
Conditioned Clifford-Steerable Kernels

Anonymous Author(s)

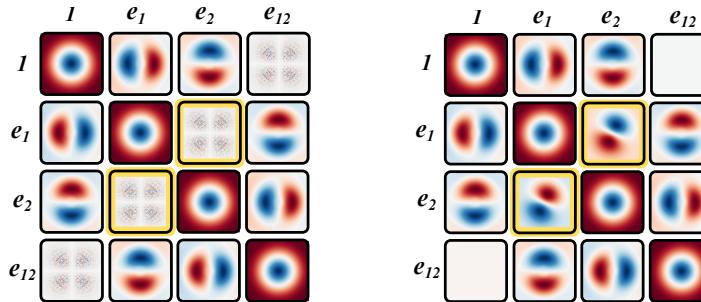
Affiliation
Address
email

Abstract

Clifford-Steerable CNNs (CSCNNs) provide a unified framework that allows incorporating equivariance to arbitrary pseudo-Euclidean groups, including $E(n)$ and Poincaré-equivariance on Minkowski spacetime. In this work, we analyze the shortcomings of the approach. We demonstrate that the kernel basis used in CSCNNs is not complete. Furthermore, we suggest to restore missing degrees of freedom by using an extra information obtained directly from data at virtually no cost. Our approach significantly and consistently outperforms baseline methods on PDE forecasting tasks, specifically fluid dynamics and relativistic electrodynamics.

1 Introduction

Physical systems are often associated with symmetries that govern their evolution. Therefore, it is desirable to respect those symmetries in the modelling process to faithfully capture the dynamics. *Equivariant Neural Networks* have proven to be strong candidates for learning such dynamics due to their ability to produce outputs that transform consistently under these symmetries. As part of the equivariant family, *Steerable Convolutional Neural Networks* ensure equivariance by enforcing a constraint over their convolution kernels. The steerability constraint is group-specific and has been solved and implemented for several common symmetry groups, such as for $O(2)$ (Weiler and Cesa, 2019), $O(3)$ (Geiger et al., 2022), and for general compact groups $E(n)$ as demonstrated by Lang and Weiler (2020); Cesa et al. (2022).



(a) Default Clifford-steerable Kernel (b) Conditioned Clifford-steerable Kernel

Figure 1: Comparison of default vs. conditioned kernels for $O(2)$. The missing angular frequency basis kernels can be observed as the sparse scatters (which are just noise) in the antidiagonal of subfigure 1a. In comparison, subfigure 1b shows how our implementation recovers these missing kernels. Due to the kernel’s input-dependent activation, the typical four-lobed quadrupole patterns may not be visible, despite being present.

20 Clifford Steerable Convolutional Neural Networks (CSCNNs), introduced by Zhdanov et al. (2024),
 21 extend steerability to a more general class of symmetries - pseudo-Euclidean group $E(p, q)$, which
 22 covers, for instance, $E(n)$ -equivariance and Poincaré-equivariance on Minkowski spacetime. By
 23 using implicit parameterization of the kernels, CSCNNs are able to process feature fields defined on
 24 pseudo-Euclidean spaces, allowing CNNs to learn system dynamics in both non-relativistic (e.g. fluid
 25 dynamics) and relativistic (e.g. electromagnetism) scenarios. However, as pointed out by the authors,
 26 the original implementation yields incomplete kernel basis, which in turn limits the expressivity of
 27 the model as it is unable to represent certain degrees of freedom and therefore capture corresponding
 28 interactions (see the missing anti-diagonals in Figure 1).

29 The authors suggested that this degree of freedom may be recovered; however, this has yet to
 30 be proven or experimentally validated. We present *Conditioned Clifford-Steerable Kernels* which
 31 augment the kernel with an independent representation constructed from the input field, recovering
 32 the missing degrees of freedom, thereby achieving a strictly more expressive parameterization. The
 33 contributions of this paper are the following:

- 34 • We show that the original parametrization of CS-CNN kernels is missing kernels of angular
 35 frequency 2.
- 36 • We prove analytically that the missing degrees of freedom can be recovered by extending the
 37 input vector of the kernel by a global representation of the feature field, at little to no cost.
- 38 • We propose an efficient implementation and evaluate it on multiple PDE simulations,
 39 outperforming other convolutional baselines.

40 This paper is organised as follows: Section 2 provides a brief introduction to Steerable CNNs and
 41 the Clifford Algebra. Section 4 uncovers the main limitation of Clifford-Steerable CNN kernels,
 42 and introduces Conditioned Clifford-Steerable Kernels as a solution. After deriving their improved
 43 properties, the new kernels are empirically validated against the original implementation in Section 5.

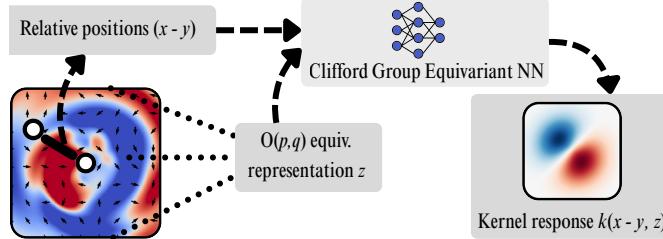


Figure 2: Illustration of Conditioned Clifford-Steerable Kernels. Similar to CSCNNs, the parameterization is implicit, with the only difference that the kernel network is augmented with an equivariant representation of the input feature field.

44 2 Theoretical Background

45 2.1 Pseudo-Euclidean Spaces

46 Pseudo-Euclidean spaces generalize Euclidean space by allowing an indefinite metric, so the squared
 47 length of a nonzero vector can be positive, zero, or negative. Every finite dimensional pseudo-
 48 Euclidean space is isometric to the *standard* pseudo-Euclidean space, to which we restrict our
 49 attention (O’Neill, 1983).

50 **Definition 2.1** (Standard pseudo-Euclidean vector spaces). Let e_1, \dots, e_{p+q} be the canonical basis
 51 of the real vector space \mathbb{R}^{p+q} . Define an *inner product of signature* (p, q) -that is, a non-degenerate
 52 symmetric bilinear form with p positive and q negative eigenvalues-by $\eta^{p,q}(v_1, v_2) := v_1^T \Delta^{p,q} v_2$
 53 where $v_1, v_2 \in \mathbb{R}^{p+q}$, and $\Delta^{p,q} := \text{diag}(\underbrace{1, \dots, 1}_{p \text{ times}}, \underbrace{-1, \dots, -1}_{q \text{ times}})$.

54 The resulting inner-product space $\mathbb{R}^{p,q} := (\mathbb{R}^{p+q}, \eta^{p,q})$ is the *standard pseudo-Euclidean vector*
 55 *space of signature* (p, q) . This is the space where our feature vectors are defined, and also the space
 56 the symmetry groups we are interested in are acting on.

57 **2.2 Clifford algebra and multivector fields**

58 In previous works Ruhe et al. (2023); Zhdanov et al. (2024), it has been demonstrated that Clifford
 59 algebra $\text{Cl}(\mathbb{R}^{p,q})$ is a representation space of the pseudo-orthogonal group $\text{O}(p, q)$. In simpler words,
 60 it means that elements of the algebra - multivectors - can be used to represent actions of the group.
 61 Below we provide the formal definition. Let V be a vector space over a field \mathbb{F} , equipped with a
 62 bilinear form: $\eta : V \times V \rightarrow \mathbb{R}$, $(v_1, v_2) \mapsto \eta(v_1, v_2)$. The *Clifford Algebra* generated by V is the
 63 unitary, associative, non-commutative algebra such that

$$v_i \bullet v_i = \eta(v_i, v_i) \cdot \mathbf{1}_{\text{Cl}(V, \eta)}, \quad \forall v_i \in V \subset \text{Cl}(V, \eta) \quad (1)$$

64 Where $\bullet := \text{Cl}(V, \eta) \times \text{Cl}(V, \eta) \rightarrow \text{Cl}(V, \eta)$ is the *geometric product*, an $\text{O}(p, q)$ -equivariant opera-
 65 tion enabling us to multiply elements of the algebra while respecting their geometric properties. Such
 66 elements are *multivectors*, which can be constructed by linearly combining products of $v_1, \dots, v_n \in V$
 67 the following way: $x = \sum_{i \in I} c_i \cdot v_{i,1} \bullet \dots \bullet v_{i,l_i}$ where index set I is finite, $v_{i,k} \in V$ and $c_i \in \mathbb{R}$.
 68 We can decompose the algebra into vector subspaces $\text{Cl}^{(n)}(V, \eta)$, $n = 0, \dots, d$, which we refer
 69 to as *grades*. One grade $n = |A|$ is composed by choosing n out of the maximum d basis vectors
 70 of the underlying vector space V and has dimensionality $\dim \text{Cl}^{(d)}(V, \eta) = \binom{d}{n}$. Grades $n = 0$ and
 71 $n = 1$ consist of scalars and vectors, respectively. For $n \geq 2$, one encounters bivectors, trivectors
 72 and higher-grade multivectors, which geometrically represent oriented points, areas and their higher-
 73 dimensional analogues. For a more in depth introduction to the Clifford Algebra, we refer the reader
 74 to Appendix D of Zhdanov et al. (2024).

75 **2.3 Clifford Steerable CNNs**

76 *Clifford-Steerable CNNs* process multivector fields, which are functions $f : \mathbb{R}^{p,q} \rightarrow \text{Cl}(\mathbb{R}^{p,q})^c$
 77 assigning a feature $f(x)$ to each point $x \in \mathbb{R}^{p,q}$ in feature vector space $\text{Cl}(\mathbb{R}^{p,q})^c$ where c denotes the
 78 feature channel dimension. They are required to commute with the pseudo-Euclidean symmetry group
 79 $\text{E}(p, q) = (\mathbb{R}^{p,q}, +) \rtimes \text{O}(p, q)$, which is the semi-direct product of the the *translation group* $(\mathbb{R}^{p,q}, +)$
 80 and the *pseudo-orthogonal group* $\text{O}(p, q) = \{g \in \text{GL}(\mathbb{R}^{p,q}) \mid g^\top \Delta^{p,q} g = \Delta^{p,q}\}$. Convolutions
 81 are by construction rotation equivariant, thus the task of constructing an $\text{E}(p, q)$ equivariant CNN
 82 reduces to an $\text{O}(p, q)$ equivariance requirement. Previous work by Weiler et al. (2023) showed that to
 83 achieve this, the convolution kernels $K : \mathbb{R}^{p,q} \rightarrow \text{Hom}_{\text{Vec}}(\text{Cl}(\mathbb{R}^{p,q})^{c_{\text{in}}}, \text{Cl}(\mathbb{R}^{p,q})^{c_{\text{out}}})$ need to obey
 84 the $\text{O}(p, q)$ -steerability constraint $K(gx) = \rho_{\text{out}}(g) K(x) \rho_{\text{in}}(g)^{-1}$, $\forall g \in \text{O}(p, q)$, $x \in \mathbb{R}^{p,q}$.
 85

86 Several approaches have been proposed to solve this constraint, for example (Lang and Weiler, 2020;
 87 Cesa et al., 2022) analytically derived steerable kernel bases. Kernels can also be implemented as
 88 channel permutations in the matrix dimensions, such as in (Cohen and Welling, 2016; Bekkers et al.,
 89 2018). Both methods however are only suitable for compact groups, leaving out $\text{O}(p, q)$. Numerical
 90 approaches have been proposed by (De Haan et al., 2020; Shutty and Wierzynski, 2023) based on Lie-
 91 algebra irreps, however only for connected subgroups of $\text{O}(p, q)$. (Zhdanov et al., 2023) solved the
 92 steerability constraint *implicitly* by parametrising the kernel using an equivariant neural network. This
 93 method was adapted to work on Clifford algebras, by using *Clifford Group Equivariant neural net-*
 94 *works* (Ruhe et al., 2023) to parametrise the kernels of *Clifford Steerable CNNs* (Zhdanov et al., 2024).
 95

96 The implicit kernel K serving as the convolution kernel of *Clifford Steerable CNN* is a composition
 97 of two operators: $K = H \circ \mathcal{K}$, where \mathcal{K} is the parametrising *Kernel Network*, and H denotes a *kernel*
 98 *head*, which transforms the kernel network output to be used as a convolution kernel on multivector
 99 fields.

100 **Definition 2.2** (Kernel network). The Kernel network $\mathcal{K} : \mathbb{R}^{p,q} \rightarrow \text{Cl}(\mathbb{R}^{p,q})^{c_{\text{out}} \times c_{\text{in}}}$ is an $\text{O}(p, q)$ -
 101 equivariant *Clifford-group equivariant neural network* (CGENN) (Ruhe et al., 2023). Let $k = 0, \dots, d$
 102 denote the multivector grades and $w_{mn}^{(k)} \in \mathbb{R}$ weights. A CGENN is built only from operations that
 103 are $\text{O}(p, q)$ -equivariant by construction:

104 1. *Grade-wise linear maps* $L_m^{(k)}(x_1, \dots, x_{c_{\text{in}}}) := \sum_{n=1}^{c_{\text{in}}} w_{mn}^k \cdot x_n^{(k)}$, which act independently
 105 inside every irreducible sub-representation (grade) of the respective Clifford metric space
 106 $\text{Cl}(\mathbb{R}^{p,q})^{(k)}$;

107 2. *Weighted geometric-product layers* $P^{(k)}(x_1, x_2) := \sum_{m=0}^d \sum_{n=0}^d w_{mn}^k \cdot (x_1^{(m)} \bullet x_2^{(n)})^{(k)}$

108 Additionally, the network contains component-wise nonlinearities and norm-based normalisations
 109 that depend only on the quadratic form and therefore remain invariant under the group action.
 110 Because every CGENN layer is a multivector *polynomial*, the universal result that “all multivector
 111 polynomials are $O(p, q)$ -equivariant” applies (Ruhe et al., 2023), which guarantees that their
 112 composition in \mathcal{K} satisfies $\mathcal{K}(gv) = \rho_{Cl(\mathbb{R}^{p,q})}(g)\mathcal{K}(v)$ for all $g \in O(p, q)$.
 113

114 **Definition 2.3** (Kernel head). The kernel head: $H : Cl(\mathbb{R}^{p,q})^{c_{\text{out}} \times c_{\text{in}}} \rightarrow$
 115 $\text{Hom}_{\text{Vec}}(Cl(\mathbb{R}^{p,q})^{c_{\text{in}}}, Cl(\mathbb{R}^{p,q})^{c_{\text{out}}})$ is a grade-projected, partially evaluated geometric product
 116 that turns each multivector component $k = [k_{ij}]$ into an \mathbb{R} -linear map between channel
 117 stacks:

$$[H(k) f]_i^{(k)} = \sum_{j \in [c_{\text{in}}]} \sum_{m,n=0}^d w_{mn,ij}^{(k)} (k_{ij}^{(m)} \bullet f_j^{(n)})^{(k)}.$$

118 The scalars $w_{mn,ij}^{(k)}$ (shared over space) merely re-weight already equivariant terms; hence H itself
 119 is $O(p, q)$ -equivariant. A proof can be found in Appendix E.1 of Zhdanov et al. (2024). Since both
 120 \mathcal{K} and H are $O(p, q)$ -equivariant, their composition $K = H \circ \mathcal{K}$ obeys the steerability constraint.
 121 Convoluting with K therefore yields an operator that is exactly $E(p, q)$ -equivariant (Theorem 3.4 in
 122 Zhdanov et al. (2024)).

123 3 Missing basis of regular CSCNN kernels

124 This section takes a representation theoretic investigation of the original CSCNN kernel implementa-
 125 tion by Zhdanov et al. (2024), uncovering the missing kernel basis. To overcome this limitation,
 126 we introduce *Conditioned-Clifford Steerable Kernels*, that recover the missing kernels without
 127 breaking equivariance, allowing for the full parametrization of the convolution kernels. Note that
 128 here we restrict our attention to $O(2)$ acting on \mathbb{R}^2 , however, the same decomposition can be used to
 129 generalise the analysis to $O(n)$ and $O(p, q)$.
 130

131 To understand what is meant by angular frequency 2 kernels, it is worth investigating the kernel
 132 functions in the frequency domain. We can do this, as it was shown by Weiler and Cesa (2019);
 133 Weiler et al. (2018), the irreducible representations of $O(2)$ are one- or two-dimensional circular
 134 harmonic modes. More precisely, every kernel obeying the steerability constraint admits a polar
 135 decomposition

$$K(r, \varphi) = \sum_{m \in \mathbb{Z}} R_m(r) Y_m(\varphi), \quad Y_m(\varphi) = e^{im\varphi}.$$

136 Where m is the grade (irrep) of the multivector output and where the angular factors Y_m are fixed and
 137 only the radial parts are freely learned R_m . As such, the kernel can be fully expressed by the different
 138 angular frequency bases and their learned radial part. Depending on the field types a convolution
 139 operator is mapping between, its kernel is required to contain specific angular frequency bases for the
 140 mapping to be fully complete.

141 3.1 Clebsch–Gordan requirements for kernels

142 To identify what angular frequencies are required for convolutions, we use the *Clebsch–Gordan*
 143 decomposition of its tensor products, which was adapted by Lang and Weiler (2020) for kernel
 144 irreps. When a kernel maps one field type (irrep j) to another field type (irrep l), the tensor product
 145 $D^{(j)} \otimes D^{(l)}$ must supply the trivial representation so that contraction with an input feature can
 146 yield an output feature. For multivector grades $Cl(\mathbb{R}^{2,0})^{(k)}$ the Clebsch–Gordan decomposition
 147 $D^{(j)} \otimes D^{(l)} \cong \bigoplus_{m=|j-l|}^{j+l} D^{(m)}$ is shown in Table 1

148 Examining $j = l = 1$ (so the vector field \rightarrow vector field mapping) yields the set of required harmonic
 149 blocks $m \in \{|1-1|, 1+1\} = \{0, 2\}$. Which means that if we want our mapping from vector
 150 \rightarrow vector to be complete, we need bases with angular frequencies $m \in \{0, 2\}$ to be represented in the
 151 kernel. The scalar part $m = 0$ can be recovered easily: it already “lives” in grades 0 and 2. However,

Table 1: Clifford field types by grade and $O(2)$ irreducible representations

grade	field type	irrep	angular frequency
$\text{Cl}(\mathbb{R}^{2,0})^{(0)}$	scalar	$D^{(0)}$	$m = 0$ (even)
$\text{Cl}(\mathbb{R}^{2,0})^{(1)}$	vector	$D^{(1)}$	$m = 1$
$\text{Cl}(\mathbb{R}^{2,0})^{(2)}$	pseudoscalar	$D_{\text{odd}}^{(0)}$	$m = 0$ (odd)

152 the quadrupole part with angular frequency $m = 2$ can appear only through the symmetric traceless
 153 component of a product of two *independent* vectors (Zou and Zheng, 2003).

154
 155 The implicit kernel of a Clifford-Steerable CNN, $K = H \circ \mathcal{K}$, receives a *single* grade-1 vector per
 156 edge, namely the displacement $v := x - y \in \mathbb{R}^2 \cong \text{Cl}(\mathbb{R}^2)^{(1)}$ (Zhdanov et al., 2023, 2024). Inside
 157 of kernel network \mathcal{K} , the transformations applied to this vector (as detailed in Definition 2.2) only do
 158 the following:

159 1. *Grade-wise linear maps* reshuffle channels but never change the spatial irrep: a vector stays
 160 a vector ($m = 1$).
 161 2. *Geometric products* take two copies of the same vector. Their product immediately projects
 162 (due to the Clifford rules in Equation (1)) into grade 0 or grade 2, i.e. frequency $m = 0$.

163 Because both operands derive from the same direction, no symmetric-traceless rank-2 object-and
 164 hence no $m = 2$ harmonic-can be formed at any depth inside a single kernel network. This is not
 165 changed by the kernel head H either, since its operation is also a (partially evaluated) geometric
 166 product. This results in the incomplete vector \rightarrow vector mapping of the convolutional layers, and
 167 the original CSCNN architecture. To complete the parametrisation, the interaction of independent
 168 vectors are needed.

169 4 Conditioned Clifford-Steerable Kernels

170 4.1 Conditioned Steerable Convolution

171 We follow the derivation of a steerable convolution from Weiler et al. (2023) (Theorem 4.3.1). Let us
 172 start with a convolutional operator

$$I_k [F] (x) := \int_{\mathbb{R}^d} dy K(x - y) F(y) \quad (2)$$

173 that is parameterized by a kernel $K : \mathbb{R} \rightarrow \mathbb{R}^{c_{\text{out}} \times c_{\text{in}}}$. For the convolutional layer to be G -equivariant,
 174 the kernel must be G -steerable, that is,

$$K(gx) = \rho_{\text{out}}(g)K(x)\rho_{\text{in}}(g)^{-1} \quad \forall x \in \mathbb{R}^d, g \in G. \quad (3)$$

175 where ρ_{in} and ρ_{out} are representations of input and output feature fields, respectively.

176 Let us now introduce an equivariant function $T : L^2(\mathbb{R}^d, \mathbb{R}^c) \rightarrow \mathbb{R}^c$ which computes a vector
 177 representation of a feature field. We will use the function to augment the kernel function, yielding the
 178 following quasi-linear transformation

$$I_k^T [F] (x) := \int_{\mathbb{R}^d} dy K(x - y, T[F]) F(y). \quad (4)$$

179 Since the conditioning function T is equivariant, we have by definition $T[g.F] = \rho(g)T[F]$.

180 *Lemma 1.* The quasi-linear map 4 is G -equivariant if and only if

$$K(gx, \rho_{\text{in}}(g)T[F]) = \rho_{\text{Hom}}(g)(K(x, T[f_{\text{in}}]))$$

181 We provide the proof of the lemma in the Appendix B.1.

182 **4.2 Conditioned Clifford-Steerable Convolution**

183 The quasi-linear map provides a framework to tackle the issue of missing degrees of freedom in
184 Clifford-Steerable convolutions. As discussed, the kernel basis of the latter is incomplete, as it relies
185 on implicit kernels that are unable to generate high-frequency components from a positional vector
186 alone. Furthermore, the key idea of our approach is to condition the implicit kernel on relative
187 position, as well as a field representation, whose interaction together will yield the otherwise missing
188 high-frequency components of the kernel basis.

189 As discussed in Section 2.3, in the framework of Clifford-Steerable CNNs, the kernel K is defined
190 as a composition $K = H \circ \mathcal{K}$. The kernel head H transforms the output of the kernel network \mathcal{K} to
191 matrix-valued kernels and is agnostic to the form of \mathcal{K} . Furthermore, we only need to adapt the kernel
192 network to be compatible with conditioned convolution. Specifically, let us define conditioned kernel
193 network $\mathcal{K}_T : \mathbb{R}^{p,q}, \text{Cl}(\mathbb{R}^{p,q})^{c_{\text{in}}} \rightarrow \text{Cl}(\mathbb{R}^{p,q})^{c_{\text{out}} \times c_{\text{in}}}$ which now additionally takes the multivector
194 representation of the input field generated by the conditioning function $T : L^2(\mathbb{R}^d, \text{Cl}(\mathbb{R}^{p,q})^c) \rightarrow$
195 $\text{Cl}(\mathbb{R}^{p,q})^c$. Since we use CGENNs to implement \mathcal{K}_T , the function is $O(p, q)$ -equivariant by definition.
196 Having the conditioned kernel network, we define Conditioned-Steerable kernels as $K = H \circ \mathcal{K}_T$.

197 *Lemma 2* (Equivariance of conditioned Clifford-steerable kernels). Every conditioned Clifford-
198 steerable kernel $K = H \circ \mathcal{K}_T$ is $O(p, q)$ -steerable w.r.t. the standard action $\rho(g) = g$ and ρ_{Hom} :

$$K(gv, \rho_{\text{in}}(g)T[F]) = \rho_{\text{Hom}}(g)(K(v, T[f_{\text{in}}])) \quad \forall g \in O(p, q), v \in \mathbb{R}^{p,q}.$$

199 We provide the proof in Appendix B.2.

200 Equipped with conditioned kernels, we are now able to tackle the issue with missing degrees of
201 freedom in the original framework of CSCNNs:

202 **Corollary 4.1** (Completeness of the kernel basis). Let the output of $T[f_{\text{in}}]$ have non-trivial grade 1
203 component. Then the kernel basis of conditioned Clifford-Steerable CNNs is complete.

204 **5 Experimental Results**

205 To test the efficacy of our new kernel implementation, we compared its performance to the original
206 CSCNN model and several strong baselines in three PDE learning tasks: the 2-dimensional (\mathbb{R}^2)
207 *Navier-Stokes* (*NS2*) equations, the 3-dimensional (\mathbb{R}^3) *Maxwell* (*MW3*) equations, and the relativistic
208 2-dimensional ($\mathbb{R}^{1,2}$) *Maxwell* (*MW2*) equations.

209 **5.1 Notes on implementation**

210 For the simplicity, we choose the conditioning function T to be the global mean pooling operation.
211 The operation is equivariant, as proven in Weiler et al. (2023). Compared to the original CSCNNs,
212 our model’s only additional computational cost comes from computing the condition. However, since
213 this step uses highly optimized operations (like mean and max pooling), the resulting overhead is
214 minimal. The implementation is done in JAX (Bradbury et al., 2018). More details can be found in
215 Appendix A.1.

216 **5.2 Experiment Setup**

217 The goal of each experiment is to empirically learn the dynamics of the systems from numerical
218 simulations, and predict future states based on previous timesteps (Gupta and Brandstetter, 2022).
219 For the non-relativistic tasks, variables describing the states in the 4 previous timesteps are taken
220 as input, with each timestep serving as an additional feature channel, and the goal is to predict the
221 subsequent state. Although improper, this setup allows us to compare our results to prior models
222 which otherwise would not be able to handle a full spacetime. For Navier-Stokes, the input state
223 is described by a velocity vector for the vector part, and the pressure field for the scalar part, with
224 the bivector part padded with zeros to keep the multivector structure consistent. For the MW3 task,
225 the inputs are vector electric fields and bivector magnetic fields, with the rest of the blades padded.
226 For the relativistic Maxwell task, time forms its own spacetime dimension of size 32, thus the task
227 becomes predicting the next 32 timesteps based on the previous 32. The input is an electromagnetic
228 field, which forms the bivector part of a multivector. This setup properly incorporates the time

229 dimension into spacetime.

230

231 5.3 Model Architectures

232 Our model is built upon the ResNet architecture, where we substitute the standard convolutional layers
 233 with our novel Conditioned Clifford-Steerable convolutions. We compare Conditioned CSCNNs
 234 against architecturally similar networks: the standard ResNet, the Clifford ResNet (Brandstetter
 235 et al., 2022), the $O(n)$ -steerable CNN (Weiler and Cesa, 2019; Weiler et al., 2023) and the original
 236 Clifford-Steerable CNN (Zhdanov et al., 2024). Additionally, we evaluate the prominent PDE learner
 237 Fourier Neural Operator (FNO) (Li et al., 2021) and its $D_4 < O(2)$ equivariant variation, G-FNO
 238 (Helwig et al., 2023). All models were scaled to match the parameter count of the basic ResNet
 239 architecture.

240 5.4 Results

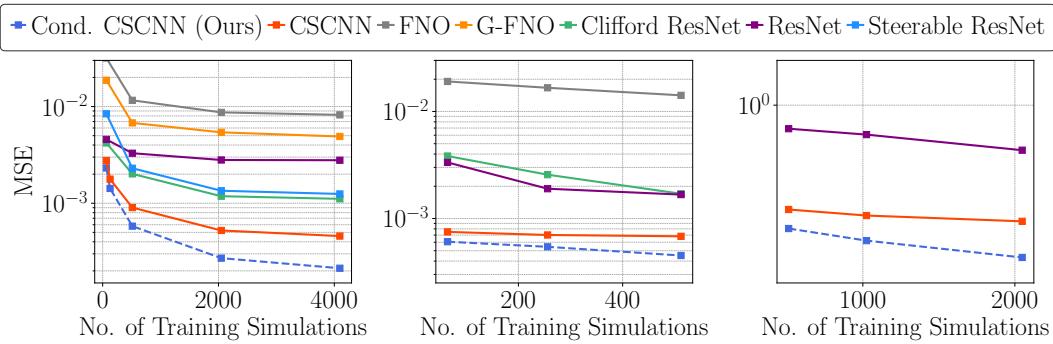


Figure 3: Mean squared errors for (1) Navier-Stokes \mathbb{R}^2 , (2) Maxwell \mathbb{R}^3 , and (3) relativistic Maxwell $\mathbb{R}^{1,2}$ simulation tasks as a function of the simulations included in the training dataset. Conditioned-CSCNNs outperform all baselines, with the gap widening as data increases.

241 Figure 3 shows the MSE on the test set. On all tasks, our *Conditioned-CSCNN* outperforms
 242 every baseline, including the original CSCNN implementation. Its advantage becomes more
 243 pronounced as more simulations are included in the training set, which is an experimental proof
 244 of the added model complexity from the recovered angular frequency 2 kernels. Indeed, as shown
 245 in Figure 4, the highest MSE improvement against the baseline CSCNN comes from the vector
 246 component of the loss, suggesting complete $\text{vector} \rightarrow \text{vector}$ mappings of the convolution layers. Note
 247 that we can only test this for the NS experiment, where we have ground truth data for the vector output.
 248

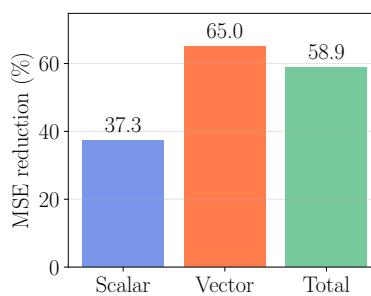


Figure 4: Per-grade MSE reduction of our model vs. baseline CSCNN for the Navier-Stokes experiment with $n = 2048$ simulations in the training dataset. We see the largest improvement in the vector component.

249 To test whether our model is still equivariant, we calculated relative equivariance error between the
 250 output of the transformed input vs the transformed output of the original input: $\text{err}(f; g, x) =$

Table 2: Relative equivariance errors for Clifford-Steerable convolutions

Kernel	Relative error (mean)	Relative error (max)
Default CS Convolution	2.4×10^{-7}	1.1×10^{-3}
Conditioned CS Convolution	3.4×10^{-7}	5.9×10^{-4}

251 $\frac{|f(gx) - gf(x)|}{|f(gx) + gf(x)|}$. Table 2 shows the $O(2)$ relative errors for the original and our conditioned convolu-
 252 tions. While the error is slightly higher for the conditioned convolutions, it is still equivariant up to
 253 numerical artifacts.

254 6 Conclusion

255 We introduce Conditioned Clifford-Steerable CNNs, a generalization of Clifford-Steerable CNNs
 256 that employs quasi-linear convolutions. By conditioning the convolution kernel on representations
 257 of the input feature field, our method recovers the degrees of freedom that were absent in the
 258 original framework. Conditioned Clifford-Steerable CNNs can therefore express complete mappings
 259 between multivector fields on (pseudo-)Euclidean spaces while maintaining consistent transformation
 260 properties under the general Euclidean group $E(p, q)$. This capability makes our approach particularly
 261 well-suited for learning the physical dynamics of both relativistic and non-relativistic systems.
 262 We demonstrate this advantage across various PDE learning tasks, where Conditioned CSCNNs
 263 consistently outperform strong baselines, including the original CSCNN implementation.

264 **Limitations** Our work has two primary limitations. First, we analytically derive the missing basis
 265 kernels only for the special case of $O(2, 0)$, leaving the proof for the general case of $O(p, q)$ as future
 266 work. However, our empirical results demonstrate that addressing this limitation yields dramatically
 267 improved performance across all tasks, indicating a more expressive underlying transformation.
 268 Second, our implementation has significantly slower runtimes than baseline CSCNNs for higher-
 269 dimensional groups, which we attribute to suboptimal optimization of conditioned CNNs. We
 270 anticipate that more efficient implementations—such as those utilizing Triton—could resolve these
 271 computational bottlenecks.

272 **Impact Statement** Our approach has broad applications in equivariant physical modeling, including
 273 neural PDE surrogates, weather forecasting, and fluid dynamics simulation. By incorporating known
 274 symmetries into deep learning models, one can improve data efficiency and robustness to domain
 275 shifts, which is important in real-world applications. Additionally, we also hope to apply the
 276 framework to robotics, where exploiting symmetries can dramatically improve sample efficiency
 277 when training a physical robotic system Wang et al. (2022).

278 **References**

279 Bekkers, E. J., Lafarge, M. W., Veta, M., Eppenhof, K. A., Pluim, J. P., and Duits, R. (2018).
280 Roto-translation covariant convolutional networks for medical image analysis. In *Medical Image
281 Computing and Computer Assisted Intervention—MICCAI 2018: 21st International Conference,
282 Granada, Spain, September 16–20, 2018, Proceedings, Part I*, pages 440–448. Springer.

283 Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke,
284 A., VanderPlas, J., Wanderman-Milne, S., and Zhang, Q. (2018). JAX: composable transformations
285 of Python+NumPy programs.

286 Brandstetter, J., Berg, R. v. d., Welling, M., and Gupta, J. K. (2022). Clifford neural layers for pde
287 modeling. *arXiv preprint arXiv:2209.04934*.

288 Cesa, G., Lang, L., and Weiler, M. (2022). A program to build e (n)-equivariant steerable cnns. In
289 *International conference on learning representations*.

290 Cohen, T. and Welling, M. (2016). Group equivariant convolutional networks. In *International
291 conference on machine learning*, pages 2990–2999. PMLR.

292 De Haan, P., Weiler, M., Cohen, T., and Welling, M. (2020). Gauge equivariant mesh cnns:
293 Anisotropic convolutions on geometric graphs. *arXiv preprint arXiv:2003.05425*.

294 Filipovich, M. J. and Hughes, S. (2022). Pycharge: An open-source python package for self-consistent
295 electrodynamics simulations of lorentz oscillators and moving point charges. *Computer Physics
296 Communications*, 274:108291.

297 Geiger, M., Smidt, T., Alby, M., Miller, B. K., Boomsma, W., Dice, B., Lapchevskyi, K., Weiler,
298 M., Tyszkiewicz, M., Batzner, S., et al. (2022). Euclidean neural networks: e3nn. *Preprint at
299 https://doi. org/10.48550/arXiv*, 2207.

300 Gupta, J. K. and Brandstetter, J. (2022). Towards multi-spatiotemporal-scale generalized pde
301 modeling. *arXiv preprint arXiv:2209.15616*.

302 Helwig, J., Zhang, X., Fu, C., Kurtin, J., Wojtowytch, S., and Ji, S. (2023). Group equivariant
303 fourier neural operators for partial differential equations. In *Proceedings of the 40th International
304 Conference on Machine Learning*, ICML’23. JMLR.org.

305 Holl, P., Koltun, V., and Thuerey, N. (2020). Learning to control pdes with differentiable physics.
306 *CoRR*, abs/2001.07457.

307 Kingma, D. P. and Ba, J. (2017). Adam: A method for stochastic optimization.

308 Lang, L. and Weiler, M. (2020). A wigner-eckart theorem for group equivariant convolution kernels.
309 *arXiv preprint arXiv:2010.10952*.

310 Li, Z., Kovachki, N. B., Azizzadenesheli, K., liu, B., Bhattacharya, K., Stuart, A., and Anandkumar,
311 A. (2021). Fourier neural operator for parametric partial differential equations. In *International
312 Conference on Learning Representations*.

313 O’Neill, B. (1983). Semi-riemannian geometry with applications to relativity. *Pure and Applied
314 Mathematics/Academic Press, Inc.*

315 Ruhe, D., Brandstetter, J., and Forré, P. (2023). Clifford group equivariant neural networks. *Advances
316 in Neural Information Processing Systems*, 36:62922–62990.

317 Shutty, N. and Wierzynski, C. (2023). Computing representations for lie algebraic networks. In
318 *NeurIPS Workshop on Symmetry and Geometry in Neural Representations*, pages 1–21. PMLR.

319 Wang, D., Jia, M., Zhu, X., Walters, R., and Platt, R. W. (2022). On-robot learning with equivariant
320 models. In *Conference on Robot Learning*.

321 Wang, R., Walters, R., and Yu, R. (2021). Incorporating symmetry into deep dynamics models for
322 improved generalization. In *International Conference on Learning Representations*.

323 Weiler, M. and Cesa, G. (2019). General e (2)-equivariant steerable cnns. *Advances in neural*
324 *information processing systems*, 32.

325 Weiler, M., Forré, P., Verlinde, E., and Welling, M. (2023). Equivariant and coordinate independent
326 convolutional networks. *A Gauge Field Theory of Neural Networks*, page 110.

327 Weiler, M., Geiger, M., Welling, M., Boomsma, W., and Cohen, T. (2018). 3d steerable cnns:
328 Learning rotationally equivariant features in volumetric data. *CoRR*, abs/1807.02547.

329 Zhdanov, M., Hoffmann, N., and Cesa, G. (2023). Implicit convolutional kernels for steerable cnns.
330 *Advances in Neural Information Processing Systems*, 36:17395–17407.

331 Zhdanov, M., Ruhe, D., Weiler, M., Lucic, A., Brandstetter, J., and Forré, P. (2024). Clifford-steerable
332 convolutional neural networks. *arXiv preprint arXiv:2402.14730*.

333 Zou, D. and Zheng, Q.-s. (2003). Maxwell’s multipole representation of traceless symmetric tensors
334 and its application to functions of high-order tensors. *Proceedings of The Royal Society A:
335 Mathematical, Physical and Engineering Sciences*, 459:527–538.

336 **A Appendix**

337 **A.1 Implementation details**

338 The basic ResNet architecture that was used for constructing the Conditioned-CSCNN, the CSCNN,
 339 the Clifford ResNet and the $O(n)$ -steerable ResNet were based on the setup of Wang et al. (2021);
 340 Brandstetter et al. (2022); Gupta and Brandstetter (2022). They consist of 8 residual blocks with 7×7
 341 and $7 \times 7 \times 7$ sized kernels for the 2D and 3D experiments respectively. We used two embedding
 342 and two output layers. For constructing the FNO (Li et al., 2021) and G-FNO (Helwig et al., 2023),
 343 we followed the original implementations. The models have approximately 7M parameters for the
 344 Navier Stokes and 1.5M for Maxwell's experiments.

345 **A.1.1 Conditioned CS kernel implementation**

346 In constructing the Conditioned Clifford-Steerable Kernels, we built on the architecture described in
 347 Appendix A of Zhdanov et al. (2024). We form the conditioning vector by a *masked spatial mean*
 348 computed for each channel c and blade/grade k . On a grid $\Omega \subset \mathbb{R}^d$ and with the indicator χ_{B_r} of the
 349 largest centered ball B_r (the circular/spherical mask, which we explain below), the pooled stack is

$$(T_{\text{pool}}[f_{\text{in}}])_c^{(k)} := \frac{1}{|\Omega|} \sum_{x \in \Omega} \chi_{B_r}(x) f_{\text{in},c}^{(k)}(x),$$

350 The resulting conditioning multivector stack is then concatenated with the relative position vector to
 351 form the input to the Kernel Network.

352 **Circular mask** : In practice, the multivector feature fields are often discretised as square shaped
 353 arrays $(c, X_1 \dots X_d, 2^d)$. Thus, an operation defined solely on this finite grid will break equivariance
 354 towards the corners of the domain, as shown in Figure 5. To overcome this, we apply circular
 355 masking (set grid values outside of the circle/sphere to 0) before the pooling operation, making it
 356 $O(n)$ -equivariant in the continuum.

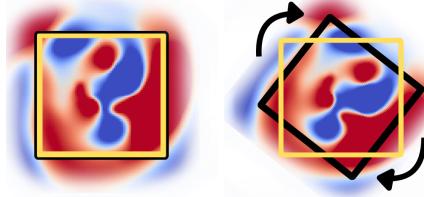


Figure 5: Illustration of how the receptive field of the finite, discretised kernel changes under rotations. The square shape causes any operation defined on it to break equivariance towards the corners.

357 **A.1.2 Training details**

358 We adopted the optimised hyperparameters from Zhdanov et al. (2024) for all of our models and
 359 experiments. We used Adam optimizer (Kingma and Ba, 2017) with cosine learning rate scheduler.
 360 Each model was trained to convergence. The models were trained on one node of Snellius, the Dutch
 361 national computing cluster with 4 NVIDIA A100 GPUs.

362 **A.1.3 Datasets**

363 **2D Navier Stokes equations** The ground truth simulations are taken from Gupta and Brandstetter
 364 (2022) and are based on Φ Flow by Holl et al. (2020). From the corresponding validation and test
 365 partitions, we randomly sampled 1024 trajectories. The simulations were generated on a 128×128
 366 pixel grid with uniform spatial spacing $\Delta x = \Delta y = 0.25$ m and a time step of $\Delta t = 1.5$ s.

367 **3D Maxwell's equations** Within the non-relativistic $\text{Cl}(\mathbb{R}^{3,0})$ setting, we represent the electric field
 368 E as a vector field and the magnetic field B as a bivector field. The dataset, drawn from Brandstetter

369 et al. (2022), consists of 3D Maxwell simulations discretized on a $32 \times 32 \times 32$ voxel grid with
 370 uniform spacing $\Delta x = \Delta y = \Delta z = 5 \times 10^{-7}$ m and time step $\Delta t = 50$ s. The validation and test
 371 splits together contain 128 simulations.

372 **2D relativistic Maxwell's equations** We generate a dataset for Maxwell's equations in 2+1D
 373 spacetime ($\mathbb{R}^{1,2}$) utilizing the PyCharge simulation package by Filipovich and Hughes (2022). The
 374 simulations model the dynamics of electromagnetic fields emitted by oscillating and orbiting point
 375 charges moving at relativistic speeds. The spacetime grid is discretized with a resolution of 128×128
 376 points, corresponding to a spatial extent of 50 nm and a temporal duration of $3.77 \cdot 10^{-14}$ s.

377 Each simulation is initialized with a unique configuration of charge sources, governed by the following
 378 randomly sampled parameters: **Source Composition:** A combination of 2 to 4 oscillating charges
 379 and 1 to 2 orbiting charges, with integer magnitudes sampled uniformly from the range $[-3e, 3e]$.
 380 **Initial Conditions:** Sources are placed uniformly on the grid with a predefined minimum separation.
 381 Each is assigned a random linear velocity and either oscillates in a random direction or orbits with a
 382 random radius. **Relativistic Constraint:** Oscillation/rotation frequencies and velocities are sampled
 383 such that the total particle velocity does not exceed $0.85c$, a necessary constraint to ensure the stability
 384 of the PyCharge solver.

385 To handle the wide dynamic range of the resulting field strengths, we apply a normalization scheme.
 386 The generated field bivectors are divided by their Minkowski norm and then multiplied by the
 387 logarithm of that norm. Although Minkowski norms can be zero or negative, we found they were
 388 consistently positive in our generated data. Finally, we filter numerical artifacts by removing any
 389 outlier simulations that exhibit a standard deviation greater than 20. The curated dataset is split into
 390 2048 training, 256 validation, and 256 test simulations.

391 B Proofs

392 B.1 Proof of Lemma 1

393 *Proof.* We seek to proof the G -equivariance of the quasi-linear convolution

$$I_k^T [F] (x) := \int_{\mathbb{R}^d} dy K(x - y, T[F]) F(y). \quad (5)$$

394 where the conditioning function T is G -equivariant. We will use the notation of Weiler et al. (2023),
 395 where the action of a group G on a feature vector field of type ρ is written as

$$[tg \triangleright_\rho F] (x) = \rho(g) F((tg)^{-1}x). \quad (6)$$

396 The G equivariance of the convolutional layer then formally means that

$$I_k^T [g \triangleright_{\rho_{\text{in}}} F] = g \triangleright_{\rho_{\text{out}}} I_k^T [F], \quad \forall g \in G \quad (7)$$

397 where ρ_{in} and ρ_{out} stand for the type of input and output feature fields, respectively.

398 By expanding the left-hand side of the equation above, we obtain

$$I_k^T [g \triangleright_{\rho_{\text{in}}} F] = \int_{\mathbb{R}^d} dy K(x - y, T[g \triangleright_\rho F]) [g \triangleright_\rho F](y) \quad (8)$$

$$\stackrel{T \text{ equiv.}}{=} \int_{\mathbb{R}^d} dy K(x - y, \rho_{\text{in}}(g) T[F]) \rho_{\text{in}} F(g^{-1}y) \quad (9)$$

$$\stackrel{g^{-1}y \rightarrow y}{=} \int_{\mathbb{R}^d} dy K(x - gy, \rho_{\text{in}}(g) T[F]) \rho_{\text{in}} F(y) \quad (10)$$

399 The right-hand side then yields

$$g \triangleright_{\rho_{\text{out}}} I_k^T [F] = \int_{\mathbb{R}^d} dy \rho_{\text{out}}(g) K(g^{-1}x - y, T[F]) F(y) \quad (11)$$

400 These expressions agree for any $g \in G$ and any feature map $F \in L^2(\mathbb{R}^d, \mathbb{R}^{c_{\text{in}}})$ if and only if

$$K(x - gy, \rho_{\text{in}}(g) T[F]) \rho_{\text{in}} = \rho_{\text{out}}(g) K(g^{-1}x - y, T[F]). \quad (12)$$

401 After substitution $g^{-1}x \rightarrow x$ and $x - y = x$, we obtain the constraint:

$$K(gx, \rho_{\text{in}}(g)T[F]) = \rho_{\text{out}}(g)K(x, T[F])\rho_{\text{in}}(g)^{-1} \equiv \rho_{\text{Hom}}(g)(K(x, T[F])) \quad (13)$$

402 \square

403 **B.2 Proof of Lemma 2**

404 *Proof.* \mathcal{K}_T is $O(p, q)$ -equivariant by definition, and the kernel head H is $O(p, q)$ -equivariant by
405 Proposition 3.2 in Zhdanov et al. (2024). The $O(p, q)$ -equivariance of their composition follows from

$$K(gv, \rho_{Cl}^{c_{\text{in}}}(g)T[f_{\text{in}}]) = H(\mathcal{K}_T(gv, \rho_{Cl}^{c_{\text{in}}}(g)T[f_{\text{in}}])) \quad (14)$$

$$\stackrel{\mathcal{K}_T \text{ equiv.}}{=} H(\rho_{Cl}^{c_{\text{out}} \times c_{\text{in}}}(g)\mathcal{K}_T(v, T[f_{\text{in}}])) \quad (15)$$

$$= \rho_{\text{Hom}}(g)H(\mathcal{K}_T(v, T[f_{\text{in}}])) \quad (16)$$

$$= \rho_{\text{Hom}}(g)K(v, T[f_{\text{in}}]) \quad (17)$$

406 \square