# Preface

**Sophia Sanborn** SANBORN@UCSB.EDU
*University of California, Santa Barbara*

**Christian Shewmake** SHEWMAKE@BERKELEY.EDU
*University of California, Berkeley*

**Simone Azeglio** SIMONE.AZEGLIO@INSERM.FR
*Sorbonne University & École Normale Supérieure, Paris*

**Arianna Di Bernardo** ARIANNA.DI.BERNARDO@ENS.PSL.EU
*École Normale Supérieure, Paris*

**Nina Miolane** NINAMIOLANE@UCSB.EDU
*University of California, Santa Barbara*

**Editors:** Sophia Sanborn, Christian Shewmake, Simone Azeglio, Arianna Di Bernardo, Nina Miolane

## 1. Symmetry and Geometry in Neural Representations

An emerging set of findings in sensory and motor neuroscience is beginning to illuminate a new paradigm for understanding the neural code. Across sensory and motor regions of the brain, neural circuits are found to mirror the geometric and topological structure of the systems they represent—either in their synaptic structure, or in the implicit manifold generated by their activity. This phenomenon can be observed in the circuit of neurons representing head direction in the fly Kim et al. (2017); Wolff et al. (2015), in the activities of grid cells Chaudhuri et al. (2019); Gardner et al. (2022), and in the low-dimensional manifold structure observed in motor cortex Gallego et al. (2017). This suggests a general computational strategy that is employed throughout the brain to preserve the geometric structure of data throughout stages of information processing.

Independently but convergently, this very same computational strategy has emerged in the field of deep learning. The nascent sub-field of Geometric Deep Learning Bronstein et al. (2021) incorporates geometric priors into artificial neural networks to preserve the geometry of signals as they are passed through layers of the network. This approach provably demonstrates gains in the computational efficiency, robustness, and generalization performance of these models.

The convergence of these findings suggests deep, substrate-agnostic principles for information processing. Symmetry and geometry were instrumental in unifying models of fundamental forces and elementary particles in 20th-century physics. Likewise, they have the potential to illuminate unifying principles for how neural systems form useful representations of the world.

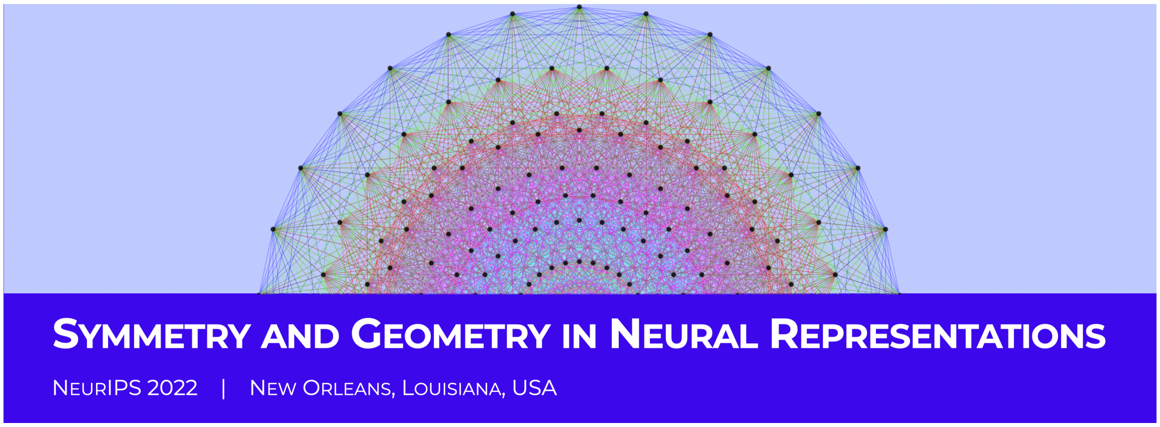


## 2. The Workshop

The first annual NeurIPS Workshop on Symmetry and Geometry in Neural Representations (NeurReps) was conceived to bring together researchers at the nexus of applied geometry, deep learning, and neuroscience, with the goal of advancing this understanding and illuminating geometric principles for neural information processing. Ultimately, we hope that this venue and associated community will support the development of the geometric approach to understanding neural representations, while strengthening ties to the mathematics community. The Neural Information Processing Systems (NeurIPS) conference historically emerged from the field of theoretical neuroscience or "connectionism." This venue thus further serves the workshop's goal of reinforcing the bond between deep learning and neuroscience.

### 2.1. Call for Papers

Our call for papers invited submissions for publication in this volume and presentation at the conference. The call requested submissions contributing novel research at the intersection of geometric deep learning, computational neuroscience, geometric statistics, and topological data analysis, which incorporate symmetry, geometry, or topology into the *design of artificial neural networks*, the *analysis of neural data*, or *theories of neural computation*. Both theoretical contributions and applied results were encouraged, as well as the use of diverse mathematical objects such as quotient spaces, fiber bundles, Lie groups, Riemannian manifolds, graphs, and group representations. Two tracks were established:

1. **Proceedings Track**: NeurReps Proceedings papers are up to 10 pages long, excluding references and appendices. This track is intended for self-contained research papers with a high degree of development. Accepted papers are published in this volume.

2. **Extended Abstract Track**: Extended abstracts are up to 5 pages long, excluding references and appendices. This track is a space for contributions such as early-stage results, insightful negative findings, opinion pieces, or novel datasets. Extended abstracts are not included in this volume.

The call for papers yielded 91 submissions, which were reviewed double-blind. Submitted work spanned from questions in statistical learning theory for geometric models, to novel

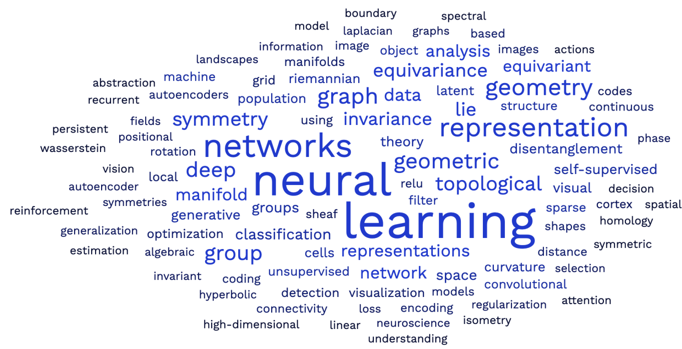


**Figure 1:** A word cloud constructed from the abstracts of all submissions.

geometric deep learning architectures, to methods for neural data analysis. Each submission received a minimum of three reviews, which were aggregated and assessed by the editors for final inclusion. This resulted in 65 accepted works: 25 full-length papers, and 41 extended abstracts. Of the 25 full-length papers, 21 are published in this volume, by choice of the authors. Reviews of the accepted works can be found on the NeurReps OpenReview Portal.

### 2.2. The Schedule

The workshop was held on December 3rd, 2022 in New Orleans, Lousiana, USA, and attracted between 300 and 400 in-person attendees. The workshop featured six invited keynote speakers with talks spanning topics in neuroscience, machine learning, and their intersection. In addition, two discussion panels were held. All accepted works were presented as posters during the workshop, with the top ten submissions selected for oral presentation. Three submissions were selected for ten-minute spotlights and the remainder selected for five-minute lightning talks. Awards were given to the best proceedings paper and best extended abstract. All talks and panels were live-streamed and recorded and can be found online on SlidesLive.

Keynotes

**In search of invariance in brains and machines**
Bruno Olshausen

**Symmetry-based representations for artificial and biological intelligence**
Irina Higgins

**From equivariance to naturality**
Taco Cohen

**Generative models of non-Euclidean neural population dynamics**
Kristopher Jensen

**Robustness of representations in artificial and biological neural networks**
Gabriel Kreiman

**Neural ideograms and equivariant representation learning**
Erik Bekkers

### Discussion Panels

**Geometric and topological principles for representation learning in ML**
*Panelists*: Irina Higgins, Taco Cohen, Erik Bekkers, Rose Yu
*Moderator*: Nina Miolane

**Geometric and topological principles for representations in the brain**
*Panelists*: Bruno Olshausen, Kristopher Jensen, Gabriel Krieman, Manu Madhav
*Moderator*: Christian Shewmake

### Spotlight Talks

**Is the information geometry of probabilistic population codes learnable?**
Vastola[†], Cohen, Drugowitsch

**Computing representations for Lie Algebraic Networks**
Shutty[†], Wierzynski
*Awarded Best Proceedings Paper*

**Kendall Shape-VAE : Learning shapes in a generative framework**
Vadgama[†], Tomczak, Bekkers

### Lightning Talks

**Equivariance with learned canonical mappings**
Kaba[†*], Mondal[*], Zhang, Bengio, Ravanbakhsh

**Capacity of group-invariant linear readouts from equivariant representations: How many objects can be linearly classified under all possible views?**
Farrell[†], Bordelon, Trivedi, Pehlevan

**Do neural networks trained with topological features learn different internal representations?**
McGuire, Jackson, Emerson, Kvinge[†]

**Expander Graph Propagation**
Deac[†], Lackenby, Veličković

**Homomorphism AutoEncoder: Learning group structured representations from observed transitions**
Keurti[†], Pan, Besserve, Grewe, Schölkopf
*Awarded Best Extended Abstract*

**Sheaf Attention Networks**
Barbero[†], Bodnar, Sáez de Ocáriz Borde, Liò

**On the expressive power of geometric graph neural networks**
Joshi[*], Bodnar[*], Mathis[†], Cohen, Liò

[†] Denotes presenting author. [*] Denotes shared first-authorship

## 3. The Community

To support the growth of this nascent research area outside of the annual workshop, we have established a digital community for NeurReps, which at the time of writing has over 800 members. Instructions for joining and contributing to the community can be found on the community page of our website.

## 4. Program Committee

We are immensely grateful to our 95-member program committee, who collectively submitted 267 high-quality reviews for the workshop's 91 submissions. Thank you to:

Adele Myers
Adrian Valente
Alessandro Sarti
Alex Williams
Alexandra Libby
Alice Le Brigant
Alison Poupin
Andrew Ligeralde
Anna Calissano
Balasubramaniam Srinivasan
Bastian Rieck
Bilal Alsallakh
Blake Bordelon
Boyan Beronov
Bruno Olshausen
Chris Kymn
Christoph Ortner
Christopher Hillar
Christopher Kim
Claire Donnat
Clementine Domine
David Klee
David Klindt
David Robin
Davide Boscaini
Donlapark Ponnoprat
Dorina Thanou
Dylan Paiton
Edouard Oyallon
Elodie Maignant
Emanuele Marconato
Emanuele Rodolà
Emanuele Rossi
Erik Bekkers
Federico Claudi
Francesco Di Giovanni
Francisco Acosta
Frank Nielsen
Frédéric Barbaresco
Geoffrey Woollard
Grégoire Sergeant-Perthuis
Hannah Lawrence
Henry Adams
Hrittik Roy
Ilyes Batatia
Jacob Zavatone-Veth
James Whittington
Joey Bose
Johan Mathe
Jonathan Huml
Justin Solomon
Kaitlin Maile
Kartik Sharma
Kathryn Hess
Khanh Dao Duc
Kristopher Jensen
Maksim Zhdanov
Manos Theodosis
Manu Madhav
Marco Fumero
Marco Pegoraro
Mathilde Papillon
Maurice Weiler
Mikail Khona
Mitchell Ostrow
Nicolas Guigui
Noah Shutty
Oded Stein
Ondrej Biza
Patrick Rubin-Delanchy
Pim de Haan
Rana Shahroz
Rongjie Lai
Rucha Joshi
Santiago Cadena
Sarah Marzen
Sharvaree Vadgama
Shayan Shekarforoush
Shubhendu Trivedi
Simon Mathis
Sina Tootoonian
Søren Hauberg
Stéphane Deny
Sylvain Chevallier
Tamar Flash
Tatyana Sharpee
Uri Cohen
Valentino Maiorca
Vincent Benenati
Wei Ye
Will Dorrell
Wolfgang Polonik
Xiangru Huang
Xiaoling Hu
Xinling Yu
Yann Thanwerdas
Yubei Chen

**Figure 2:** The NeurReps poster session.

## 5. Moving Forward

We believe it is both timely and important to create a research venue and supportive community for the exchange of knowledge at the intersection of differential geometry, topology, machine learning and neuroscience. Moving forward, we will continue to create opportunities for dialogue and discussion on these themes at NeurIPS and other meetings. Furthering our broader aim of community-building, we have also established an active community of students and researchers which we believe will act as a gathering place to organize related events, such as seminars and hackathons.

## References

Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.

Rishidev Chaudhuri, Berk Gerçek, Biraj Pandey, Adrien Peyrache, and Ila Fiete. The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nature neuroscience*, 22(9):1512–1520, 2019.

Juan A Gallego, Matthew G Perich, Lee E Miller, and Sara A Solla. Neural manifolds for the control of movement. *Neuron*, 94(5):978–984, 2017.

Richard J Gardner, Erik Hermansen, Marius Pachitariu, Yoram Burak, Nils A Baas, Benjamin A Dunn, May-Britt Moser, and Edvard I Moser. Toroidal topology of population activity in grid cells. *Nature*, pages 1–6, 2022.

Sung Soo Kim, Hervé Rouault, Shaul Druckmann, and Vivek Jayaraman. Ring attractor dynamics in the drosophila central brain. *Science*, 356(6340):849–853, 2017.

Tanya Wolff, Nirmala A Iyer, and Gerald M Rubin. Neuroarchitecture and neuroanatomy of the drosophila central complex: A gal4-based dissection of protocerebral bridge neurons and circuits. *Journal of Comparative Neurology*, 523(7):997–1037, 2015.