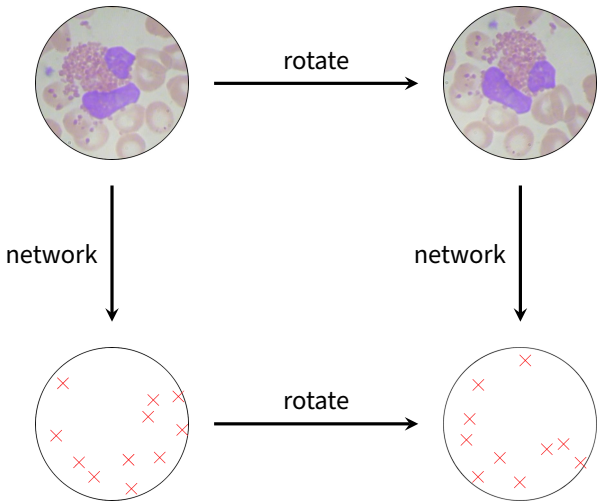

Symmetries in deep learning

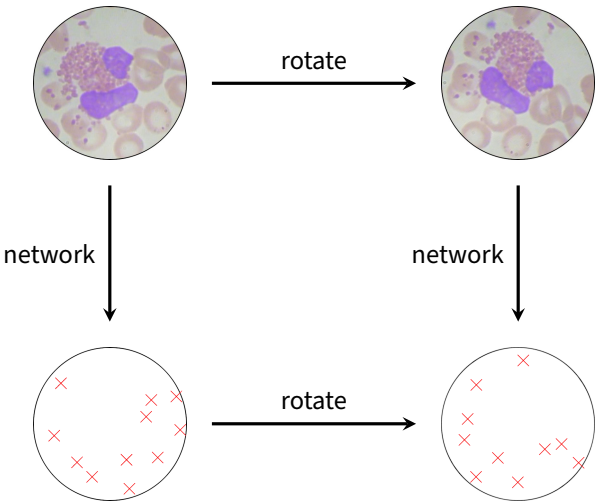

$$\iff \mathcal{N}(\rho_{\text{in}}(g)x) = \rho_{\text{out}}(g)\mathcal{N}(x)$$

Symmetries in deep learning

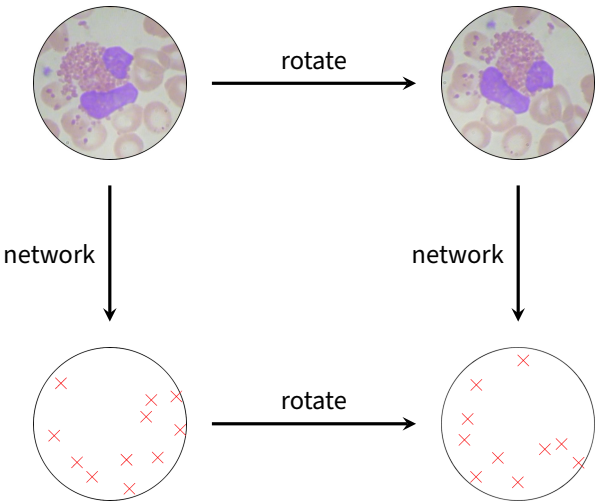

$$\mathcal{N}(\rho_{\text{in}}(g)x) = \rho_{\text{out}}(g)\mathcal{N}(x)$$

network

rotation

image

Equivariance

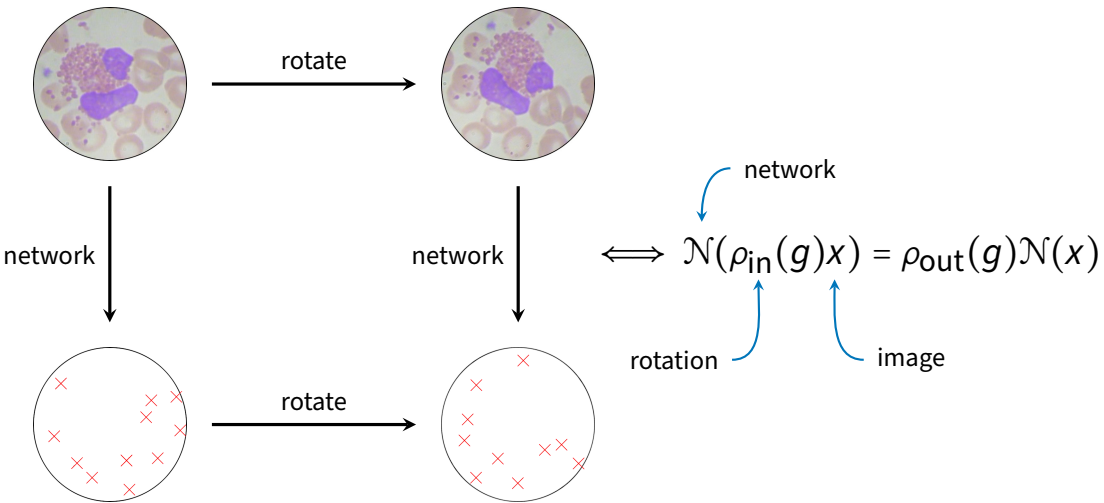

Equivariant neural networks

Equivariant neural networks

Group Equivariant Convolutional Networks

Taco S. Cohen T.S.COHEN@UVA.NL
University of Amsterdam

Max Welling M.WELLING@UVA.NL
University of Amsterdam
University of California Irvine
Canadian Institute for Advanced Research

Abstract

We introduce Group equivariant Convolutional Neural Networks (G-CNNs), a natural generalization of convolutional neural networks that reduces sample complexity by exploiting symmetries. Convolution layers can be used effectively in a *deep* network because all the layers in such a network are *translation equivariant*: shifting the image and then feeding it through a number of layers is the same as feeding the original image through the same layers and then shifting the resultinze feature maps (at least un to edge-effects). In

Equivariant neural networks

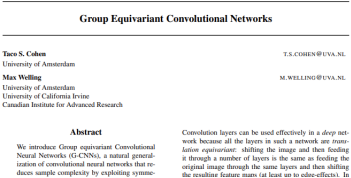

Group Equivariant Convolutional Networks

Taco S. Cohen T.S.COHEN@UVA.NL
University of Amsterdam

Max Welling M.WELLING@UVA.NL
University of Amsterdam
University of California Irvine
Canadian Institute for Advanced Research

Abstract

We introduce Group equivariant Convolutional Neural Networks (G-CNNs), a natural generalization of convolutional neural networks that reduces sample complexity by exploiting symme-

Convolution layers can be used effectively in a *deep* network because all the layers in such a network are *translation equivariant*: shifting the image and then feeding it through a number of layers is the same as feeding the original image through the same layers and then shifting the resultine feature maps (at least up to edge-effects). In

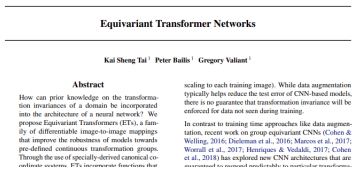

Equivariant Transformer Networks

Kai Sheng Tai [1] **Peter Bailis** [1] **Gregory Valiant** [1]

Abstract

How can prior knowledge on the transformation invariances of a domain be incorporated into the architecture of a neural network? We propose Equivariant Transformers (ETs), a family of differentiable image-to-image mappings that improve the robustness of models towards pre-defined continuous transformation groups. Through the use of specially-derived canonical coordinate systems, ETs incorporate functions that

scaling to each training image). While data augmentation typically helps reduce the test error of CNN-based models, there is no guarantee that transformation invariance will be enforced for data not seen during training.

In contrast to training time approaches like data augmentation, recent work on group equivariant CNNs (Cohen & Welling, 2016; Dieleman et al., 2016; Marcos et al., 2017; Worrall et al., 2017; Henriques & Vedaldi, 2017; Cohen et al., 2018) has explored new CNN architectures that are

Equivariant neural networks

Group Equivariant Convolutional Networks

Taco S. Cohen T.S.COHEN@UVA.NL
University of Amsterdam

Max Welling M.WELLING@UVA.NL
University of Amsterdam
University of California Irvine
Canadian Institute for Advanced Research

Abstract

We introduce Group equivariant Convolutional Neural Networks (G-CNNs), a natural generalization of convolutional neural networks that reduces sample complexity by exploiting symmetries. Convolution layers can be used effectively in a *deep* network because all the layers in such a network are *translation equivariant*: shifting the image and then feeding it through a number of layers is the same as feeding the original image through the same layers and then shifting the resultline feature maps (at least up to edge-effects). In

Equivariant Transformer Networks

Kai Sheng Tai [1] **Peter Bailis** [1] **Gregory Valiant** [1]

Abstract

How can prior knowledge on the transformation invariances of a domain be incorporated into the architecture of a neural network? We propose Equivariant Transformers (ETs), a family of differentiable image-to-image mappings that improve the robustness of models towards pre-defined continuous transformation groups. Through the use of specially-derived canonical coordinate systems, ETs incorporate functions that

scaling to each training image). While data augmentation typically helps reduce the test error of CNN-based models, there is no guarantee that transformation invariance will be enforced for data not seen during training.

In contrast to training time approaches like data augmentation, recent work on group equivariant CNNs (Cohen & Welling, 2016; Dieleman et al., 2016; Marcos et al., 2017; Worrall et al., 2017; Henriques & Vedaldi, 2017; Cohen et al., 2018) has explored new CNN architectures that are

Theory for Equivariant Quantum Neural Networks

Quynh T. Nguyen,[1,2] Louis Schatzki,[3,4] Paolo Braccia,[1,5] Michael Ragone,[1,6] Patrick J. Coles,[1] Frédéric Sauvage,[1] Martín Larocca,[1,7] and M. Cerezo[3]

[1] *Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*
[2] *School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts 02138, USA*
[3] *Information Sciences, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*
[4] *Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA*
[5] *Dipartimento di Fisica e Astronomia, Università di Firenze, Sesto Fiorentino (FI), 50019 , Italy*
[6] *Department of Mathematics, University of California Davis, Davis, California 95616, USA*
[7] *Center for Nonlinear Studies, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*

Quantum neural network architectures that have little-to-no inductive biases are known to face trainability and generalization issues. Inspired by a similar problem, recent breakthroughs in machine learning address this challenge by creating models encoding the symmetries of the learning task. This is materialized through the usage of equivariant neural networks whose action commutes with that of the symmetries. In this work, we instead draw ideas to the quantum realm by

Equivariant neural networks

Group Equivariant Convolutional Networks

Taco S. Cohen T.S.COHEN@UVA.NL
University of Amsterdam

Max Welling M.WELLING@UVA.NL
University of Amsterdam
University of California Irvine
Canadian Institute for Advanced Research

Abstract

We introduce Group equivariant Convolutional Neural Networks (G-CNNs), a natural generalization of convolutional neural networks that reduces sample complexity by exploiting symmetries. Convolution layers can be used effectively in a *deep* network because all the layers in such a network are *translation equivariant*: shifting the image and then feeding it through a number of layers is the same as feeding the original image through the same layers and then shifting the resultline feature maps (at least up to edge-effects). In

Equivariant Transformer Networks

Kai Sheng Tai [1] **Peter Bailis** [1] **Gregory Valiant** [1]

Abstract

How can prior knowledge on the transformation invariances of a domain be incorporated into the architecture of a neural network? We propose Equivariant Transformers (ETs), a family of differentiable image-to-image mappings that improve the robustness of models towards pre-defined continuous transformation groups. Through the use of specially-derived canonical coordinate systems, ETs incorporate functions that

scaling to each training image). While data augmentation typically helps reduce the test error of CNN-based models, there is no guarantee that transformation invariance will be enforced for data not seen during training.

In contrast to training time approaches like data augmentation, recent work on group equivariant CNNs (Cohen & Welling, 2016; Dieleman et al., 2016; Marcos et al., 2017; Worrall et al., 2017; Henriques & Vedaldi, 2017; Cohen et al., 2018) has explored new CNN architectures that are guaranteed to respond predictably to particular transforma-

Theory for Equivariant Quantum Neural Networks

Quynh T. Nguyen,[1,2] Louis Schatzki,[3,4] Paolo Braccia,[1,5] Michael Ragone,[1,6] Patrick J. Coles,[1] Frédéric Sauvage,[1] Martín Larocca,[1,7] and M. Cerezo[3]

[1] *Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*
[2] *School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts 02138, USA*
[3] *Information Sciences, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*
[4] *Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA*
[5] *Dipartimento di Fisica e Astronomia, Università di Firenze, Sesto Fiorentino (FI), 50019 , Italy*
[6] *Department of Mathematics, University of California Davis, Davis, California 95616, USA*
[7] *Center for Nonlinear Studies, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*

Quantum neural network architectures that have little-to-no inductive biases are known to face trainability and generalization issues. Inspired by a similar problem, recent breakthroughs in machine learning address this challenge by creating models encoding the symmetries of the learning task. This is materialized through the usage of equivariant neural networks whose action commutes with that of the symmetries. In this work, we instead draw ideas to the quantum realm by

An Efficient Lorentz Equivariant Graph Neural Network for Jet Tagging

Shiqi Gong[a,c,1] **Qi Meng**[b] **Jue Zhang**[b] **Huilin Qu**[c] **Congqiao Li**[d] **Sitian Qian**[d] **Weitao Du**[a] **Zhi-Ming Ma**[a] **Tie-Yan Liu**[b]

[a] *Academy of Mathematics and Systems Science, Chinese Academy of Sciences, Zhongguancun East Road, Beijing 100190, China*
[b] *Microsoft Research Asia, Danling Street, Beijing 100080, China*
[c] *CERN, EP Department, CH-1211 Geneva 23, Switzerland*
[d] *School of Physics, Peking University, Chengfu Road, Beijing 100871, China*

Equivariant neural networks

Group Equivariant Convolutional Networks

Taco S. Cohen T.S.COHEN@UVA.NL
University of Amsterdam

Max Welling M.WELLING@UVA.NL
University of Amsterdam
University of California Irvine
Canadian Institute for Advanced Research

Abstract

We introduce Group equivariant Convolutional Neural Networks (G-CNNs), a natural generalization of convolutional neural networks that reduces sample complexity by exploiting symmetry. Convolution layers can be used effectively in a *deep* network because all the layers in such a network are *translation equivariant*: shifting the image and then feeding it through a number of layers is the same as feeding the original image through the same layers and then shifting the resultiine feature maps (at least up to edge-effects). In

Equivariant Transformer Networks

Kai Sheng Tai [1] **Peter Bailis** [1] **Gregory Valiant** [1]

Abstract

How can prior knowledge on the transformation invariances of a domain be incorporated into the architecture of a neural network? We propose Equivariant Transformers (ETs), a family of differentiable image-to-image mappings that improve the robustness of models towards pre-defined continuous transformation groups. Through the use of specially-derived canonical coordinate systems, ETs incorporate functions that

scaling to each training image). While data augmentation typically helps reduce the test error of CNN-based models, there is no guarantee that transformation invariance will be enforced for data not seen during training.

In contrast to training time approaches like data augmentation, recent work on group equivariant CNNs (Cohen & Welling, 2016; Dieleman et al., 2016; Marcos et al., 2017; Worrall et al., 2017; Henriques & Vedaldi, 2017; Cohen et al., 2018) has explored new CNN architectures that are guaranteed to respond predictably to particular transforma-

Theory for Equivariant Quantum Neural Networks

Quynh T. Nguyen,[1,2] Louis Schatzki,[3,4] Paolo Braccia,[1,5] Michael Ragone,[1,6] Patrick J. Coles,[1] Frédéric Sauvage,[1] Martín Larocca,[1,7] and M. Cerezo[3]

[1] *Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*
[2] *School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts 02138, USA*
[3] *Information Sciences, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*
[4] *Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA*
[5] *Dipartimento di Fisica e Astronomia, Università di Firenze, Sesto Fiorentino (FI), 50019 , Italy*
[6] *Department of Mathematics, University of California Davis, Davis, California 95616, USA*
[7] *Center for Nonlinear Studies, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*

Quantum neural network architectures that have little-to-no inductive biases are known to face trainability and generalization issues. Inspired by a similar problem, recent breakthroughs in machine learning address this challenge by creating models encoding the symmetries of the learning task. This is materialized through the usage of equivariant neural networks whose action com-

An Efficient Lorentz Equivariant Graph Neural Network for Jet Tagging

Shiqi Gong[a,c,1] **Qi Meng**[b] **Jue Zhang**[b] **Huilin Qu**[c] **Congqiao Li**[d] **Sitian Qian**[d] **Weitao Du**[e] **Zhi-Ming Ma**[a] **Tie-Yan Liu**[c]

[a] *Academy of Mathematics and Systems Science, Chinese Academy of Sciences, Zhongguancun East Road, Beijing 100190, China*
[b] *Microsoft Research Asia, Danling Street, Beijing 100080, China*
[c] *CERN, EP Department, CH-1211 Geneva 23, Switzerland*
[d] *School of Physics, Peking University, Chengfu Road, Beijing 100871, China*

E(3)-Equivariant Graph Neural Networks for Data-Efficient and Accurate Interatomic Potentials

Simon Batzner*,[1] Albert Musaelian,[1] Lixin Sun,[1] Mario Geiger,[2] Jonathan P. Mailoa,[3] Mordechai Kornbluth,[3] Nicola Molinari,[1] Tess E. Smidt,[4,5] and Boris Kozinsky*[1,3]

[1] *John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, USA*
[2] *École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland*
[3] *Robert Bosch Research and Technology Center, Cambridge, MA 02139, USA*
[4] *Computational Research Division and Center for Advanced Mathematics for Energy Research Applications, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA*
[5] *Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, Cambridge, MA 02142, USA*

This work presents Neural Equivariant Interatomic Potentials (NequIP), an E(3)-equivariant neural network approach for learning interatomic potentials from ab-*initio* calculations for molecular dynamics simulations. While most contemporary symmetry-aware models use invariant convolutions and only act on scalars, NequIP employs E(3)-equivariant convolutions for interactions of geometric tensors, resulting in a more information-rich and faithful representation of atomic environments. The method achieves state-of-the-art accuracy on a challenging and diverse set of molecules and

Equivariant neural networks

Group Equivariant Convolutional Networks

Taco S. Cohen T.S.COHEN@UVA.NL
University of Amsterdam

Max Welling M.WELLING@UVA.NL
University of Amsterdam
University of California Irvine
Canadian Institute for Advanced Research

Abstract

We introduce Group equivariant Convolutional Neural Networks (G-CNNs), a natural generalization of convolutional neural networks that reduces sample complexity by exploiting symme-

Convolution layers can be used effectively in a *deep* network because all the layers in such a network are *translation equivariant*: shifting the image and then feeding it through a number of layers is the same as feeding the original image through the same layers and then shifting the resultine feature maps (at least up to edge-effects). In

Equivariant Transformer Networks

Kai Sheng Tai [1] **Peter Bailis** [1] **Gregory Valiant** [1]

Abstract

How can prior knowledge on the transformation invariances of a domain be incorporated into the architecture of a neural network? We propose Equivariant Transformers (ETs), a family of differentiable image-to-image mappings that improve the robustness of models towards pre-defined continuous transformation groups. Through the use of specially-derived canonical coordinate systems, ETs incorporate functions that

scaling to each training image). While data augmentation typically helps reduce the test error of CNN-based models, there is no guarantee that transformation invariance will be enforced for data not seen during training.

In contrast to training time approaches like data augmentation, recent work on group equivariant CNNs (Cohen & Welling, 2016; Dieleman et al., 2016; Marcos et al., 2017; Worrall et al., 2017; Henriques & Vedaldi, 2017; Cohen et al., 2018) has explored new CNN architectures that are

Theory for Equivariant Quantum Neural Networks

Quynh T. Nguyen,[1,2] Louis Schatzki,[3,4] Paolo Braccia,[1,5] Michael Ragone,[1,6] Patrick J. Coles,[1] Frédéric Sauvage,[1] Martín Larocca,[1,7] and M. Cerezo[3]

[1]*Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*
[2]*School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts 02138, USA*
[3]*Information Sciences, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*
[4]*Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA*
[5]*Dipartimento di Fisica e Astronomia, Università di Firenze, Sesto Fiorentino (FI), 50019 , Italy*
[6]*Department of Mathematics, University of California Davis, Davis, California 95616, USA*
[7]*Center for Nonlinear Studies, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*

Quantum neural network architectures that have little-to-no inductive biases are known to face trainability and generalization issues. Inspired by a similar problem, recent breakthroughs in machine learning address this challenge by creating models encoding the symmetries of the learning task. This is materialized through the usage of equivariant neural networks whose action commutes with that of the symmetry. In this work, we instead draw ideas to the quantum realm by

An Efficient Lorentz Equivariant Graph Neural Network for Jet Tagging

Shiqi Gong[a,c,1] **Qi Meng**[b] **Jue Zhang**[b] **Huilin Qu**[c] **Congqiao Li**[d] **Sitian Qian**[d] **Weitao Du**[a] **Zhi-Ming Ma**[a] **Tie-Yan Liu**[c]

[a]*Academy of Mathematics and Systems Science, Chinese Academy of Sciences, Zhongguancun East Road, Beijing 100190, China*
[b]*Microsoft Research Asia, Danling Street, Beijing 100080, China*
[c]*CERN, EP Department, CH-1211 Geneva 23, Switzerland*
[d]*School of Physics, Peking University, Chengfu Road, Beijing 100871, China*

E(3)-Equivariant Graph Neural Networks for Data-Efficient and Accurate Interatomic Potentials

Simon Batzner*,[1] Albert Musaelian,[1] Lixin Sun,[1] Mario Geiger,[2] Jonathan P. Mailoa,[3] Mordechai Kornbluth,[3] Nicola Molinari,[1] Tess E. Smidt,[4,5] and Boris Kozinsky*[1,3]

[1]*John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, USA*
[2]*École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland*
[3]*Robert Bosch Research and Technology Center, Cambridge, MA 02139, USA*
[4]*Computational Research Division and Center for Advanced Mathematics for Energy Research Applications, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA*
[5]*Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, Cambridge, MA 02142, USA*

This work presents Neural Equivariant Interatomic Potentials (NequIP), an E(3)-equivariant neural network approach for learning interatomic potentials from ab-initio calculations for molecular dynamics simulations. While most contemporary symmetry-aware models use invariant convolutions and only act on scalars, NequIP employs E(3)-equivariant convolutions for interactions of geometric tensors, resulting in a more information-rich and faithful representation of atomic environments. The method achieves state-of-the-art accuracy on a challenging and diverse set of molecules and

Hierarchical, rotation-equivariant neural networks to select structural models of protein complexes

Stephan Eismann* Department of Applied Physics Stanford University seismann@stanford.edu

Raphael J.L. Townshend* Department of Computer Science Stanford University raphael@cs.stanford.edu

Nathaniel Thomas* Department of Physics Stanford University nthomas103@gmail.com

Milind Jagota Department of Electrical Engineering Stanford University mjagota@stanford.edu

Bowen Jing Department of Computer Science Stanford University bjing@stanford.edu

Ron O. Dror Department of Computer Science Stanford University rondror@cs.stanford.edu

ABSTRACT

Predicting the structure of multi-protein complexes is a grand challenge in biochemistry, with major implications for basic science and drug discovery. Computational structure prediction methods generally leverage pre-defined structural features to distinguish accurate structural models from less accurate ones. This raises the question of whether it is possible to learn characteristics of accurate models directly from atomic coordinates of protein complexes, with no prior assumptions. Here we introduce a machine learning method that learns directly from the 3D positions of all atoms to

Equivariant neural networks

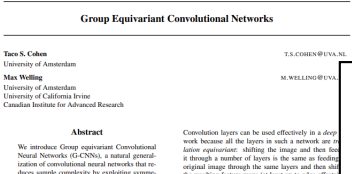

Group Equivariant Convolutional Networks

Taco S. Cohen T.S.COHEN@UVA.NL
University of Amsterdam

Max Welling M.WELLING@UVA.
University of Amsterdam
University of California Irvine
Canadian Institute for Advanced Research

Abstract

We introduce Group equivariant Convolutional Neural Networks (G-CNNs), a natural generalization of convolutional neural networks that reduces sample complexity by exploiting symme-

Convolution layers can be used effectively in a *deep* work because all the layers in such a network are *m* *lative equivariant*: shifting the image and then fee it through a number of layers is the same as feeding original image through the same layers and then shi the resultine feature maps (at least up to edge-effects

An Efficient Lorentz Equivariant Graph Neural Network for Jet Tagging

Shiqi Gong[a,c,1] **Qi Meng**[b] **Jue Zhang**[b] **Huilin Qu**[c] **Congqiao Li**[d] **Sitian Qian**[d] **Weitao Du**[e] **Zhi-Ming Ma**[a] **Tie-Yan Liu**[b]

[a] *Academy of Mathematics and Systems Science, Chinese Academy of Sciences, Zhongguancun East Road, Beijing 100190, China*
[b] *Microsoft Research Asia, Danling Street, Beijing 100080, China*
[c] *CERN, EP Department, CH-1211 Geneva 23, Switzerland*
[d] *School of Physics, Peking University, Chengfu Road, Beijing 100871, China*

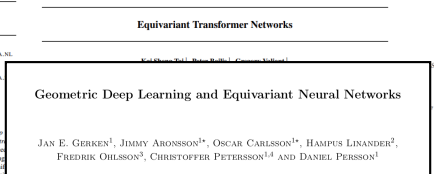

Equivariant Transformer Networks

Kai Sheng Tai[1] Peter Bailis[1] Gregory Valiant[1]

Geometric Deep Learning and Equivariant Neural Networks

Jan E. Gerken[1], Jimmy Aronsson[1*], Oscar Carlsson[1*], Hampus Linander[2], Fredrik Ohlsson[3], Christoffer Petersson[1,4] and Daniel Persson[1]

[1] *Chalmers University of Technology, Department of Mathematical Sciences SE-41296 Gothenburg, Sweden*

[2] *Gothenburg University, Department of Physics SE-41296, Gothenburg, Sweden*

[3] *Umeå University, Department of Mathematics and Mathematical Statistics SE-901 87, Umeå, Sweden*

[4] *Zenseact SE-41756, Gothenburg, Sweden*

[*] *equal contribution*

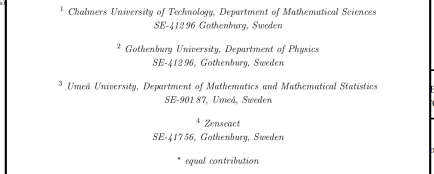

Engineering and Computer Science, Cambridge, MA 02142, USA

This work presents Neural Equivariant Interatomic Potentials (NequIP), an E(3)-equivariant neural network approach for learning interatomic potentials from ab-*initio* calculations for molecular dynamics simulations. While most contemporary symmetry-aware models use invariant convolutions and only act on scalars, NequIP employs E(3)-equivariant convolutions for interactions of geometric tensors, resulting in a more information-rich and faithful representation of atomic environments. The method achieves state-of-the-art accuracy on a challenging and diverse set of molecules and

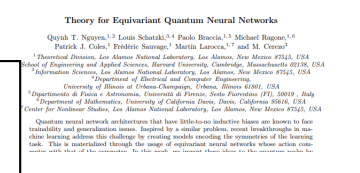

Theory for Equivariant Quantum Neural Networks

Quynh T. Nguyen,[1,2] Louis Schatzki,[3,4] Paolo Braccia,[1,5] Michael Ragone,[1,6] Patrick J. Coles,[1] Frédéric Sauvage,[1] Martín Larocca,[1,7] and M. Cerezo[3]

[1] *Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*
[2] *School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts 02138, USA*
[3] *Information Sciences, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*
[4] *Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA*
[5] *Dipartimento di Fisica e Astronomia, Università di Firenze, Sesto Fiorentino (FI), 50019 , Italy*
[6] *Department of Mathematics, University of California Davis, Davis, California 95616, USA*
[7] *Center for Nonlinear Studies, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*

Quantum neural network architectures that have little-to-no inductive biases are known to face trainability and generalization issues. Inspired by a similar problem, recent breakthroughs in machine learning address this challenge by creating models encoding the symmetries of the learning task. This is materialized through the usage of equivariant neural networks whose action com-

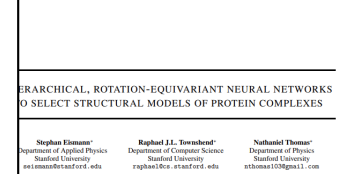

ERARCHICAL, ROTATION-EQUIVARIANT NEURAL NETWORKS O SELECT STRUCTURAL MODELS OF PROTEIN COMPLEXES

Stephan Eismann[*] Department of Applied Physics Stanford University seismann@stanford.edu

Raphael J.L. Townshend[*] Department of Computer Science Stanford University raphael@cs.stanford.edu

Nathaniel Thomas[*] Department of Physics Stanford University nthomas103@gmail.com

Milind Jagota Department of Electrical Engineering Stanford University mjagota@stanford.edu

Bowen Jing Department of Computer Science Stanford University bjing@stanford.edu

Ron O. Dror Department of Computer Science Stanford University rondror@cs.stanford.edu

ABSTRACT

Predicting the structure of multi-protein complexes is a grand challenge in biochemistry, with major implications for basic science and drug discovery. Computational structure prediction methods generally leverage pre-defined structural features to distinguish accurate structural models from less accurate ones. This raises the question of whether it is possible to learn characteristics of accurate models directly from atomic coordinates of protein complexes, with no prior assumptions. Here we introduce a machine learning method that learns directly from the 3D positions of all atoms to

Data augmentation

Data augmentation

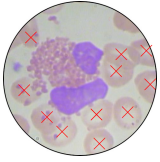

Data augmentation

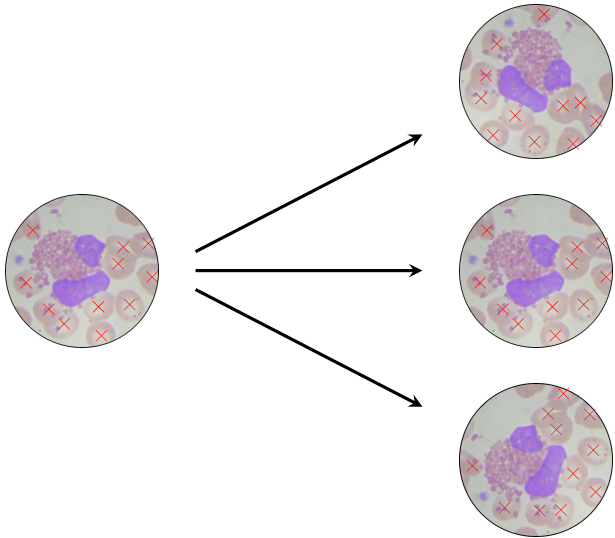

Data augmentation

Data augmentation

Article

Accurate structure prediction of biomolecular interactions with AlphaFold 3

https://doi.org/10.1038/s41586-024-07487-w

Received: 19 December 2023

Accepted: 29 April 2024

Published online: 8 May 2024

Open access

Check for updates

Josh Abramson[1,*], Jonas Adler[1,*], Jack Dunger[1,*], Richard Evans[1,*], Tim Green[1,*], Alexander Pritzel[1,*], Olaf Ronneberger[1,*], Lindsay Willmore[1,*], Andrew J. Ballard[1], Joshua Bambrick[1], Sebastian W. Bodenstein[1], David A. Evans[1], Chia-Chun Hung[1], Michael O'Neill[1], David Reiman[1], Kathryn Tunyasuvunakool[1], Zachary Wu[1], Akvilė Žemgulytė[1], Eirini Arvaniti[2], Charles Beattie[2], Ottavia Bertolli[2], Alex Bridgland[2], Alexey Cherepanov[2], Miles Congreve[2], Alexander I. Cowen-Rivers[2], Andrew Cowie[2], Michael Figurnov[2], Fabian B. Fuchs[2], Hannah Gladman[2], Rishub Jain[2], Yousuf A. Khan[2], Caroline M. R. Low[2], Kuba Perlin[2], Anna Potapenko[2], Pascal Savy[2], Sukhdeep Singh[2], Adrian Stecula[2], Ashok Thillaisundaram[2], Catherine Tong[2], Sergei Yakneen[2], Ellen D. Zhong[2,3], Michal Zielinski[2], Augustin Žídek[2], Victor Bapst[1,4], Pushmeet Kohli[1,4], Max Jaderberg[1,4,5], Demis Hassabis[1,4,5] ✉, John M. Jumper[1,4,5] ✉

The introduction of AlphaFold 2 has spurred a revolution in modelling the structure of proteins and their interactions, enabling a huge range of applications in protein modelling and design[1–6]. Here we describe our AlphaFold 3 model with a substantially updated diffusion-based architecture that is capable of predicting the joint structure of complexes including proteins, nucleic acids, small molecules, ions and modified residues. The new AlphaFold model demonstrates substantially improved accuracy

Data augmentation

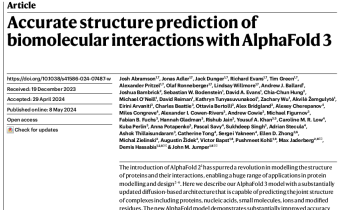

Article

Accurate structure prediction of biomolecular interactions with AlphaFold 3

https://doi.org/10.1038/s41586-024-07487-w

Received: 19 December 2023

Accepted: 29 April 2024

Published online: 8 May 2024

Open access

Check for updates

Josh Abramson[1,7], Jonas Adler[1,7], Jack Dunger[1,7], Richard Evans[1,7], Tim Green[1,7], Alexander Pritzel[1,7], Olaf Ronneberger[1,7], Lindsay Willmore[1,7], Andrew J. Ballard[1], Joshua Bambrick[1], Sebastian W. Bodenstein[1], David A. Evans[1], Chia-Chun Hung[1], Michael O'Neill[1], David Reiman[1], Kathryn Tunyasuvunakool[1], Zachary Wu[1], Akvilė Žemgulytė[1], Eirini Arvaniti[2], Charles Beattie[1], Ottavia Bertolli[2], Alex Bridgland[2], Alexey Cherepanov[1], Miles Congreve[2], Alexander I. Cowen-Rivers[1], Andrew Cowie[1], Michael Figurnov[1], Fabian B. Fuchs[1], Hannah Gladman[1], Rishub Jain[1], Yousuf A. Khan[1], Caroline M. R. Low[1], Kuba Perlin[1], Anna Potapenko[1], Pascal Savy[1], Sukhdeep Singh[1], Adrian Stecula[1], Ashok Thillaisundaram[1], Catherine Tong[1], Sergei Yakneen[1], Ellen D. Zhong[1,3], Michal Zielinski[1], Augustin Žídek[1], Victor Bapst[1,4], Pushmeet Kohli[1,4], Max Jaderberg[1,4,5], Demis Hassabis[1,4,6] ✉ & John M. Jumper[1,4,6] ✉

The introduction of AlphaFold 2 has spurred a revolution in modelling the structure of proteins and their interactions, enabling a huge range of applications in protein modelling and design[1–6]. Here we describe our AlphaFold 3 model with a substantially updated diffusion-based architecture that is capable of predicting the joint structure of complexes including proteins, nucleic acids, small molecules, ions and modified residues. The new AlphaFold model demonstrates substantially improved accuracy

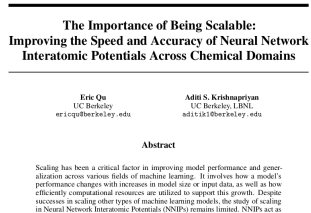

The Importance of Being Scalable: Improving the Speed and Accuracy of Neural Network Interatomic Potentials Across Chemical Domains

Eric Qu
UC Berkeley
ericqu@berkeley.edu

Aditi S. Krishnapriyan
UC Berkeley, LBNL
aditik1@berkeley.edu

Abstract

Scaling has been a critical factor in improving model performance and generalization across various fields of machine learning. It involves how a model's performance changes with increases in model size or input data, as well as how efficiently computational resources are utilized to support this growth. Despite successes in scaling other types of machine learning models, the study of scaling in Neural Network Interatomic Potentials (NNIPs) remains limited. NNIPs act as

Data augmentation

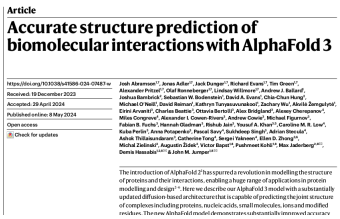

Article

Accurate structure prediction of biomolecular interactions with AlphaFold 3

https://doi.org/10.1038/s41586-024-07487-w

Received: 19 December 2023

Accepted: 29 April 2024

Published online: 8 May 2024

Open access

Check for updates

Josh Abramson[1,7], Jonas Adler[1,7], Jack Dunger[1,7], Richard Evans[1,7], Tim Green[1,7], Alexander Pritzel[1,7], Olaf Ronneberger[1,7], Lindsay Willmore[1,7], Andrew J. Ballard[1], Joshua Bambrick[1], Sebastian W. Bodenstein[1], David A. Evans[1], Chia-Chun Hung[1], Michael O'Neill[1], David Reiman[1], Kathryn Tunyasuvunakool[1], Zachary Wu[1], Akvilė Žemgulytė[1], Eirini Arvaniti[2], Charles Beattie[1], Ottavia Bertolli[2], Alex Bridgland[2], Alexey Cherepanov[1], Miles Congreve[2], Alexander I. Cowen-Rivers[1], Andrew Cowie[1], Michael Figurnov[1], Fabian B. Fuchs[1], Hannah Gladman[1], Rishub Jain[1], Yousuf A. Khan[2], Caroline M. R. Low[1], Kuba Perlin[1], Anna Potapenko[1], Pascal Savy[1], Sukhdeep Singh[1], Adrian Stecula[1], Ashok Thillaisundaram[1], Catherine Tong[1], Sergei Yakneen[1], Ellen D. Zhong[1,8], Michal Zielinski[1], Augustin Žídek[1], Victor Bapst[1,9], Pushmeet Kohli[1,9], Max Jaderberg[1,9,10], Demis Hassabis[1,9,10] ✉ & John M. Jumper[1,9,10] ✉

The introduction of AlphaFold 2 has spurred a revolution in modelling the structure of proteins and their interactions, enabling a huge range of applications in protein modelling and design[1–6]. Here we describe our AlphaFold 3 model with a substantially updated diffusion-based architecture that is capable of predicting the joint structure of complexes including proteins, nucleic acids, small molecules, ions and modified residues. The new AlphaFold model demonstrates substantially improved accuracy

The Importance of Being Scalable: Improving the Speed and Accuracy of Neural Network Interatomic Potentials Across Chemical Domains

Eric Qu
UC Berkeley
ericqu@berkeley.edu

Aditi S. Krishnapriyan
UC Berkeley, LBNL
aditik1@berkeley.edu

Abstract

Scaling has been a critical factor in improving model performance and generalization across various fields of machine learning. It involves how a model's performance changes with increases in model size or input data, as well as how efficiently computational resources are utilized to support this growth. Despite successes in scaling other types of machine learning models, the study of scaling in Neural Network Interatomic Potentials (NNIPs) remains limited. NNIPs act as

Swallowing the Bitter Pill: Simplified Scalable Conformer Generation

Yuyang Wang [1] **Ahmed A. Elhag** [1 2] **Navdeep Jaitly** [1] **Joshua M. Susskind** [1] **Miguel Ángel Bautista** [1]

Abstract

We present a novel way to predict molecular conformers through a simple formulation that sidesteps many of the heuristics of prior works and achieves state of the art results by using the advantages of scale. By training a diffusion generative model directly on 3D atomic positions without making assumptions about the explicit structure of molecules (e.g. modelling torsional angles) we are able to radically simplify stochastic learning and make it easier to scale up the

is the vast complexity of the 3D structure space, encompassing factors such as bond lengths and torsional angles. Despite the molecular graph dictating potential 3D conformers through specific constraints, such as bond types and spatial arrangements determined by chiral centers, the conformational space experiences exponential growth with the expansion of the graph size and the number of rotatable bonds (Axelrod & Gomez-Bombarelli, 2022). This complicates brute force and exhaustive approaches, making them virtually unfeasible for even moderately small molecules.

Systematic methods like FMFFGA (Hawkins et al., 2010)

Data augmentation

Article

Accurate structure prediction of biomolecular interactions with AlphaFold 3

https://doi.org/10.1038/s41586-024-07487-w

Received: 19 December 2023

Accepted: 29 April 2024

Published online: 8 May 2024

Open access

Check for updates

Josh Abramson[1,7], Jonas Adler[1,7], Jack Dunger[1,7], Richard Evans[1,7], Tim Green[1,7], Alexander Pritzel[1,7], Olaf Ronneberger[1,7], Lindsay Willmore[1,7], Andrew J. Ballard[1], Joshua Bambrick[1], Sebastian W. Bodenstein[1], David A. Evans[1], Chia-Chun Hung[1], Michael O'Neill[1], David Reiman[1], Kathryn Tunyasuvunakool[1], Zachary Wu[1], Akvilė Žemgulytė[1], Eirini Arvaniti[2], Charles Beattie[2], Ottavia Bertolli[2], Alex Bridgland[2], Alexey Cherepanov[2], Miles Congreve[2], Alexander I. Cowen-Rivers[2], Andrew Cowie[2], Michael Figurnov[2], Fabian B. Fuchs[2], Hannah Gladman[2], Rishub Jain[2], Yousuf A. Khan[2], Caroline M. R. Low[2], Kuba Perlin[2], Anna Potapenko[2], Pascal Savy[2], Sukhdeep Singh[2], Adrian Stecula[2], Ashok Thillaisundaram[2], Catherine Tong[2], Sergei Yakneen[2], Ellen D. Zhong[3,4], Michal Zielinski[2], Augustin Žídek[2], Victor Bapst[1,4], Pushmeet Kohli[1,4], Max Jaderberg[1,4,5], Demis Hassabis[1,4,6] ✉, John M. Jumper[1,4,6] ✉

The introduction of AlphaFold 2 has spurred a revolution in modelling the structure of proteins and their interactions, enabling a huge range of applications in protein modelling and design[1–6]. Here we describe our AlphaFold 3 model with a substantially updated diffusion-based architecture that is capable of predicting the joint structure of complexes including proteins, nucleic acids, small molecules, ions and modified residues. The new AlphaFold model demonstrates substantially improved accuracy

Probing the effects of broken symmetries in machine learning

Marcel F Langer, Sergey N Pozdnyakov and Michele Ceriotti*

Laboratory of Computational Science and Modeling and National Centre for Computational Design and Discovery of Novel Materials MARVEL, Institute of Materials, Ecole Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland

* Author to whom any correspondence should be addressed.

E-mail: michele.ceriotti@epfl.ch

Keywords: machine learning, symmetry-constrained models, atomistic modeling, molecular simulations

Supplementary material for this article is available online

Symmetry is one of the most central concepts in physics, and it is no surprise that it has also been widely adopted as an inductive bias for machine-learning models applied to the physical sciences. This is especially true for models targeting the properties of matter at the atomic scale. Both established and state-of-the-art approaches, with almost no exceptions, are built to be exactly equivariant to translations, permutations, and rotations of the atoms. Incorporating symmetries—rotations in particular—constrains the model design space and implies more complicated architectures that are often also computationally demanding. There are indications

The Importance of Being Scalable: Improving the Speed and Accuracy of Neural Network Interatomic Potentials Across Chemical Domains

Eric Qu
UC Berkeley
ericqu@berkeley.edu

Aditi S. Krishnapriyan
UC Berkeley, LBNL
aditik1@berkeley.edu

Abstract

Scaling has been a critical factor in improving model performance and generalization across various fields of machine learning. It involves how a model's performance changes with increases in model size or input data, as well as how efficiently computational resources are utilized to support this growth. Despite successes in scaling other types of machine learning models, the study of scaling in Neural Network Interatomic Potentials (NNIPs) remains limited. NNIPs act as

Swallowing the Bitter Pill: Simplified Scalable Conformer Generation

Yuyang Wang[1] Ahmed A. Elhag[1,2] Navdeep Jaitly[1] Joshua M. Susskind[1] Miguel Ángel Bautista[1]

Abstract

We present a novel way to predict molecular conformers through a simple formulation that sidesteps many of the heuristics of prior works and achieves state of the art results by using the advantages of scale. By training a diffusion generative model directly on 3D atomic positions without making assumptions about the explicit structure of molecules (e.g. modeling torsional angles) we are able to radically simplify [illegible] and make it easier to scale up the

is the vast complexity of the 3D structure space, encompassing factors such as bond lengths and torsional angles. Despite the molecular graph dictating potential 3D conformers through specific constraints, such as bond types and spatial arrangements determined by chiral centers, the conformational space experiences exponential growth with the expansion of the graph size and the number of rotatable bonds (Axelrod & Gomez-Bombarelli, 2022). This complicates brute force and exhaustive approaches, making them virtually unfeasible for even moderately small molecules.

Systematic methods like OMEGA (Hawkins et al., 2010)

Data augmentation

Article

Accurate structure prediction of biomolecular interactions with AlphaFold 3

https://doi.org/10.1038/s41586-024-07487-w

Received: 19 December 2023

Accepted: 29 April 2024

Published online: 8 May 2024

Open access

Check for updates

Josh Abramson[1,7], Jonas Adler[1,7], Jack Dunger[1,7], Richard Evans[1,7], Tim Green[1,7], Alexander Pritzel[1,7], Olaf Ronneberger[1,7], Lindsay Willmore[1,7], Andrew J. Ballard[1], Joshua Bambrick[1], Sebastian W. Bodenstein[1], David A. Evans[1], Chia-Chun Hung[1], Michael O'Neill[1], David Reiman[1], Kathryn Tunyasuvunakool[1], Zachary Wu[1], Akvilė Žemgulytė[1], Eirini Arvaniti[2], Charles Beattie[1], Ottavia Bertolli[2], Alex Bridgland[2], Alexey Cherepanov[3], Miles Congreve[4], Alexander I. Cowen-Rivers[5], Andrew Cowie[1], Michael Figurnov[1], Fabian B. Fuchs[1], Hannah Gladman[1], Rishub Jain[1], Yousuf A. Khan[2], Caroline M. R. Low[1], Kuba Perlin[1], Anna Potapenko[1], Pascal Savy[1], Sukhdeep Singh[2], Adrian Stecula[4], Ashok Thillaisundaram[1], Catherine Tong[1], Sergei Yakneen[6], Ellen D. Zhong[1,8], Michal Zielinski[1], Augustin Žídek[1], Victor Bapst[1,9], Pushmeet Kohli[1,9], Max Jaderberg[1,9,10], Demis Hassabis[1,9,11] ✉, John M. Jumper[1,9,11] ✉

The introduction of AlphaFold 2 has spurred a revolution in modelling the structure of proteins and their interactions, enabling a huge range of applications in protein modelling and design[1–6]. Here we describe our AlphaFold 3 model with a substantially updated diffusion-based architecture that is capable of predicting the joint structure of complexes including proteins, nucleic acids, small molecules, ions and modified residues. The new AlphaFold model demonstrates substantially improved accuracy

Probing the effects of broken symmetries in machine learning

Marcel F Langer, Sergey N Pozdnyakov and Michele Ceriotti ✉

Laboratory of Computational Science and Modeling and National Centre for Computational Design and Discovery of Novel Materials MARVEL, Institute of Materials, Ecole Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland

✉ Author to whom any correspondence should be addressed.

E-mail: michele.ceriotti@epfl.ch

Keywords: machine learning, symmetry-constrained models, atomistic modelling, molecular simulations

Supplementary material for this article is available online

Abstract

Symmetry is one of the most central concepts in physics, and it is no surprise that it has also been widely adopted as an inductive bias for machine-learning models applied to the physical sciences. This is especially true for models targeting the properties of matter at the atomic scale. Both established and state-of-the-art approaches, with almost no exceptions, are built to be exactly equivariant to translations, permutations, and rotations of the atoms. Incorporating symmetries—rotations in particular—constrains the model design space and implies more complicated architectures that are often also computationally demanding. There are indications

The Importance of Being Scalable: Improving the Speed and Accuracy of Neural Network Interatomic Potentials Across Chemical Domains

Eric Qu
UC Berkeley
ericqu@berkeley.edu

Aditi S. Krishnapriyan
UC Berkeley, LBNL
aditik1@berkeley.edu

Abstract

Scaling has been a critical factor in improving model performance and generalization across various fields of machine learning. It involves how a model's performance changes with increases in model size or input data, as well as how efficiently computational resources are utilized to support this growth. Despite successes in scaling other types of machine learning models, the study of scaling in Neural Network Interatomic Potentials (NNIPs) remains limited. NNIPs act as

Two for One: Diffusion Models and Force Fields for Coarse-Grained Molecular Dynamics

Marloes Arts,*,†,‡,∥ Victor Garcia Satorras,*,¶,# Chin-Wei Huang,¶ Daniel Zügner,§ Marco Federici,†,‡ Cecilia Clementi,⊥,△ Frank Noé,§ Robert Pinsler,# and Rianne van den Berg¶

†*Work done during an internship at Microsoft Research (Amsterdam).*
‡*University of Copenhagen, Department of Computer Science, Universitetsparken 1, Copenhagen, 2100, Denmark.*
¶*AI4Science, Microsoft Research, Evert van de Beekstraat 354, Amsterdam, 1118 CZ, The Netherlands.*
§*AI4Science, Microsoft Research, Karl-Liebknecht-Straße 32, Berlin, 10178, Germany.*
∥*University of Amsterdam, Informatics Institute, Science Park 904, Amsterdam, 1098 XH, The Netherlands.*
⊥*Freie Universität Berlin, Department of Physics, Arnimallle 12, Berlin, 14195, Germany.*
#*AI4Science, Microsoft Research, 21 Station Road, Cambridge, CB1 2FB, United Kingdom.*
△*Equal contribution.*

* E-mail: ma@di.ku.dk; victorga@microsoft.com

Abstract

Coarse-grained (CG) molecular dynamics enables the study of biological processes at temporal and spatial scales that would be intractable at an atomistic resolution. However, accurately learning a CG force field remains a challenge. In this work, we leverage connections between score-based generative models, force fields and molecular

Swallowing the Bitter Pill: Simplified Scalable Conformer Generation

Yuyang Wang[1] Ahmed A. Elhag[1,2] Navdeep Jaitly[1] Joshua M. Susskind[1] Miguel Ángel Bautista[1]

Abstract

We present a novel way to predict molecular conformers through a simple formulation that sidesteps many of the heuristics of prior works and achieves state of the art results by using the advantages of scale. By training a diffusion generative model directly on 3D atomic positions without making assumptions about the explicit structure of molecules (e.g. modeling torsional angles) we are able to radically simplify simu-

is the vast complexity of the 3D structure space, encompassing factors such as bond lengths and torsional angles. Despite the molecular graph dictating potential 3D conformers through specific constraints, such as bond types and spatial arrangements determined by chiral centers, the conformational space experiences exponential growth with the expansion of the graph size and the number of rotatable bonds (Axelrod & Gomez-Bombarelli, 2022). This complicates brute force and exhaustive approaches, making them virtually unfeasible for even moderately small molecules.

Systematic methods like RMSEA (Hawkins et al., 2010)

Data augmentation

Article

Accurate structure prediction of biomolecular interactions with AlphaFold 3

https://doi.org/10.1038/s41586-024-07487-w

Received: 19 December 2023

Accepted: 29 April 2024

Published online: 8 May 2024

Open access

Check for updates

Josh Abramson[1,7], Jonas Adler[1,7], Jack Dunger[1,7], Richard Evans[1,7], Tim Green[1,7], Alexander Pritzel[1,7], Olaf Ronneberger[1,7], Lindsay Willmore[1,7], Andrew J. Ballard[1], Joshua Bambrick[1], Sebastian W. Bodenstein[1], David A. Evans[1], Chia-Chun Hung[1], Michael O'Neill[1], David Reiman[1], Kathryn Tunyasuvunakool[1], Zachary Wu[1], Akvilė Žemgulytė[1], Eirini Arvaniti[2], Charles Beattie[2], Ottavia Bertolli[2], Alex Bridgland[2], Alexey Cherepanov[2], Miles Congreve[2], Alexander I. Cowen-Rivers[2], Andrew Cowie[2], Michael Figurnov[2], Fabian B. Fuchs[2], Hannah Gladman[2], Rishub Jain[2], Yousuf A. Khan[2], Caroline M. R. Low[2], Kuba Perlin[2], Anna Potapenko[2], Pascal Savy[2], Sukhdeep Singh[2], Adrian Stecula[2], Ashok Thillaisundaram[2], Catherine Tong[2], Sergei Yakneen[2], Ellen D. Zhong[2,8], Michal Zielinski[2], Augustin Žídek[2], Victor Bapst[1,9], Pushmeet Kohli[1,9], Max Jaderberg[1,9,10], Demis Hassabis[1,9,10] ✉, John M. Jumper[1,9,11] ✉

The introduction of AlphaFold 2 has spurred a revolution in modelling the structure of proteins and their interactions, enabling a huge range of applications in protein modelling and design[1–5]. Here we describe our AlphaFold 3 model with a substantially updated diffusion-based architecture that is capable of predicting the joint structure of complexes including proteins, nucleic acids, small molecules, ions and modified residues. The new AlphaFold model demonstrates substantially improved accuracy

The Importance of Being Scalable: Improving the Speed and Accuracy of Neural Network Interatomic Potentials Across Chemical Domains

Eric Qu
UC Berkeley
ericqu@berkeley.edu

Aditi S. Krishnapriyan
UC Berkeley, LBNL
aditik1@berkeley.edu

Abstract

Scaling has been a critical factor in improving model performance and generalization across various fields of machine learning. It involves how a model's performance changes with increases in model size or input data, as well as how efficiently computational resources are utilized to support this growth. Despite successes in scaling other types of machine learning models, the study of scaling in Neural Network Interatomic Potentials (NNIPs) remains limited. NNIPs act as

Two for One: Diffusion Models and Force Fields for Coarse-Grained Molecular Dynamics

Marloes Arts,*,†,‡,∥ Victor Garcia Satorras,*,¶,# Chin-Wei Huang,¶ Daniel Zügner,§ Marco Federici,†,‡ Cecilia Clementi,⊥,○ Frank Noé,§ Robert Pinsler,# and Rianne van den Berg¶

†*Work done during an internship at Microsoft Research (Amsterdam).*
‡*University of Copenhagen, Department of Computer Science, Universitetsparken 1, Copenhagen, 2100, Denmark.*
¶*AI4Science, Microsoft Research, Evert van de Beekstraat 354, Amsterdam, 1118 CZ, The Netherlands.*
§*AI4Science, Microsoft Research, Karl-Liebknecht-Straße 32, Berlin, 10178, Germany.*
∥*University of Amsterdam, Informatics Institute, Science Park 904, Amsterdam, 1098 XH, The Netherlands.*
⊥*Freie Universität Berlin, Department of Physics, Arnimallle 12, Berlin, 14195, Germany.*
#*AI4Science, Microsoft Research, 21 Station Road, Cambridge, CB1 2FB, United Kingdom.*
○*Equal contribution.*
* E-mail: ma@di.ku.dk; victorgar@microsoft.com

Abstract

Coarse-grained (CG) molecular dynamics enables the study of biological processes at temporal and spatial scales that would be intractable at an atomistic resolution. However, accurately learning a CG force field remains a challenge. In this work, we leverage connections between score-based generative models, force fields and molecular

Swallowing the Bitter Pill: Simplified Scalable Conformer Generation

Yuyang Wang[1] Ahmed A. Elhag[1,2] Navdeep Jaitly[1] Joshua M. Susskind[1] Miguel Ángel Bautista[1]

Abstract

We present a novel way to predict molecular conformers through a simple formulation that sidesteps many of the heuristics of prior works and achieves state of the art results by using the advantages of scale. By training a diffusion generative model directly on 3D atomic positions without making assumptions about the explicit structure of molecules (e.g. modeling torsional angles) we are able to radically simplify [illegible] and make it easier to scale up the

Probing the effects of broken symmetries in machine learning

Marcel F Langer ✉, **Sergey N Pozdnyakov** and **Michele Ceriotti** ✉

Laboratory of Computational Science and Modeling and National Centre for Computational Design and Discovery of Novel Materials MARVEL, Institute of Materials, Ecole Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland
✉ Author to whom any correspondence should be addressed.
E-mail: michele.ceriotti@epfl.ch

Keywords: machine learning, symmetry-constrained models, atomistic modelling, molecular simulations

Supplementary material for this article is available online

Abstract

Symmetry is one of the most central concepts in physics, and it is no surprise that it has also been widely adopted as an inductive bias for machine-learning models applied to the physical sciences. This is especially true for models targeting the properties of matter at the atomic scale. Both established and state-of-the-art approaches, with almost no exceptions, are built to be exactly equivariant to translations, permutations, and rotations of the atoms. Incorporating symmetries—rotations in particular—constrains the model design space and implies more complicated architectures that are often also computationally demanding. There are indications

Does equivariance matter at scale?

Johann Brehmer * **Sönke Behrends** **Pim de Haan** * **Taco Cohen** *
Qualcomm AI Research†
hail@johannbrehmer.de

Abstract

Given large data sets and sufficient compute, is it beneficial to design neural architectures for the structure and symmetries of each problem? Or is it more efficient to learn them from data? We study empirically how equivariant and non-equivariant networks scale with compute and training samples. Focusing on a benchmark problem of rigid-body interactions and on general-purpose transformer architectures, we perform a series of experiments, varying the model size, training steps, and dataset size. We find evidence for three conclusions. First, equivariance improves data efficiency, but training non-equivariant models with data augmentation can close this gap given sufficient epochs. Second, scaling with compute follows a power law, with equivariant models outperforming non-equivariant ones at each tested compute budget. Finally, the optimal allocation of a compute budget onto model size and training duration differs between equivariant and non-equivariant models.

Data augmentation

- Easy to implement
- No specialized architecture necessary

Data augmentation

- Easy to implement
- No specialized architecture necessary
- No exact equivariance

Data augmentation

- Easy to implement
- No specialized architecture necessary
- No exact equivariance

Can we understand data augmentation theoretically?

Empirical NTK

Training dynamics under continuous gradient descent:

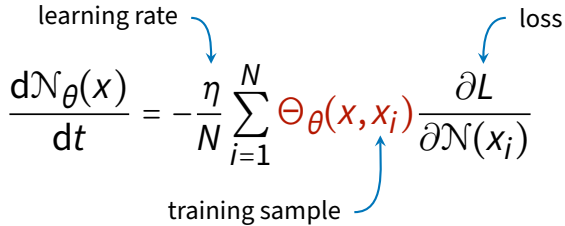

Empirical NTK

Training dynamics under continuous gradient descent:

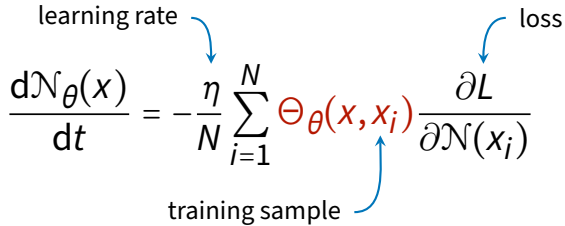

with the empirical neural tangent kernel (NTK)

$$\Theta_\theta(x,x') = \sum_\mu \frac{\partial \mathcal{N}(x)}{\partial \theta_\mu}\frac{\partial \mathcal{N}(x')}{\partial \theta_\mu}$$

Infinite width limit

[Jacot et al. 2018]

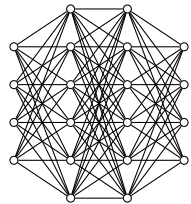

Infinite width limit

[Jacot et al. 2018]

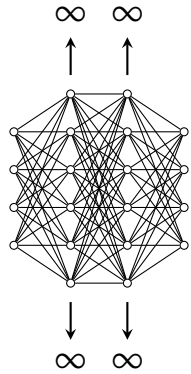

Infinite width limit

[Jacot et al. 2018]

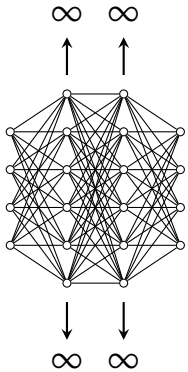

- NTK becomes independent of initialization

Infinite width limit

[Jacot et al. 2018]

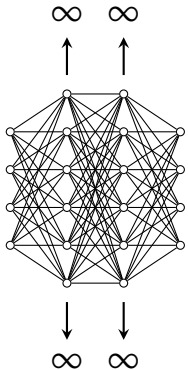

- NTK becomes independent of initialization
- NTK becomes constant in training

Infinite width limit

[Jacot et al. 2018]

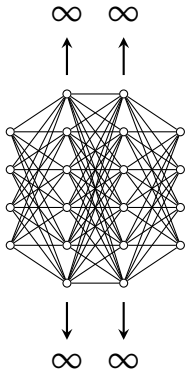

- NTK becomes independent of initialization
- NTK becomes constant in training
- NTK can be computed for most networks

Infinite width limit

[Jacot et al. 2018]

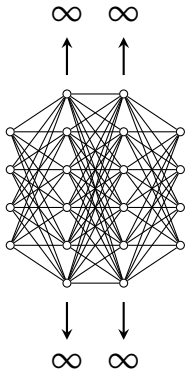

- 👍 NTK becomes independent of initialization
- 👍 NTK becomes constant in training
- 👍 NTK can be computed for most networks
- ✓ Training dynamics can be solved

Mean prediction from NTK

[Jacot et al. 2018]

ⓘ At infinite width, the mean prediction is given by

$$\mu_t(x) = \Theta(x, X)\Theta(X, X)^{-1}(\mathbb{I} - e^{-\eta\Theta(X,X)t})Y$$

Mean prediction from NTK

[Jacot et al. 2018]

ⓘ At infinite width, the mean prediction is given by

neural tangent kernel

$$\mu_t(x) = \Theta(x, X)\Theta(X, X)^{-1}(\mathbb{I} - e^{-\eta\Theta(X,X)t})Y$$

Mean prediction from NTK

[Jacot et al. 2018]

ⓘ At infinite width, the mean prediction is given by

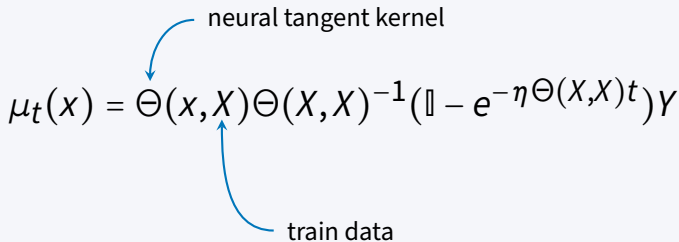

Mean prediction from NTK

[Jacot et al. 2018]

ⓘ At infinite width, the mean prediction is given by

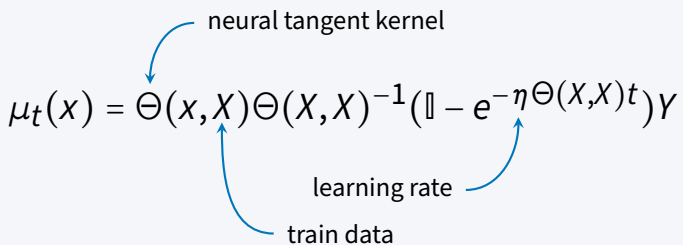

Mean prediction from NTK

[Jacot et al. 2018]

ⓘ At infinite width, the mean prediction is given by

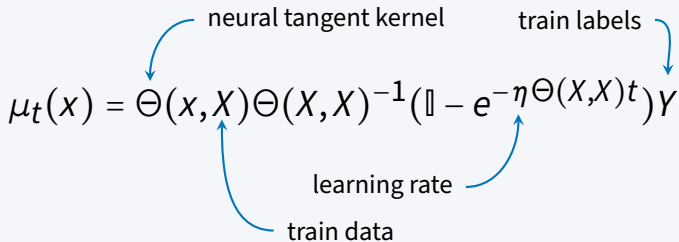

Data augmentation

Data augmentation at infinite width

$$\mu_t(x) = \Theta(x, X)\Theta(X, X)^{-1}(\mathbb{I} - e^{-\eta\Theta(X,X)t})Y$$

Data augmentation at infinite width

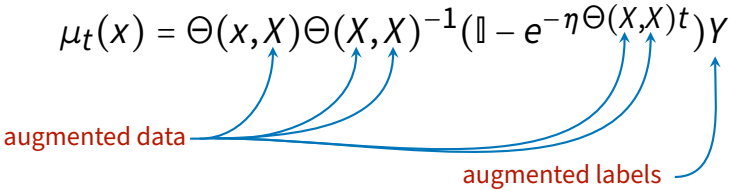

Data augmentation at infinite width

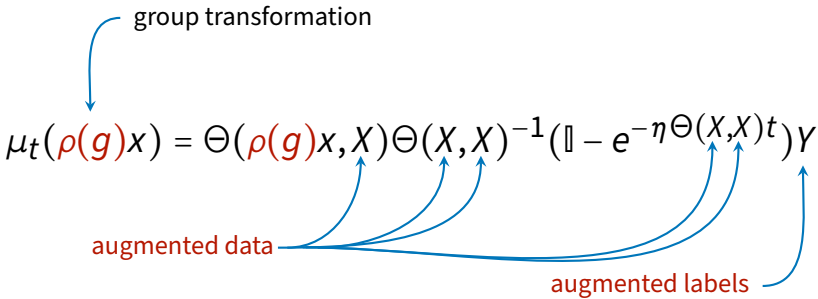

Data augmentation at infinite width

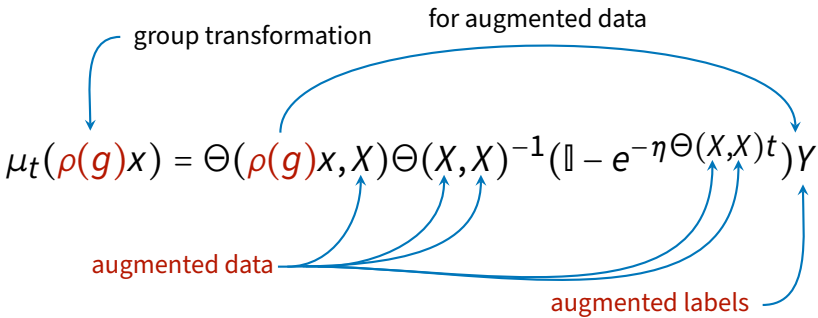

Data augmentation at infinite width

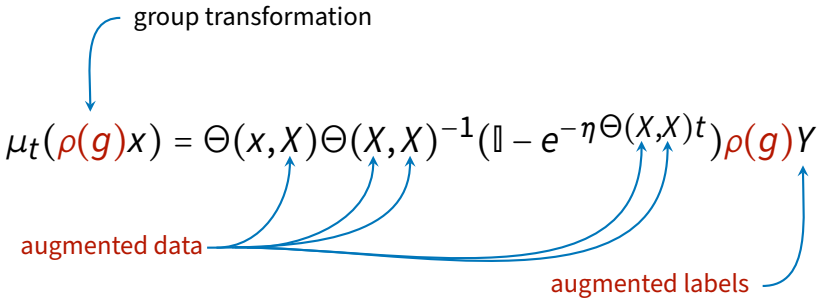

Data augmentation at infinite width

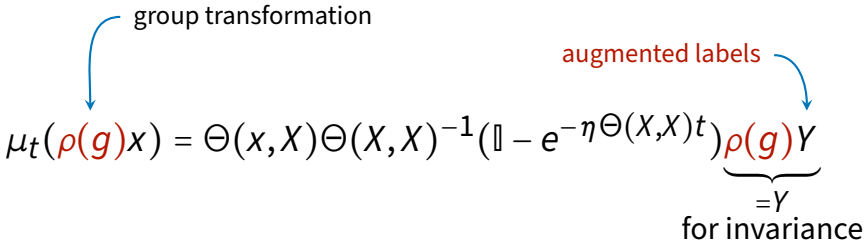

Data augmentation at infinite width

group transformation

$$\mu_t(\rho(g)x) = \Theta(x,X)\Theta(X,X)^{-1}(\mathbb{I} - e^{-\eta\Theta(X,X)t})\underbrace{\rho(g)Y}_{=Y}$$
$$= \mu_t(x)$$

for invariance

Mean prediction

$$\mu_t(x)$$

Mean prediction

$$\mu_t(x) = \mathbb{E}_{\theta_0 \sim \text{initializations}}[\mathcal{N}_{\theta_t}(x)]$$

Mean prediction

$$\mu_t(x) = \mathbb{E}_{\theta_0 \sim \text{initializations}}[\mathcal{N}_{\theta_t}(x)] = \lim_{n \to \infty} \frac{1}{n} \sum_{\theta_0 = \text{init}_1}^{\text{init}_n} \mathcal{N}_{\theta_t}(x)$$

Mean prediction

$$\mu_t(x) = \mathbb{E}_{\theta_0 \sim \text{initializations}}[\mathcal{N}_{\theta_t}(x)] = \underbrace{\lim_{n \to \infty} \frac{1}{n} \sum_{\theta_0 = \text{init}_1}^{\text{init}_n} \mathcal{N}_{\theta_t}(x)}_{\text{mean prediction of deep ensemble}}$$

Main conclusion

Deep ensembles trained with data augmentation are equivariant.

Main conclusion

Deep ensembles trained with data augmentation are equivariant.

- ✓ Proof of exact equivariance for
  - full data augmentation
  - infinite ensembles

Main conclusion

Deep ensembles trained with data augmentation are equivariant.

- ✓ Proof of exact equivariance for
  - full data augmentation
  - infinite ensembles
- ✓ Equivariance holds for all training times

Main conclusion

Deep ensembles trained with data augmentation are equivariant.

- ✓ Proof of exact equivariance for
  - full data augmentation
  - infinite ensembles
- ✓ Equivariance holds for all training times
- ✓ Equivariance holds away from the training data

Main conclusion

Deep ensembles trained with data augmentation are equivariant.

- ✓ Proof of exact equivariance for
  - full data augmentation
  - infinite ensembles
- ✓ Equivariance holds for all training times
- ✓ Equivariance holds away from the training data
- ✓ Holds also for finite-width networks

[Nordenfors, Flinth 2024]

Intuitive explanation

- ✓ Equivariance holds for all training times
- ✓ Equivariance holds away from the training data

Intuitive explanation

- ✓ Equivariance holds for all training times
- ✓ Equivariance holds away from the training data

ⓘ At infinite width, the mean output at initialization is zero everywhere.

Intuitive explanation

- ✓ Equivariance holds for all training times
- ✓ Equivariance holds away from the training data

⊙ At infinite width, the mean output at initialization is zero everywhere.

⇨ Training with full data augmentation leads to an equivariant function.

Toy example

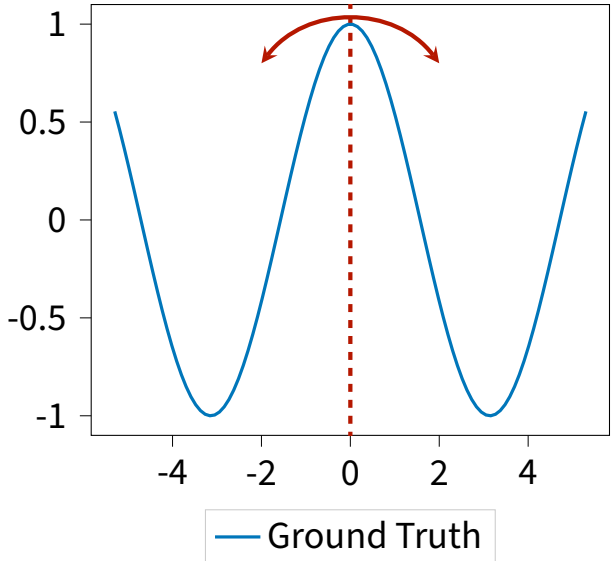
1
0.5
0
-0.5
-1
-4
-2
0
2
4
Ground Truth

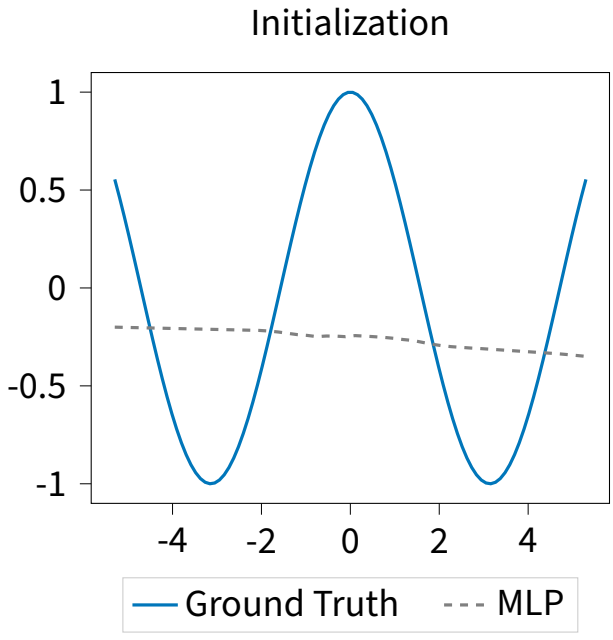
Initialization
1
0.5
0
-0.5
-1
-4
-2
0
2
4
Ground Truth
MLP

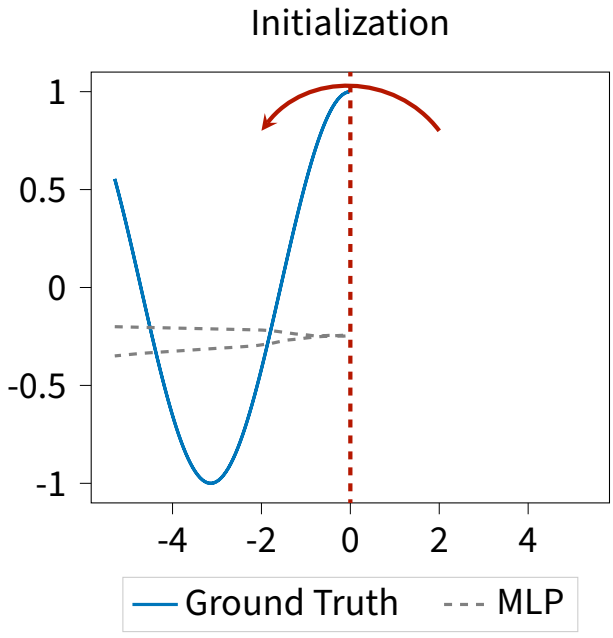
Initialization
1
0.5
0
-0.5
-1
-4
-2
0
2
4
Ground Truth
MLP

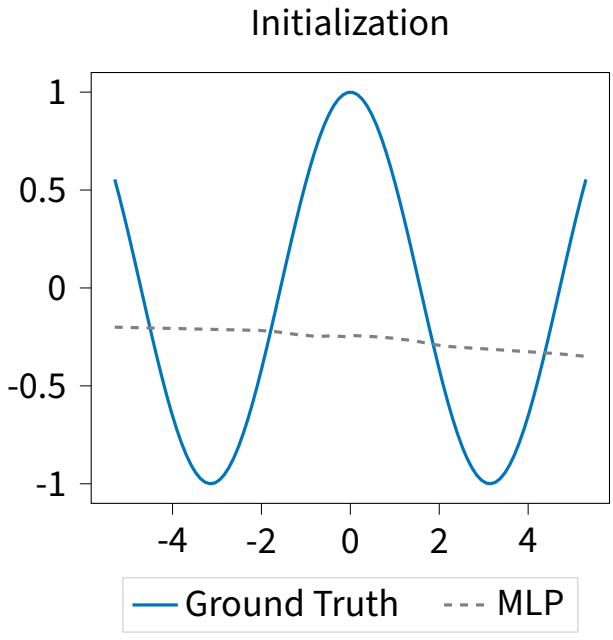
Initialization
1
0.5
0
-0.5
-1
-4
-2
0
2
4
Ground Truth
MLP

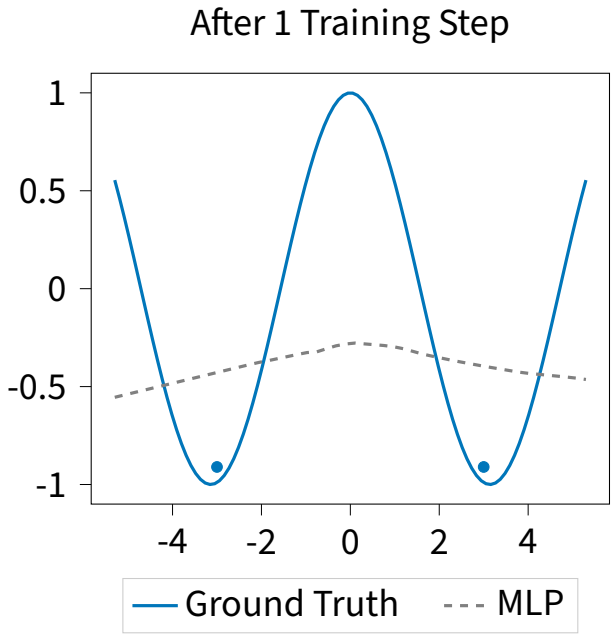
After 1 Training Step
1
0.5
0
-0.5
-1
-4
-2
0
2
4
Ground Truth
MLP

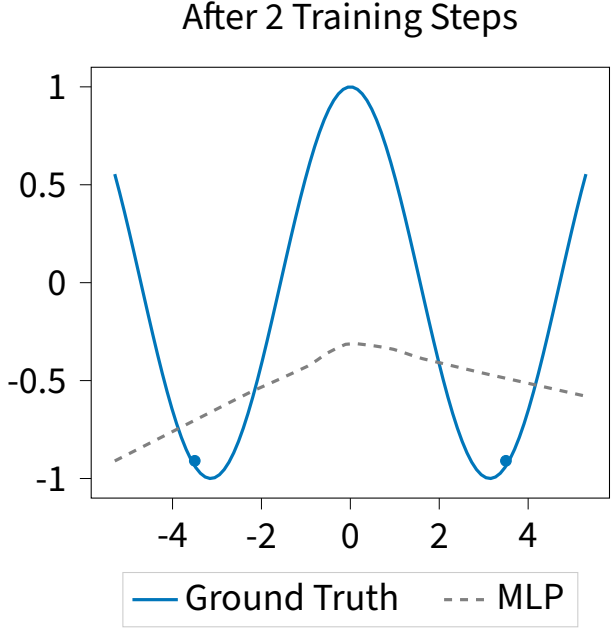
After 2 Training Steps
1
0.5
0
-0.5
-1
-4
-2
0
2
4
Ground Truth
MLP

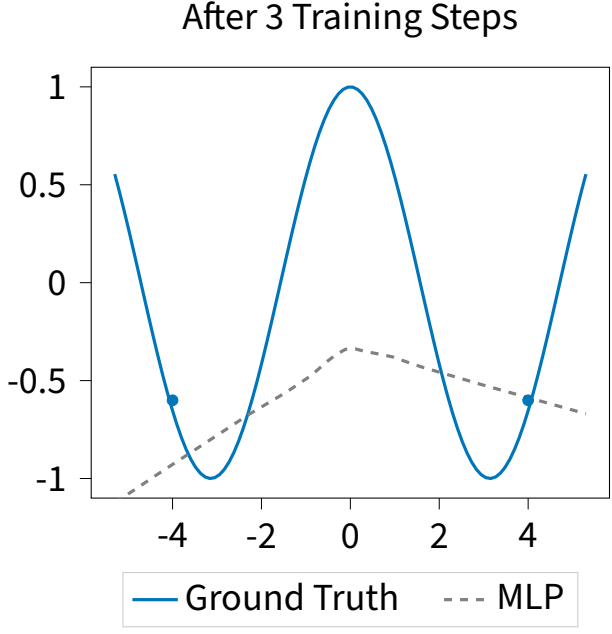
After 3 Training Steps
1
0.5
0
-0.5
-1
-4
-2
0
2
4
Ground Truth
MLP

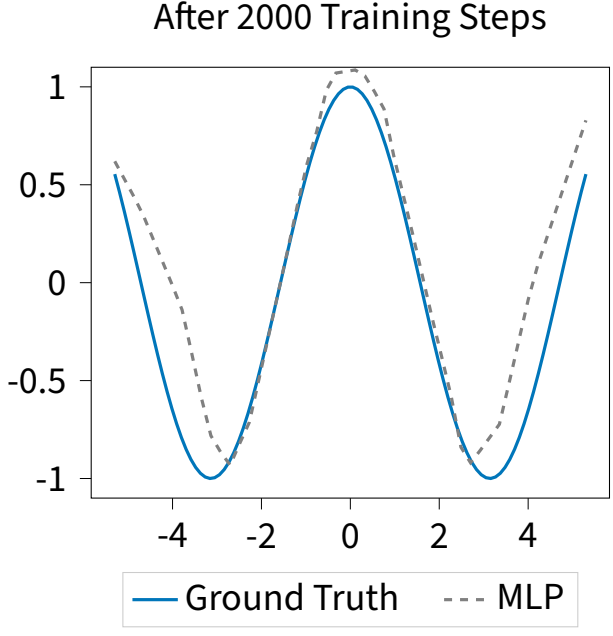
After 2000 Training Steps
1
0.5
0
-0.5
-1
-4
-2
0
2
4
Ground Truth
MLP

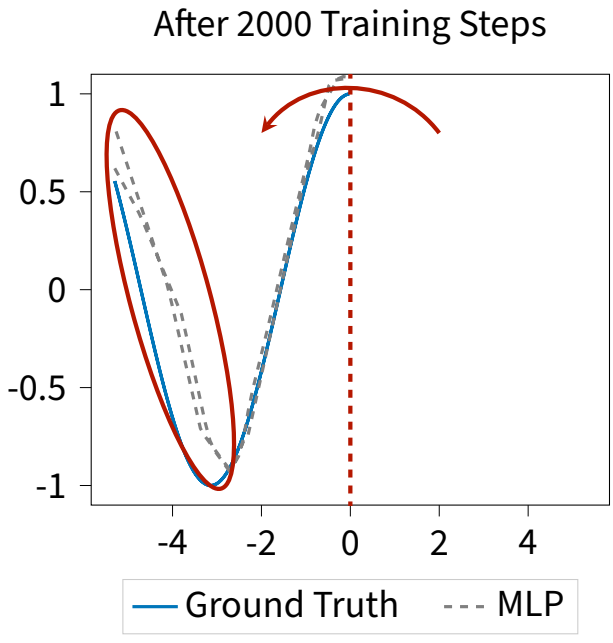
After 2000 Training Steps
1
0.5
0
-0.5
-1
-4
-2
0
2
4
Ground Truth
MLP

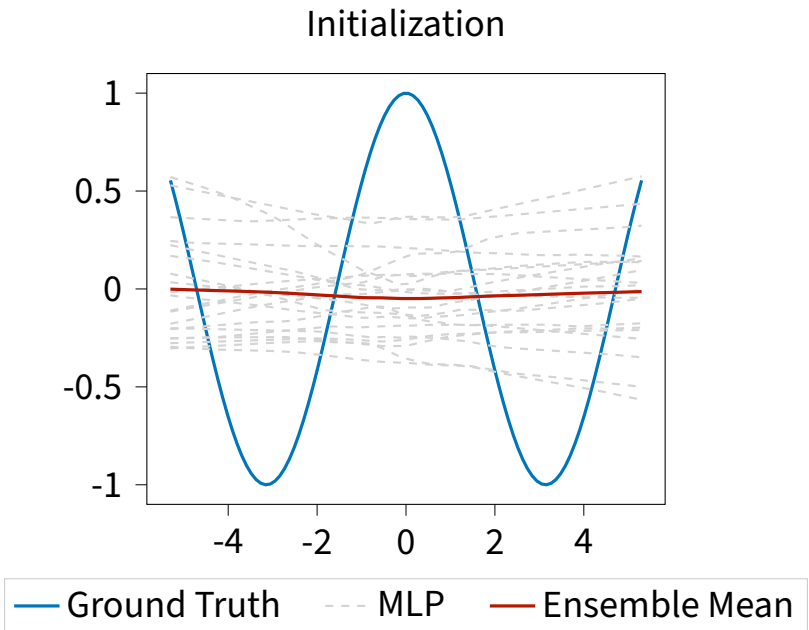
Initialization
1
0.5
0
-0.5
-1
-4
-2
0
2
4
Ground Truth
MLP
Ensemble Mean

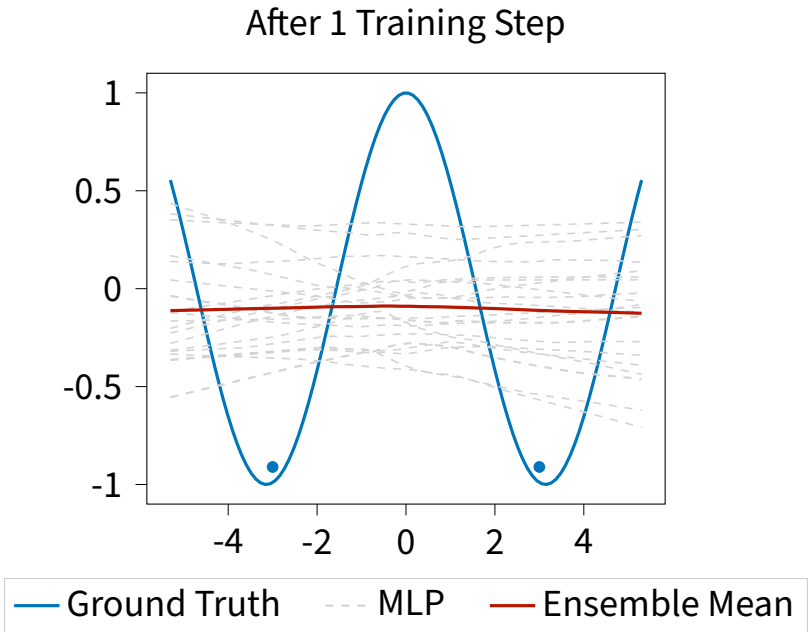
After 1 Training Step
1
0.5
0
-0.5
-1
-4
-2
0
2
4
Ground Truth
MLP
Ensemble Mean

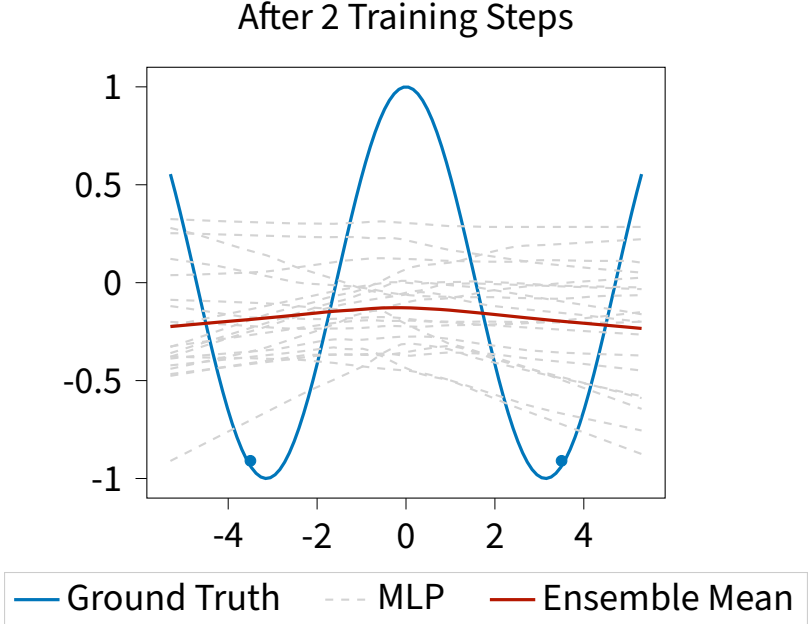
After 2 Training Steps
1
0.5
0
-0.5
-1
-4
-2
0
2
4
Ground Truth
MLP
Ensemble Mean

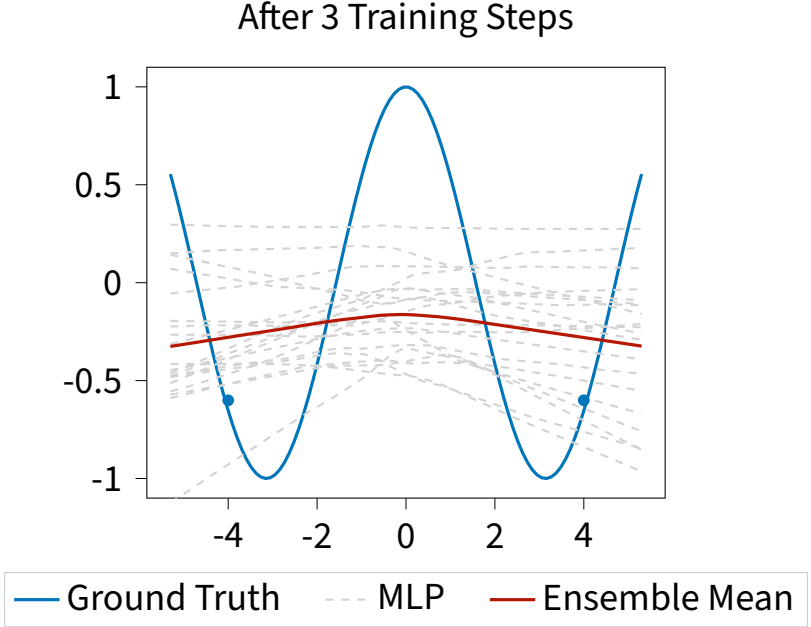
After 3 Training Steps
1
0.5
0
-0.5
-1
-4
-2
0
2
4
Ground Truth
MLP
Ensemble Mean

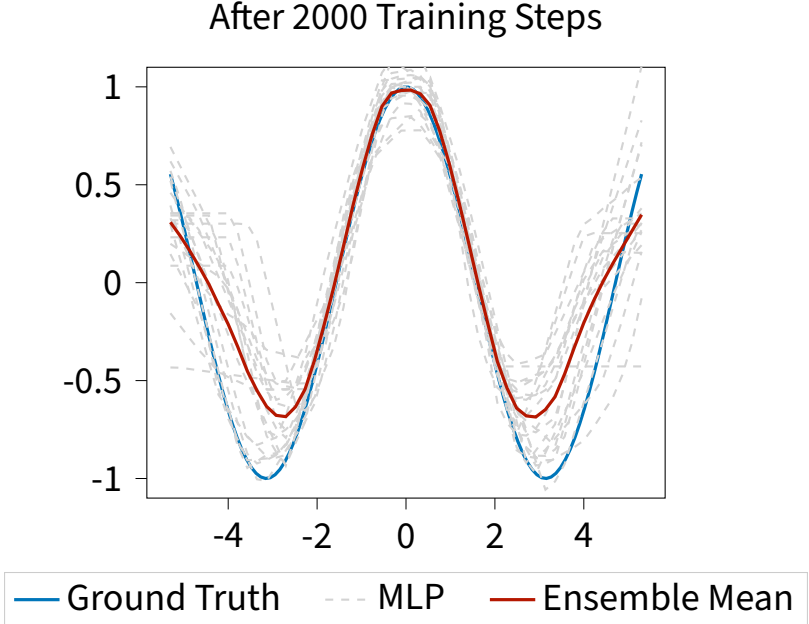
After 2000 Training Steps
1
0.5
0
-0.5
-1
-4
-2
0
2
4
Ground Truth
MLP
Ensemble Mean

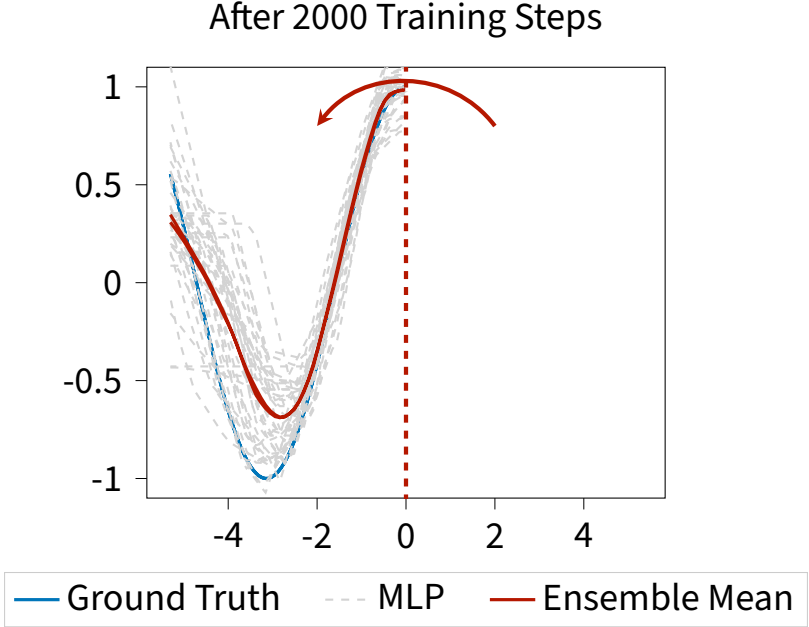
After 2000 Training Steps
1
0.5
0
-0.5
-1
-4
-2
0
2
4
Ground Truth
MLP
Ensemble Mean

What Does An Augmented Ensemble Converge To?

Rotating images

Rotating images

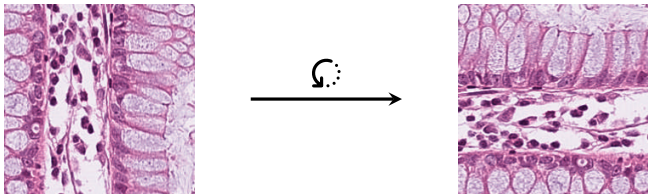

Rotating images

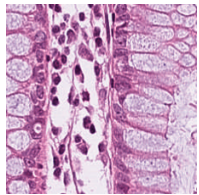

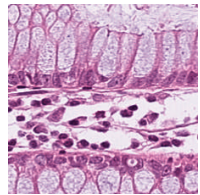

$f(x)$

f : pixels $\rightarrow$ colors

Rotating images

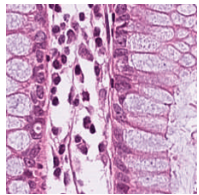

$f(x)$

f : pixels $\to$ colors

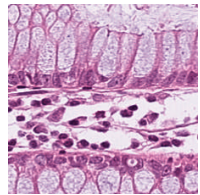

$f(\rho(g^{-1})x)$

$= [\rho_{\text{reg}}(g)f](x)$

Data augmentation and NTKs

Data augmentation and NTKs

Consider two ensembles:

trained without data augmentation

trained with data augmentation

Data augmentation and NTKs

Consider two ensembles:

trained without data augmentation

trained with data augmentation

If

$$\Theta^{\text{non-aug}}(f, f') = \frac{1}{|G|} \sum_{g \in G} \Theta^{\text{aug}}(f, \rho_{\text{reg}}(g) f')$$

Data augmentation and NTKs

Consider two ensembles:

trained without data augmentation

trained with data augmentation

If

$$\Theta^{\text{non-aug}}(f, f') = \frac{1}{|G|} \sum_{g \in G} \Theta^{\text{aug}}(f, \rho_{\text{reg}}(g) f')$$

Then

$$\mu_t^{\text{non-aug}}(x) = \mu_t^{\text{aug}}(x)$$

at infinite width.

Data augmentation and NTKs

Consider two ensembles:

trained without data augmentation

trained with data augmentation

If

$$\Theta^{\text{non-aug}}(f, f') = \frac{1}{|G|} \sum_{g \in G} \Theta^{\text{aug}}(f, \rho_{\text{reg}}(g) f')$$

Then

$$\mu_t^{\text{non-aug}}(x) = \mu_t^{\text{aug}}(x) \quad \forall t$$

at infinite width.

Data augmentation and NTKs

Consider two ensembles:

trained without data augmentation

trained with data augmentation

If

$$\Theta^{\text{non-aug}}(f, f') = \frac{1}{|G|} \sum_{g \in G} \Theta^{\text{aug}}(f, \rho_{\text{reg}}(g) f')$$

Then

$$\mu_t^{\text{non-aug}}(x) = \mu_t^{\text{aug}}(x) \quad \forall t \quad \forall x$$

at infinite width.

Data augmentation and NTKs

$$\Theta^{\text{non-aug}}(f, f') = \frac{1}{|G|} \sum_{g \in G} \Theta^{\text{aug}}(f, \rho_{\text{reg}}(g) f')$$

Data augmentation and NTKs

$$\Theta^{\text{non-aug}}(f, f') = \frac{1}{|G|} \sum_{g \in G} \Theta^{\text{aug}}(f, \rho_{\text{reg}}(g) f')$$

- Given an architecture with NTK Θ^{aug}, find an architecture with NTK $\Theta^{\text{non-aug}}$

Group convolutions

[Cohen, Welling 2016]

Group convolutions

[Cohen, Welling 2016]

Group conv's are the (unique) linear layers equivariant wrt ρ_{reg}

Group convolutions

[Cohen, Welling 2016]

Group conv's are the (unique) linear layers equivariant wrt ρ_{reg}

- Ordinary convolutions

$$f'(y) = \int_X \mathrm{d}x \, \kappa(x - y) \, f(x)$$

Group convolutions

[Cohen, Welling 2016]

Group conv's are the (unique) linear layers equivariant wrt ρ_{reg}

- Ordinary convolutions

$$f'(y) = \int_X \mathrm{d}x\, \kappa(x - y)\, f(x)$$

- Group convolutions

$$f'(g) = \int_X \mathrm{d}x\, \kappa(\rho(g^{-1})x)\, f(x) \qquad \text{lifting}$$

Group convolutions

[Cohen, Welling 2016]

Group conv's are the (unique) linear layers equivariant wrt ρ_{reg}

- Ordinary convolutions

$$f'(y) = \int_X \mathrm{d}x\, \kappa(x - y)\, f(x)$$

- Group convolutions

$$f'(g) = \int_X \mathrm{d}x\, \kappa(\rho(g^{-1})x)\, f(x) \qquad \text{lifting}$$

$$f'(g) = \int_G \mathrm{d}g\, \kappa(g^{-1}h)\, f(h) \qquad \text{group convolution}$$

Group convolutions

[Cohen, Welling 2016]

Group conv's are the (unique) linear layers equivariant wrt ρ_{reg}

- Ordinary convolutions

$$f'(y) = \int_X \mathrm{d}x\, \kappa(x - y)\, f(x)$$

- Group convolutions

$$f'(g) = \int_X \mathrm{d}x\, \kappa(\rho(g^{-1})x)\, f(x) \qquad \text{lifting}$$

$$f'(g) = \int_G \mathrm{d}g\, \kappa(g^{-1}h)\, f(h) \qquad \text{group convolution}$$

$$f' = \frac{1}{\mathrm{vol}(G)} \int_G \mathrm{d}g\, f(g) \qquad \text{group pooling}$$

GCNNs

Stack GConv-layers to obtain an invariant network

GCNNs

Stack GConv-layers to obtain an invariant network

GCNNs

Stack GConv-layers to obtain an invariant network

lifting

GCNNs

Stack GConv-layers to obtain an invariant network

GCNNs

Stack GConv-layers to obtain an invariant network

GCNNs

Stack GConv-layers to obtain an invariant network

$\longrightarrow$ lifting $\longrightarrow$ σ $\longrightarrow$ GConv $\longrightarrow$ σ

GCNNs

Stack GConv-layers to obtain an invariant network

$\longrightarrow$ lifting $\longrightarrow$ σ $\longrightarrow$ GConv $\longrightarrow$ σ $\longrightarrow$ GConv

GCNNs

Stack GConv-layers to obtain an invariant network

lifting → σ → GConv → σ → GConv → GPool

NTKs for GCNNs

Stack GConv-layers to obtain an invariant network

Compute NTK with layer-wise recursion

[Misof, Kessel, JG, 2024]

NTKs for GCNNs

Stack GConv-layers to obtain an invariant network

lifting → σ → GConv → σ → GConv → GPool

Compute NTK with layer-wise recursion

[Misof, Kessel, JG, 2024]

0

NTKs for GCNNs

Stack GConv-layers to obtain an invariant network

$\longrightarrow$ lifting $\longrightarrow$ σ $\longrightarrow$ GConv $\longrightarrow$ σ $\longrightarrow$ GConv $\longrightarrow$ GPool

Compute NTK with layer-wise recursion

[Misof, Kessel, JG, 2024]

$0 \longrightarrow \Theta^{(1)}_{g,g'}(f,f')$

NTKs for GCNNs

Stack GConv-layers to obtain an invariant network

$\longrightarrow$ lifting $\longrightarrow$ σ $\longrightarrow$ GConv $\longrightarrow$ σ $\longrightarrow$ GConv $\longrightarrow$ GPool

Compute NTK with layer-wise recursion

[Misof, Kessel, JG, 2024]

$$0 \longrightarrow \Theta^{(1)}_{g,g'}(f,f') \longrightarrow \Theta^{(2)}_{g,g'}(f,f')$$

NTKs for GCNNs

Stack GConv-layers to obtain an invariant network

$\longrightarrow$ lifting $\longrightarrow$ σ $\longrightarrow$ GConv $\longrightarrow$ σ $\longrightarrow$ GConv $\longrightarrow$ GPool

Compute NTK with layer-wise recursion

[Misof, Kessel, JG, 2024]

$$0 \longrightarrow \Theta^{(1)}_{g,g'}(f,f') \longrightarrow \Theta^{(2)}_{g,g'}(f,f') \longrightarrow \Theta^{(3)}_{g,g'}(f,f')$$

NTKs for GCNNs

Stack GConv-layers to obtain an invariant network

$\longrightarrow$ lifting $\longrightarrow$ σ $\longrightarrow$ GConv $\longrightarrow$ σ $\longrightarrow$ GConv $\longrightarrow$ GPool

Compute NTK with layer-wise recursion

[Misof, Kessel, JG, 2024]

$$0 \longrightarrow \Theta^{(1)}_{g,g'}(f,f') \longrightarrow \Theta^{(2)}_{g,g'}(f,f') \longrightarrow \Theta^{(3)}_{g,g'}(f,f') \longrightarrow \Theta^{(4)}_{g,g'}(f,f')$$

NTKs for GCNNs

Stack GConv-layers to obtain an invariant network

$\longrightarrow$ lifting $\longrightarrow$ σ $\longrightarrow$ GConv $\longrightarrow$ σ $\longrightarrow$ GConv $\longrightarrow$ GPool

Compute NTK with layer-wise recursion

[Misof, Kessel, JG, 2024]

$$0 \longrightarrow \Theta^{(1)}_{g,g'}(f,f') \longrightarrow \Theta^{(2)}_{g,g'}(f,f') \longrightarrow \Theta^{(3)}_{g,g'}(f,f') \longrightarrow \Theta^{(4)}_{g,g'}(f,f') \longrightarrow \Theta^{(5)}_{g,g'}(f,f')$$

NTKs for GCNNs

Stack GConv-layers to obtain an invariant network

$\longrightarrow$ lifting $\longrightarrow$ σ $\longrightarrow$ GConv $\longrightarrow$ σ $\longrightarrow$ GConv $\longrightarrow$ GPool

Compute NTK with layer-wise recursion

[Misof, Kessel, JG, 2024]

$$0 \longrightarrow \Theta^{(1)}_{g,g'}(f,f') \longrightarrow \Theta^{(2)}_{g,g'}(f,f') \longrightarrow \Theta^{(3)}_{g,g'}(f,f') \longrightarrow \Theta^{(4)}_{g,g'}(f,f') \longrightarrow \Theta^{(5)}_{g,g'}(f,f') \longrightarrow \Theta(f,f')$$

NTKs of MLPs and GCNNs

NTKs of MLPs and GCNNs

- Consider two neural networks

NTKs of MLPs and GCNNs

- Consider two neural networks

  An MLP

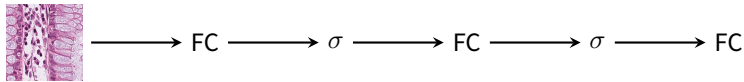

NTKs of MLPs and GCNNs

- Consider two neural networks

An MLP

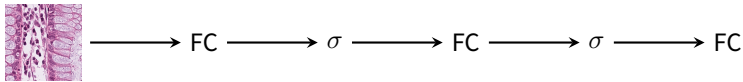

A GCNN

NTKs of MLPs and GCNNs

- Consider two neural networks

An MLP

FC → σ → FC → σ → FC

A GCNN

→ lifting

NTKs of MLPs and GCNNs

- Consider two neural networks

An MLP

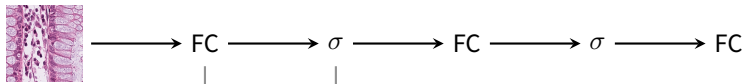

A GCNN

NTKs of MLPs and GCNNs

- Consider two neural networks

An MLP

FC → σ → FC → σ → FC

A GCNN

lifting → σ → GConv

NTKs of MLPs and GCNNs

- Consider two neural networks

An MLP

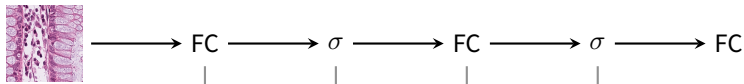

A GCNN

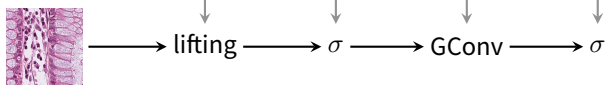

NTKs of MLPs and GCNNs

- Consider two neural networks

An MLP

FC → σ → FC → σ → FC

A GCNN

lifting → σ → GConv → σ → GConv

NTKs of MLPs and GCNNs

- Consider two neural networks

An MLP

FC → σ → FC → σ → FC

A GCNN

lifting → σ → GConv → σ → GConv → GPool

NTKs of MLPs and GCNNs

- Consider two neural networks

An MLP

FC → σ → FC → σ → FC

A GCNN

lifting → σ → GConv → σ → GConv → GPool

- Then

$$\Theta^{\mathrm{GCNN}}(f, f') = \frac{1}{|G|} \sum_{g \in G} \Theta^{\mathrm{MLP}}(f, \rho_{\mathrm{reg}}(g) f')$$

Data augmentation of MLPs

$$\Theta^{\text{GCNN}}(f, f') = \frac{1}{|G|} \sum_{g \in G} \Theta^{\text{MLP}}(f, \rho_{\text{reg}}(g) f')$$

Data augmentation of MLPs

before: non-aug

$$\Theta^{\text{GCNN}}(f, f') = \frac{1}{|G|} \sum_{g \in G} \Theta^{\text{MLP}}(f, \rho_{\text{reg}}(g) f')$$

Data augmentation of MLPs

before: non-aug

before: aug

$$\Theta^{\text{GCNN}}(f, f') = \frac{1}{|G|} \sum_{g \in G} \Theta^{\text{MLP}}(f, \rho_{\text{reg}}(g) f')$$

Data augmentation of MLPs

before: non-aug

before: aug

$$\Theta^{\text{GCNN}}(f, f') = \frac{1}{|G|} \sum_{g \in G} \Theta^{\text{MLP}}(f, \rho_{\text{reg}}(g) f')$$

⇨ training the MLP on G-augmented data

Data augmentation of MLPs

before: non-aug

before: aug

$$\Theta^{\text{GCNN}}(f, f') = \frac{1}{|G|} \sum_{g \in G} \Theta^{\text{MLP}}(f, \rho_{\text{reg}}(g) f')$$

⇨ training the MLP on G-augmented data = training the GCNN on unaugmented data

Data augmentation of MLPs

before: non-aug

before: aug

$$\Theta^{\mathrm{GCNN}}(f, f') = \frac{1}{|G|} \sum_{g \in G} \Theta^{\mathrm{MLP}}(f, \rho_{\mathrm{reg}}(g) f')$$

⇨ training the MLP on G-augmented data = training the GCNN on unaugmented data

in the ensemble mean

Data augmentation of MLPs

before: non-aug

before: aug

$$\Theta^{\text{GCNN}}(f, f') = \frac{1}{|G|} \sum_{g \in G} \Theta^{\text{MLP}}(f, \rho_{\text{reg}}(g) f')$$

⇨ training the MLP on G-augmented data = training the GCNN on unaugmented data

in the ensemble mean, $\forall t$, $\forall x$

Data augmentation of CNNs

Data augmentation of CNNs

- Consider a CNN

Data augmentation of CNNs

- Consider a CNN and a GCNN invariant wrt. roto-translations

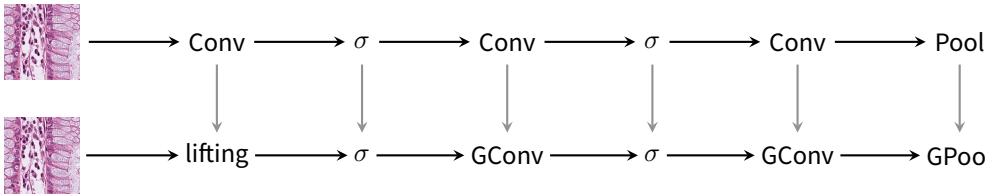

Data augmentation of CNNs

- Consider a CNN and a GCNN invariant wrt. roto-translations

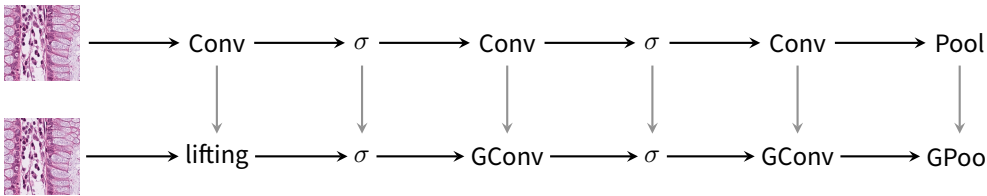

- Then

$$\Theta^{\text{GCNN}}(f, f') = \frac{1}{n} \sum_{r \in C_n} \Theta^{\text{CNN}}(f, \rho_{\text{reg}}(r) f')$$

Data augmentation of CNNs

- Consider a CNN and a GCNN invariant wrt. roto-translations

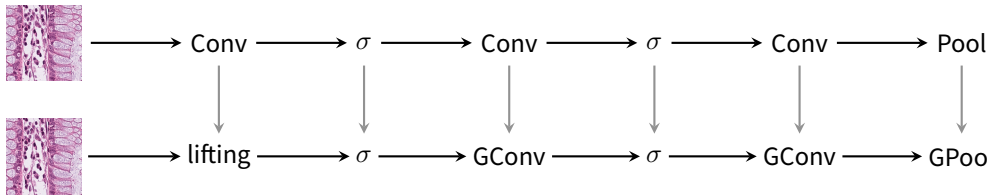

- Then

$$\Theta^{\text{GCNN}}(f, f') = \frac{1}{n} \sum_{r \in C_n} \Theta^{\text{CNN}}(f, \rho_{\text{reg}}(r) f')$$

⇨ By training the CNN on rotated images, one obtains a roto-translation invariant GCNN

Experiments

Ising model

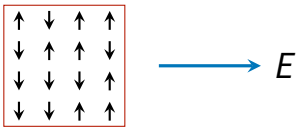

Ising model

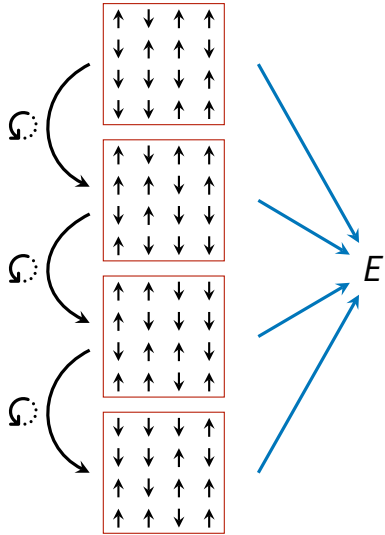

Ising model

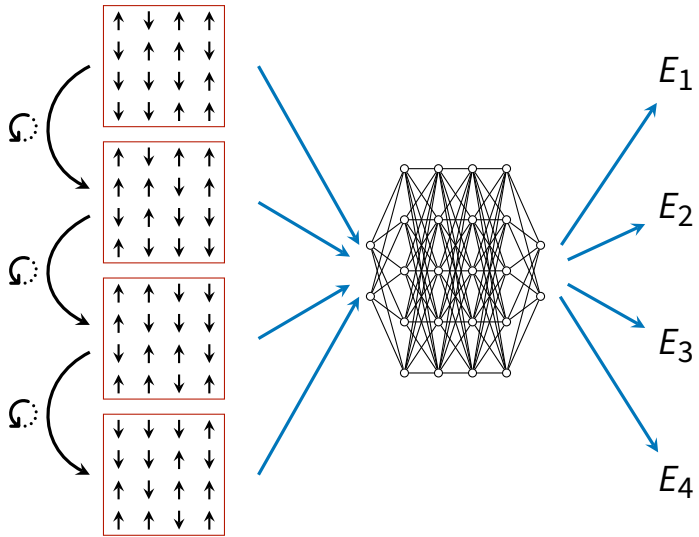

Ising model

Relative Standard Deviation

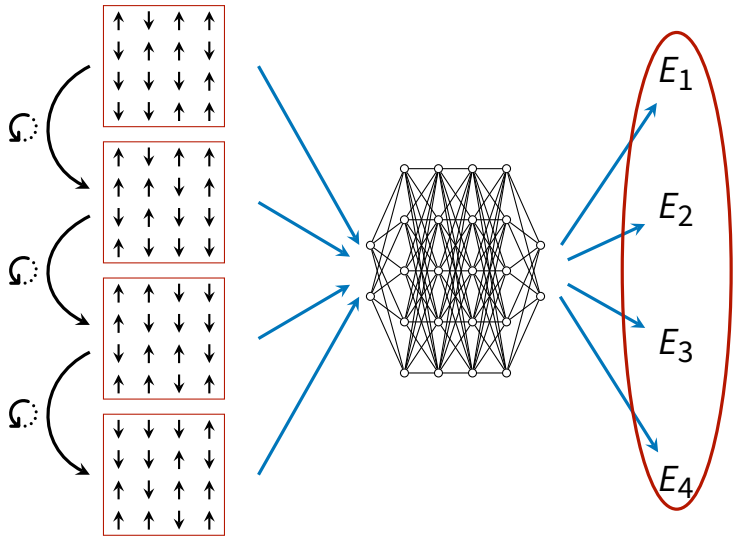

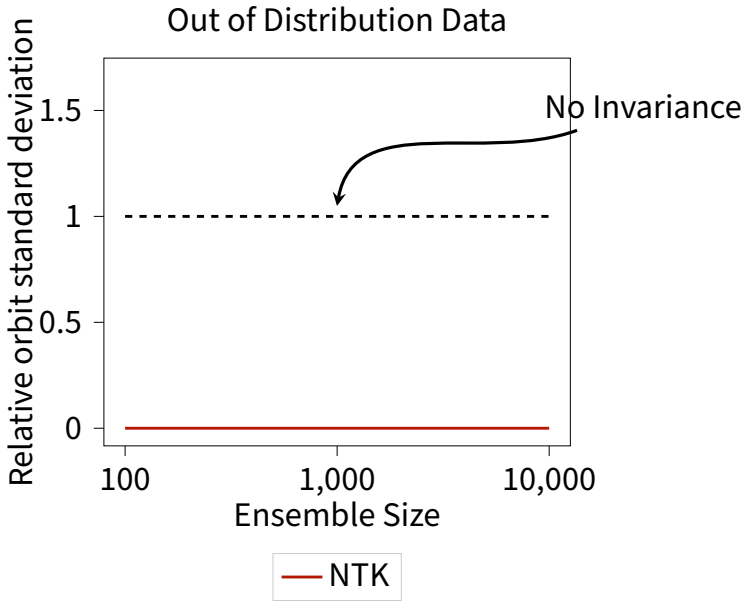
Out of Distribution Data
Relative orbit standard deviation
1.5
1
0.5
0
No Invariance
100
1,000
10,000
Ensemble Size
NTK

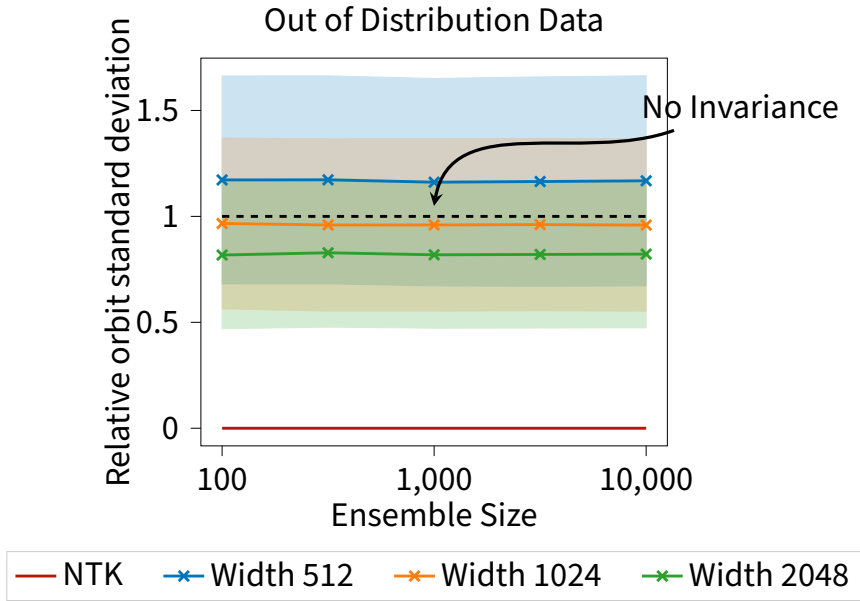
Out of Distribution Data
Relative orbit standard deviation
1.5
1
0.5
0
No Invariance
100
1,000
10,000
Ensemble Size
NTK
Width 512
Width 1024
Width 2048

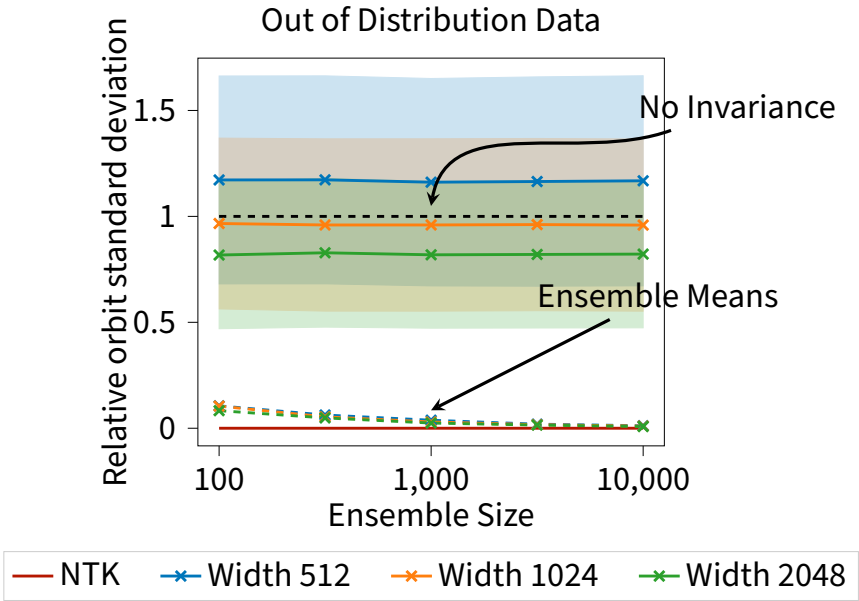
Out of Distribution Data
Relative orbit standard deviation
1.5
1
0.5
0
No Invariance
Ensemble Means
100
1,000
10,000
Ensemble Size
NTK
Width 512
Width 1024
Width 2048

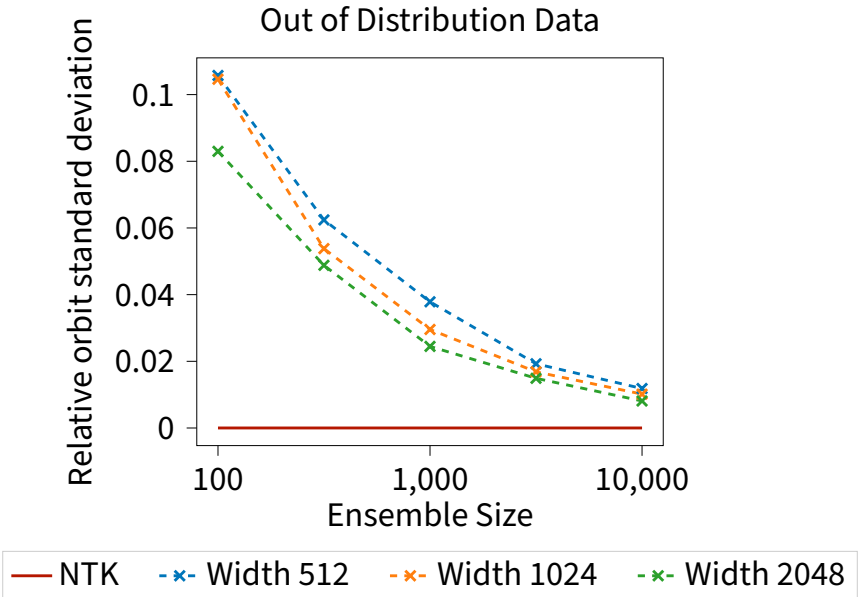
Out of Distribution Data
Relative orbit standard deviation
0.1
0.08
0.06
0.04
0.02
0
100
1,000
10,000
Ensemble Size
NTK
Width 512
Width 1024
Width 2048

Histological slices

[Kather et al. 2018]

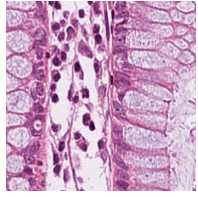

Histological slices

[Kather et al. 2018]

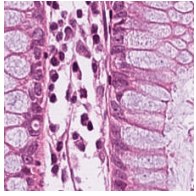

class 6

Histological slices

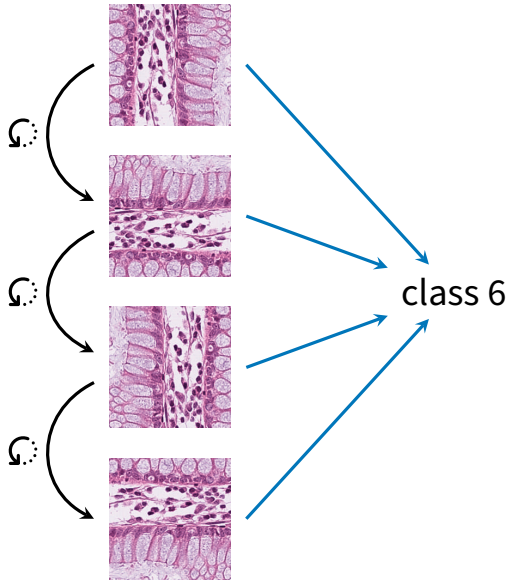

Histological slices

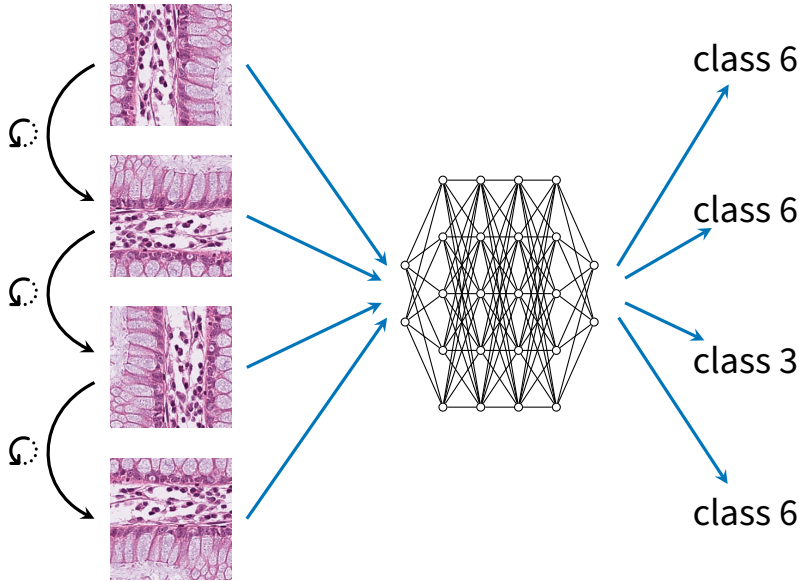

Histological slices

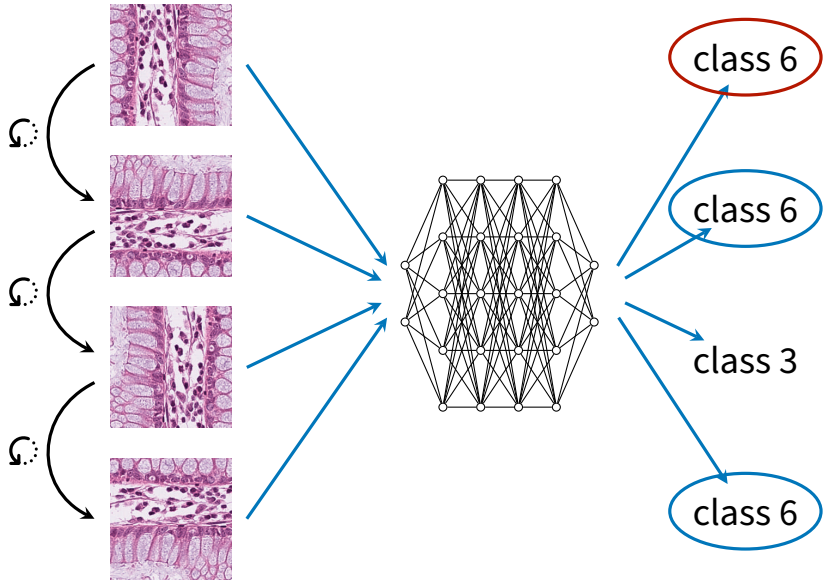

Histological slices

Orbit Same Predictions = 3

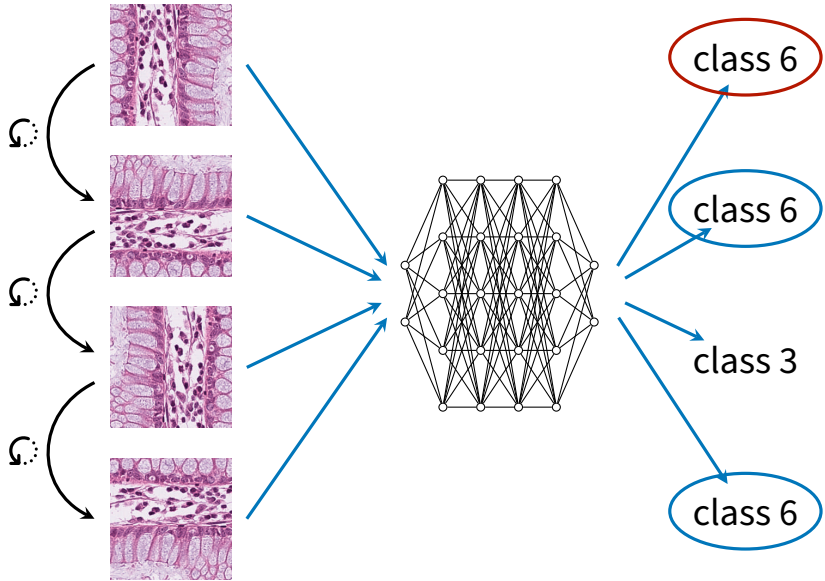

Out of distribution results

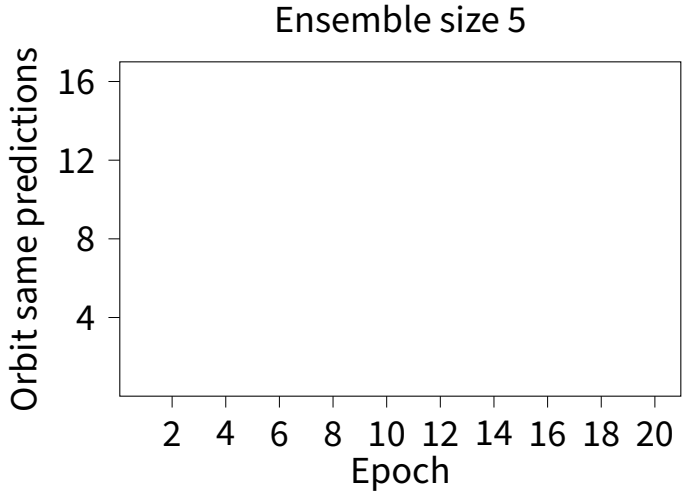

Out of distribution results

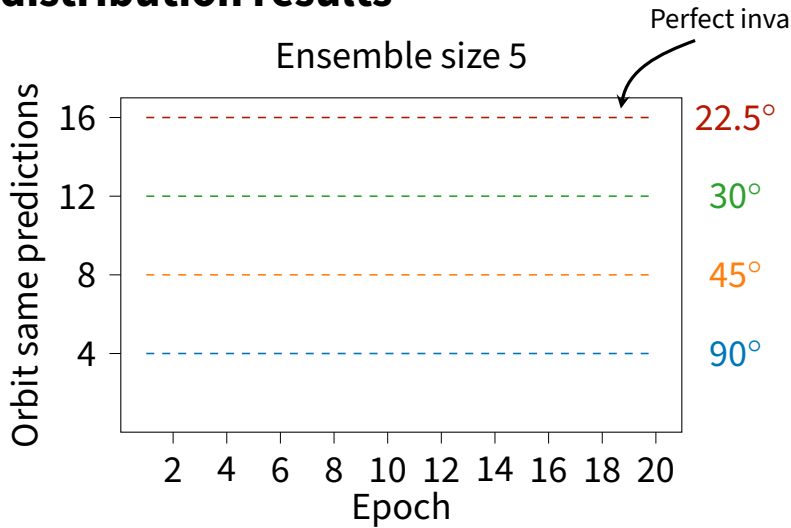

Out of distribution results

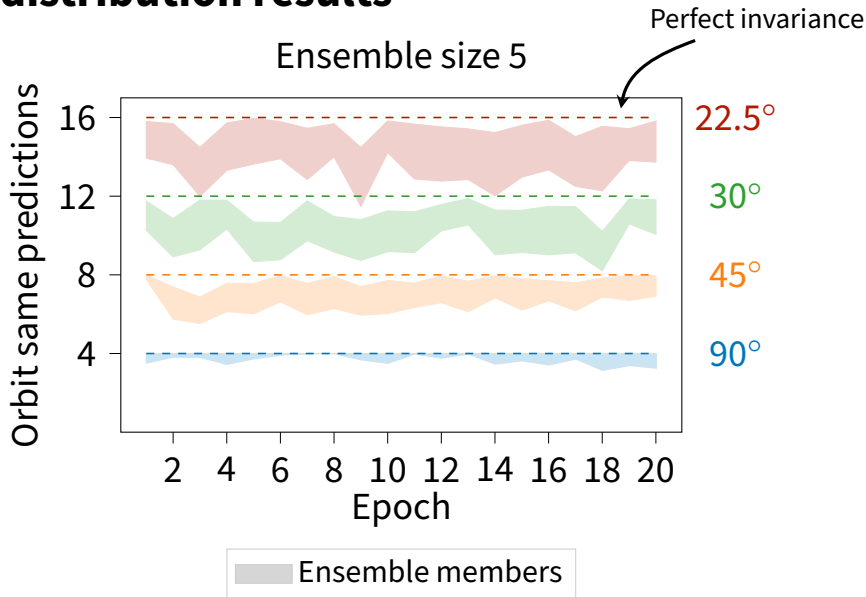

Out of distribution results

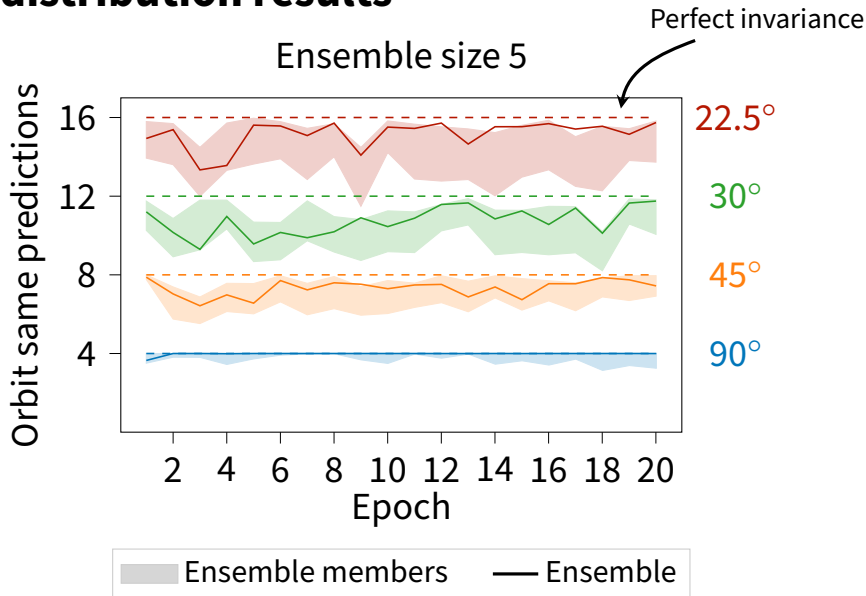

Out of distribution results

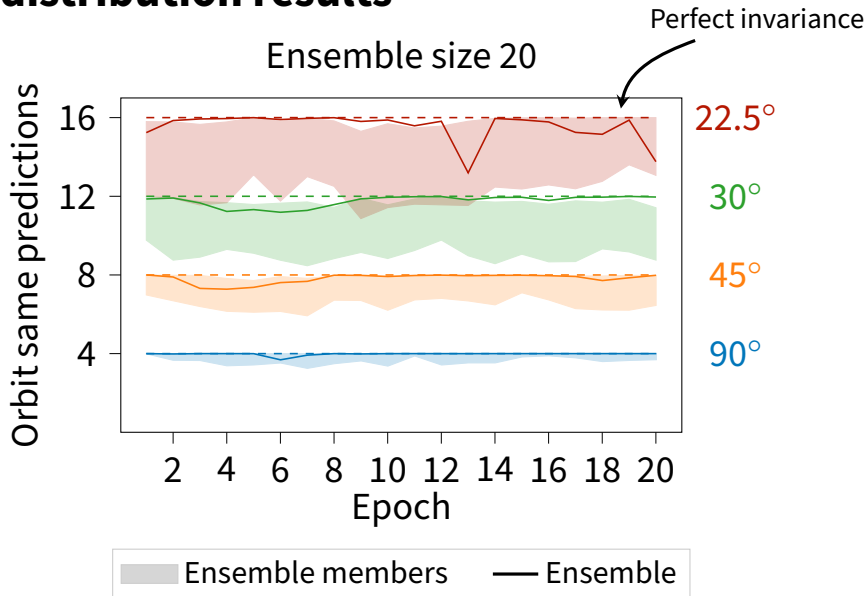

Further experimental results

Further experimental results

- ✓ Emergent invariance for rotated FashionMNIST

Further experimental results

- ✓ Emergent invariance for rotated FashionMNIST
- ✓ Partial augmentation for continuous symmetries

Further experimental results

- ✓ Emergent invariance for rotated FashionMNIST
- ✓ Partial augmentation for continuous symmetries
- ✓ Emergent equivariance (as opposed to invariance)

Comparison to other methods

Comparison to other methods

- Models trained on rotated FashionMNIST

Comparison to other methods

➪ Models trained on rotated FashionMNIST

Orbit same predictions out of distribution:

	C_4	C_8	C_{16}
DeepEns+DA	3.85±0.12	**7.72±0.34**	**15.24±0.69**
only DA	3.41±0.18	6.73±0.24	12.77±0.71
E2CNN[1]	**4±0.0**	**7.71±0.21**	**15.08±0.34**
Canon[2]	**4±0.0**	**7.45±0.14**	12.41±0.85

[1] [Weiler et al. 2019], [2] [Kaba et al. 2022]

Convergence of augmented CNNs to GCNNs

Convergence of augmented CNNs to GCNNs

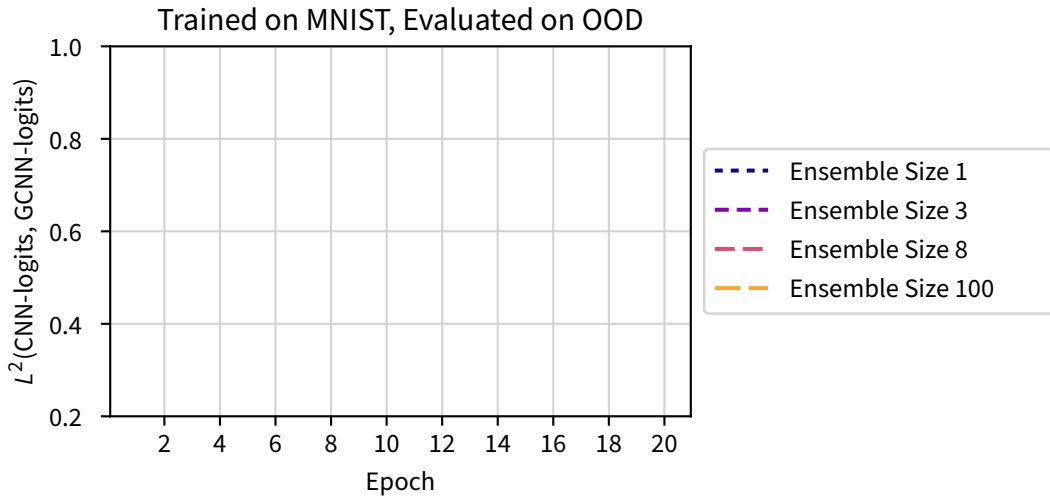

Convergence of augmented CNNs to GCNNs

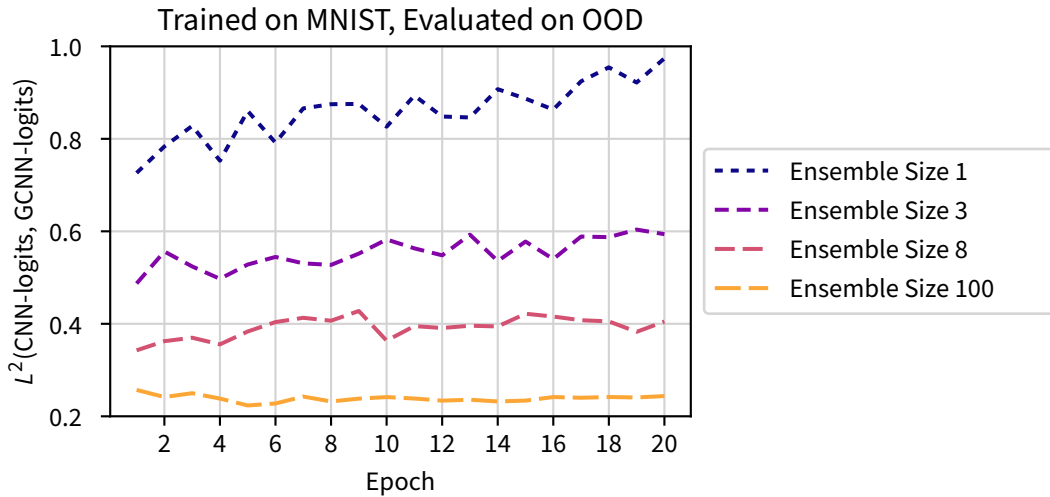

Convergence of augmented CNNs to GCNNs

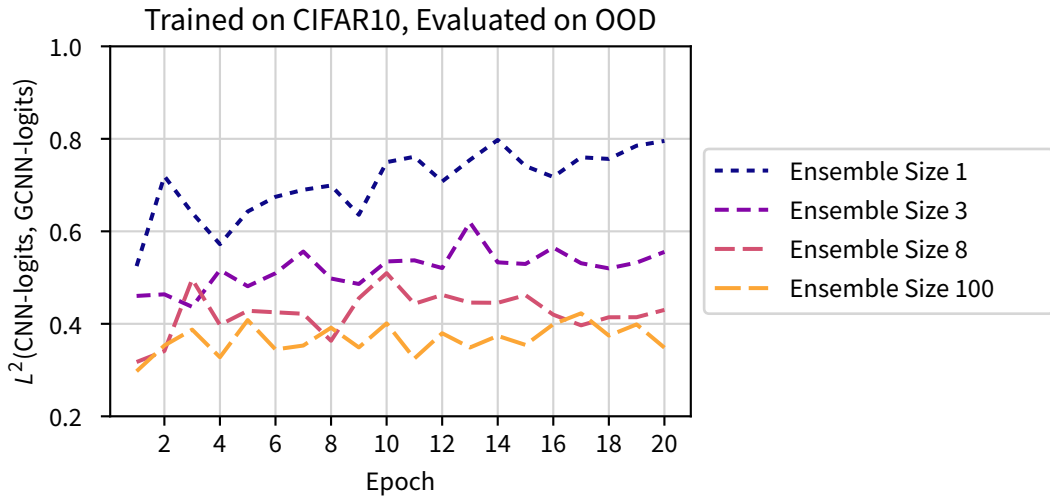

Key takeaways

Key takeaways

If you need ensembles

👍 use data augmentation to obtain an equivariant model.

Key takeaways

If you need ensembles

👍 use data augmentation to obtain an equivariant model.

If you need data augmentation

👍 use an ensemble to boost the equivariance.

Papers

- Emergent Equivariance in Deep Ensembles
  Jan E. Gerken*, Pan Kessel*
  ICML 2024 (Oral)
  * Equal contribution
- Equivariant Neural Tangent Kernels
  Philipp Misof, Pan Kessel, Jan E. Gerken
  ICML 2025

Thank you

Backup

Empirical NTK

- Consider continuous gradient descent

$$\frac{\mathrm{d}\theta_\mu}{\mathrm{d}t} = -\eta \frac{\partial \mathcal{L}(\mathcal{N}_\theta, \mathcal{D})}{\partial \theta_\mu} = -\frac{\eta}{N} \sum_{i=1}^{N} \frac{\partial L(\mathcal{N}_\theta(x_i), y_i)}{\partial \theta_\mu}$$

Empirical NTK

- Consider continuous gradient descent

$$\frac{\mathrm{d}\theta_\mu}{\mathrm{d}t} = -\eta \frac{\partial \mathcal{L}(\mathcal{N}_\theta, \mathcal{D})}{\partial \theta_\mu} = -\frac{\eta}{N} \sum_{i=1}^{N} \frac{\partial L(\mathcal{N}_\theta(x_i), y_i)}{\partial \theta_\mu}$$

- Then, the network evolves according to

$$\frac{\mathrm{d}\mathcal{N}_\theta(x)}{\mathrm{d}t} = \sum_\mu \frac{\partial \mathcal{N}_\theta(x)}{\partial \theta_\mu} \frac{\mathrm{d}\theta_\mu}{\mathrm{d}t}$$

Empirical NTK

- Consider continuous gradient descent

$$\frac{\mathrm{d}\theta_\mu}{\mathrm{d}t} = -\eta \frac{\partial \mathcal{L}(\mathcal{N}_\theta, \mathcal{D})}{\partial \theta_\mu} = -\frac{\eta}{N} \sum_{i=1}^{N} \frac{\partial L(\mathcal{N}_\theta(x_i), y_i)}{\partial \theta_\mu}$$

- Then, the network evolves according to

$$\frac{\mathrm{d}\mathcal{N}_\theta(x)}{\mathrm{d}t} = \sum_\mu \frac{\partial \mathcal{N}_\theta(x)}{\partial \theta_\mu} \frac{\mathrm{d}\theta_\mu}{\mathrm{d}t} = -\frac{\eta}{N} \sum_{i=1}^{N} \sum_\mu \frac{\partial \mathcal{N}(x)}{\partial \theta_\mu} \frac{\partial \mathcal{N}(x_i)}{\partial \theta_\mu} \frac{\partial L(\mathcal{N}_\theta(x_i), y_i)}{\partial \mathcal{N}(x_i)}$$

Empirical NTK

- Consider continuous gradient descent

$$\frac{\mathrm{d}\theta_\mu}{\mathrm{d}t} = -\eta \frac{\partial \mathcal{L}(\mathcal{N}_\theta, \mathcal{D})}{\partial \theta_\mu} = -\frac{\eta}{N} \sum_{i=1}^{N} \frac{\partial L(\mathcal{N}_\theta(x_i), y_i)}{\partial \theta_\mu}$$

- Then, the network evolves according to

$$\frac{\mathrm{d}\mathcal{N}_\theta(x)}{\mathrm{d}t} = \sum_\mu \frac{\partial \mathcal{N}_\theta(x)}{\partial \theta_\mu} \frac{\mathrm{d}\theta_\mu}{\mathrm{d}t} = -\frac{\eta}{N} \sum_{i=1}^{N} \sum_\mu \frac{\partial \mathcal{N}(x)}{\partial \theta_\mu} \frac{\partial \mathcal{N}(x_i)}{\partial \theta_\mu} \frac{\partial L(\mathcal{N}_\theta(x_i), y_i)}{\partial \mathcal{N}(x_i)}$$

With the empirical neural tangent kernel (NTK)

$$\Theta_{ij}^{\theta}(x, x') = \sum_\mu \frac{\partial \mathcal{N}_i(x)}{\partial \theta_\mu} \frac{\partial \mathcal{N}_j(x')}{\partial \theta_\mu}$$

Infinite width limit

Consider an MLP in NTK parametrization

$$z^{(\ell)} = \frac{1}{\sqrt{n_{\ell-1}}} W^{(\ell)} \sigma(z^{(\ell-1)}(x)), \quad W^{(\ell)} \in \mathbb{R}^{n_\ell \times n_{\ell-1}}, \quad W_{ij}^{(\ell)} \sim \mathcal{N}(0,1)$$

Infinite width limit

Consider an MLP in NTK parametrization

$$z^{(\ell)} = \frac{1}{\sqrt{n_{\ell-1}}} W^{(\ell)} \sigma(z^{(\ell-1)}(x)), \quad W^{(\ell)} \in \mathbb{R}^{n_\ell \times n_{\ell-1}}, \quad W_{ij}^{(\ell)} \sim \mathcal{N}(0,1)$$

At infinite width

$$z_i^{(\ell)}(x) = \sqrt{n_{\ell-1}} \frac{1}{n_{\ell-1}} \sum_{j=1}^{n_\ell} W_{ij}^{(\ell)} \sigma(z_j^{(\ell-1)}(x))$$

Infinite width limit

Consider an MLP in NTK parametrization

$$z^{(\ell)} = \frac{1}{\sqrt{n_{\ell-1}}} W^{(\ell)} \sigma(z^{(\ell-1)}(x)), \quad W^{(\ell)} \in \mathbb{R}^{n_\ell \times n_{\ell-1}}, \quad W_{ij}^{(\ell)} \sim \mathcal{N}(0,1)$$

At infinite width

$$z_i^{(\ell)}(x) = \sqrt{n_{\ell-1}} \underbrace{\frac{1}{n_{\ell-1}} \sum_{j=1}^{n_\ell}}_{\text{mean}} \underbrace{W_{ij}^{(\ell)} \sigma(z_j^{(\ell-1)}(x))}_{\text{i.i.d.}}$$

Infinite width limit

Consider an MLP in NTK parametrization

$$z^{(\ell)} = \frac{1}{\sqrt{n_{\ell-1}}} W^{(\ell)} \sigma(z^{(\ell-1)}(x)), \quad W^{(\ell)} \in \mathbb{R}^{n_\ell \times n_{\ell-1}}, \quad W_{ij}^{(\ell)} \sim \mathcal{N}(0,1)$$

At infinite width

$$z_i^{(\ell)}(x) = \sqrt{n_{\ell-1}} \underbrace{\frac{1}{n_{\ell-1}} \sum_{j=1}^{n_\ell}}_{\text{mean}} \underbrace{W_{ij}^{(\ell)} \sigma(z_j^{(\ell-1)}(x))}_{\text{i.i.d.}} \sim \mathcal{N}(0, \underbrace{\mathrm{Cov}(z_i^{(\ell)}, z_j^{(\ell)})}_{\text{NNGP}})$$

NTK at initialization

[Jacot et al. 2020]

When taking the layer widths to infinity sequentially, the empirical NTK $\Theta_{ij}^{\theta}(x, x')$ at initialization converges in probability to a deterministic kernel $\Theta(x, x')\delta_{ij}$

NTK at initialization

[Jacot et al. 2020]

When taking the layer widths to infinity sequentially, the empirical NTK $\Theta^{\theta}_{ij}(x, x')$ at initialization converges in probability to a deterministic kernel $\Theta(x, x')\delta_{ij}$

- The deterministic kernel is given in terms of a recursion over layers

NTK at initialization

[Jacot et al. 2020]

When taking the layer widths to infinity sequentially, the empirical NTK $\Theta_{ij}^{\theta}(x, x')$ at initialization converges in probability to a deterministic kernel $\Theta(x, x')\delta_{ij}$

- The deterministic kernel is given in terms of a recursion over layers
- For most common architectures, this recursion can be performed explicitly,
  e.g. using `neural-tangents` Python package [Novak et al. 2020]

Frozen NTK

[Jacot et al. 2020]

For a nonlinearity which is Lipschitz, twice differentiable and has bounded second derivative,

$$\Theta_{ij}^{\theta_t}(x, x') \to \Theta(x, x')\delta_{ij}$$

uniformly in t as the layer widths go to infinity sequentially.

Frozen NTK

[Jacot et al. 2020]

For a nonlinearity which is Lipschitz, twice differentiable and has bounded second derivative,

$$\Theta_{ij}^{\theta_t}(x, x') \to \Theta(x, x')\delta_{ij}$$

uniformly in t as the layer widths go to infinity sequentially.

- Intuitively, this happens because the weight updates vanish in the limit $n \to \infty$

Frozen NTK

[Jacot et al. 2020]

For a nonlinearity which is Lipschitz, twice differentiable and has bounded second derivative,

$$\Theta_{ij}^{\theta_t}(x, x') \to \Theta(x, x')\delta_{ij}$$

uniformly in t as the layer widths go to infinity sequentially.

- Intuitively, this happens because the weight updates vanish in the limit $n \to \infty$
- However, the network still learns

Kernel transformation

The neural tangent kernel Θ as well as the NNGP kernel K transform according to

$$\Theta(\rho(g)x, \rho(g)x') = \rho_K(g)\Theta(x, x')\rho_K^\top(g)\,,$$
$$K(\rho(g)x, \rho(g)x') = \rho_K(g)K(x, x')\rho_K^\top(g)\,,$$

for all $g \in G$ and $x, x' \in X$.

Kernel transformation

The neural tangent kernel Θ as well as the NNGP kernel K transform according to

$$\Theta(\rho(g)x, \rho(g)x') = \rho_K(g)\Theta(x, x')\rho_K^\top(g)\,,$$
$$K(\rho(g)x, \rho(g)x') = \rho_K(g)K(x, x')\rho_K^\top(g)\,,$$

for all $g \in G$ and $x, x' \in X$.

Hence, for MLPs,

$$\Theta(\rho(g)x, \rho(g)x') = \Theta(x, x') \quad \Rightarrow \quad \Theta(\rho(g)x, x') = \Theta(x, \rho^{-1}(g)x')$$

Permutation shift

- On the training data, group transformations permute the samples

$$\rho(g)x_i = x_{\pi_g(i)}, \qquad \pi_g \in S_N$$

Permutation shift

- On the training data, group transformations permute the samples

$$\rho(g)x_i = x_{\pi_g(i)}, \qquad \pi_g \in S_N$$

- Therefore, for a permutation of training samples associate to g

$$\begin{aligned}\Pi(g)\Theta(X,X) &= \Theta(\rho(g)X,X)\\ &= \Theta(X,\rho^{-1}(g)X)\\ &= \Theta(X,X)(\Pi^{-1}(g))^\top\\ &= \Theta(X,X)\Pi(g)\end{aligned}$$

NTKs for GCNNs

For GCNN-layers, define the NNGP and NTK via

$$K^{(\ell)}_{g,g'}(f,f') = \mathbb{E}\left[[\mathcal{N}^{(\ell)}(f)](g)\left([\mathcal{N}^{(\ell)}(f')](g')\right)^{\top}\right]$$

NTKs for GCNNs

For GCNN-layers, define the NNGP and NTK via

$$K^{(\ell)}_{g,g'}(f,f') = \mathbb{E}\left[[\mathcal{N}^{(\ell)}(f)](g)\left([\mathcal{N}^{(\ell)}(f')](g')\right)^{\top}\right]$$

$$\Theta^{(\ell)}_{g,g'}(f,f') = \mathbb{E}\left[\sum_{\ell'=1}^{\ell}\frac{\partial[\mathcal{N}^{(\ell)}(f)](g)}{\partial\theta^{(\ell')}}\left(\frac{\partial[\mathcal{N}^{(\ell)}(f')](g')}{\partial\theta^{(\ell')}}\right)^{\top}\right]$$

NTKs for GCNNs

$$[\mathcal{N}^{(\ell)}(f)](g) = \int_G \mathrm{d}g\, \kappa(g^{-1}h)\, [\mathcal{N}^{(\ell-1)}(f)](h)$$

The layer-recursion for a GCNN-layer is given by

$$K^{(\ell+1)}_{g,g'}(f,f') = \frac{1}{|S_\kappa|} \int_{S_\kappa} \mathrm{d}h\, K^{(\ell)}_{gh,g'h}(f,f')$$

NTKs for GCNNs

$$[\mathcal{N}^{(\ell)}(f)](g) = \int_G \mathrm{d}g\, \kappa(g^{-1}h)\, [\mathcal{N}^{(\ell-1)}(f)](h)$$

The layer-recursion for a GCNN-layer is given by

$$K^{(\ell+1)}_{g,g'}(f,f') = \frac{1}{|S_\kappa|} \int_{S_\kappa} \mathrm{d}h\, K^{(\ell)}_{gh,g'h}(f,f')$$

$$\Theta^{(\ell+1)}_{g,g'}(f,f') = K^{(\ell+1)}_{g,g'}(f,f') + \frac{1}{|S_\kappa|} \int_{S_\kappa} \mathrm{d}h\, \Theta^{(\ell)}_{gh,g'h}(f,f')$$

Emergent equivariance of cross products

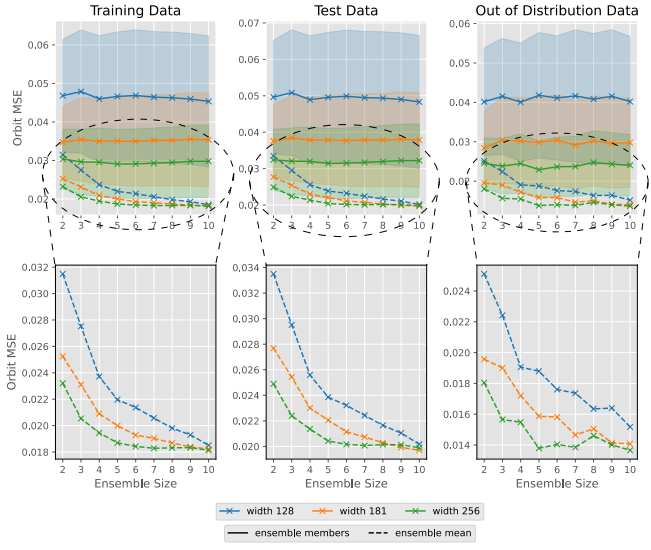

Histological Data – OOD samples

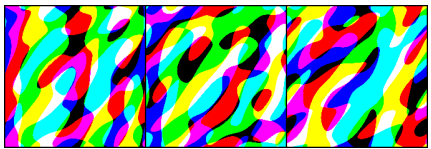

Emergent continuous symmetry on FashionMNIST

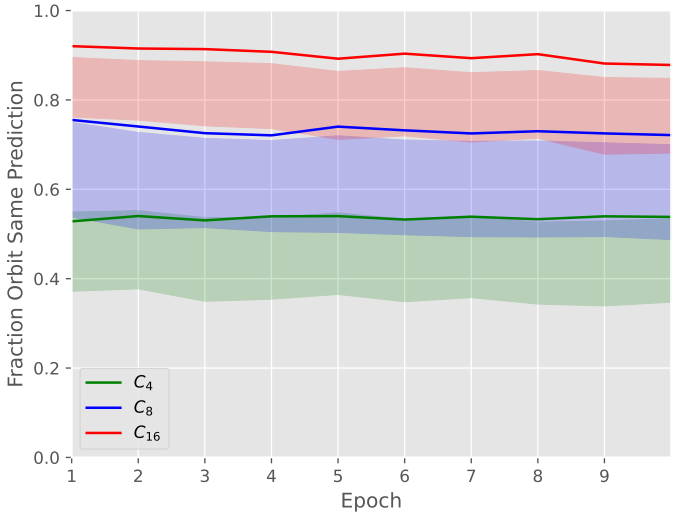

Kernel convergence

NNGP

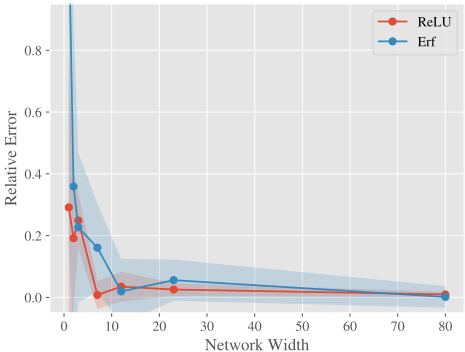

NTK

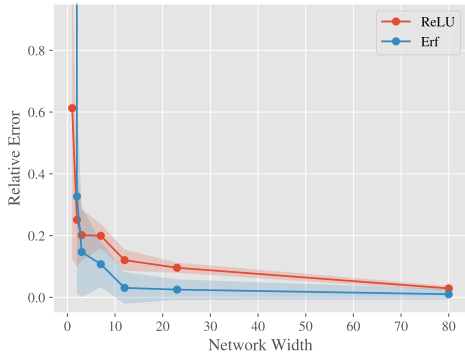

Equivariant NTKs for medical image classification

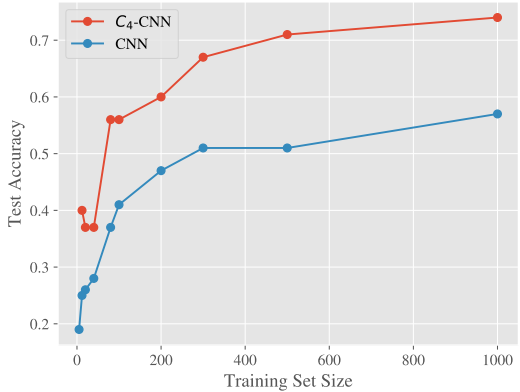

Equivariant NTKs for molecular property regression

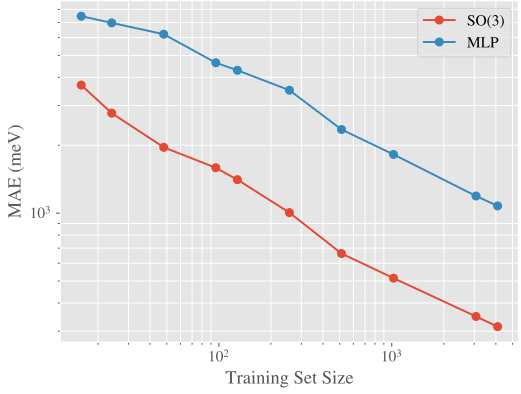

OOD samples for CNN to GCNN convergence

MNIST

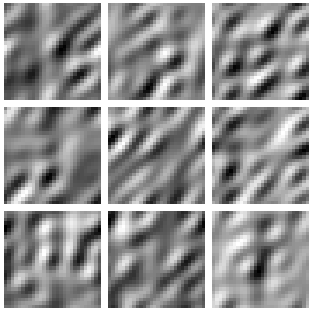

CIFAR10