

Quaternion Product Units for Deep Learning on 3D Rotation Groups

Xuan Zhang¹, Shaofei Qin¹, Yi Xu^{*1}, and Hongteng Xu²

¹MoE Key Lab of Artificial Intelligence, AI Institute ²Infinia ML, Inc.

¹Shanghai Jiao Tong University, Shanghai, China ²Duke University, Durham, NC, USA

{floatlazer, 1105042987, xuyi}@sjtu.edu.cn, hongteng.xu@duke.edu

Abstract

We propose a novel quaternion product unit (QPU) to represent data on 3D rotation groups. The QPU leverages quaternion algebra and the law of 3D rotation group, representing 3D rotation data as quaternions and merging them via a weighted chain of Hamilton products. We prove that the representations derived by the proposed QPU can be disentangled into “rotation-invariant” features and “rotation-equivariant” features, respectively, which supports the rationality and the efficiency of the QPU in theory. We design quaternion neural networks based on our QPUs and make our models compatible with existing deep learning models. Experiments on both synthetic and real-world data show that the proposed QPU is beneficial for the learning tasks requiring rotation robustness.

1. Introduction

Representing 3D data like point clouds and skeletons is essential for many real-world applications, such as autonomous driving [17, 37], robotics [30], and gaming [2, 16]. In practice, we often model these 3D data as the collection of points on a 3D rotation group $\mathbb{SO}(3)$, *e.g.*, a skeleton can be represented by the rotations between adjacent joints. Accordingly, the uncertainty hidden in these data is usually caused by the randomness on rotations. For example, in action recognition, the human skeletons with different orientations may represent the same action [27, 35]; in autonomous driving, the point clouds with different directions may capture the same vehicle [37]. Facing to the 3D data with such rotation discrepancies, we often require the corresponding representation methods to be robust to rotations.

Currently, many deep learning-based representation

methods have made efforts to enhance their robustness to rotations [4, 6, 29, 3, 26]. However, the architectures of their models are tailored for the data in the Euclidean space, rather than 3D rotation data. In particular, the 3D rotation data are not closed under the algebraic operations used in their models (*i.e.*, additions and multiplications).¹ The mismatching between data and model makes these methods difficult to analyze the influence of rotations on their outputs quantitatively. Although this mismatching problem can be mitigated by augmenting training data with additional rotations [18], this solution gives us no theoretical guarantee.

To overcome the challenge mentioned above, we proposed a novel quaternion product unit (QPU) for 3D rotation data, which establishes an efficient and interpretable mechanism to enhance the robustness of deep learning models to rotations. As illustrated in Figure 1a, for each 3D rotation, we represent it as a unit quaternion, whose imaginary part indicates the direction of its rotation axis, and the real part corresponds to the cosine of its rotation angle, respectively. Taking N quaternions as its inputs, the QPU first applies quaternion power operation to each input, scaling their rotation angles and rotation axes, respectively. Then, it applies a chain of Hamilton products to merge the weighted quaternions and output a rotation accordingly. The parameters of the QPU consists of the scaling coefficients and the bias introduced to the rotation angles.

The proposed QPU leverages quaternion algebra and the law of 3D rotation groups, merging the inputs as one quaternion. Because 3D rotation group is closed under the operations used in our QPU, the output is still a valid rotation. Moreover, given a dataset $\mathcal{X} \subset \mathbb{SO}(3)^N$, where each $\mathbf{x} = [x_1, \dots, x_N] \in \mathcal{X}$ contains N rotations, we can define two kinds of rotation robustness for a mapping function $f : \mathbb{SO}(3)^N \mapsto \mathbb{SO}(3)$:

Rotation-invariance $f(R(\mathbf{x})) = f(\mathbf{x}), \forall \mathbf{x} \in \mathcal{X}$.

Rotation-equivariance $f(R(\mathbf{x})) = R(f(\mathbf{x})), \forall \mathbf{x} \in \mathcal{X}$.

^{*}The corresponding author of this paper is Yi Xu (xuyi@sjtu.edu.cn). This work was supported in part by NSFC 61671298, STCSM 18DZ2270700, and 111 Plan B07022.

¹For example, adding two 3D rotation matrices together or multiplying a scalar with a 3D rotation matrix will not result in a valid rotation matrix.

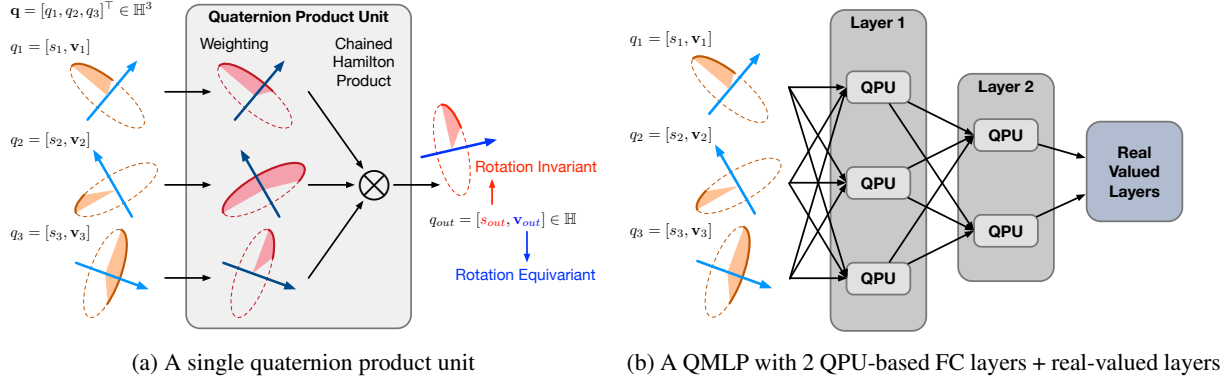


Figure 1: (a) An illustration of our QPU model. Each rotation in $\mathbb{SO}(3)$ can be represented via a unit quaternion in a hypercomplex space \mathbb{H} , whose real part and imaginary part indicates the rotation angle and the direction of the rotation axis, respectively. The QPU merges input rotations via a weighted chain of Hamilton products and derives a rotation as its output, whose real part and imaginary part yield to the rotation-invariance and the rotation-equivariance, respectively. (b) We can further construct a fully-connection (FC) layer using multiple QPUs. The stacking of these QPU-based FC layers leads to a quaternion multi-layer perceptron (QMLP) model, and this model can be followed by existing real-valued layers.

Here, R is an arbitrary rotation operation and $R(x) = [R(x_1), \dots, R(x_N)]$. We prove that the quaternion derived by QPU is disentangled naturally into a *rotation-invariant* real part and a *rotation-equivariant* imaginary part.

The proposed QPU is highly compatible with existing deep learning models. As shown in Figure 1b, we build fully-connected (FC) layers based on QPUs and stack them as a quaternion multi-layer perceptron model (QMLP). We also design a quaternion graph convolutional layer based on QPUs. These QPU-based layers can be combined with standard real-valued layers. Experiments on both synthetic and real-world data show that the proposed QPU-based models can be trained efficiently with standard back-propagation and outperform traditional real-valued models in many applications requiring rotation robustness.

2. Proposed Quaternion Product Units

Many 3D data can be represented as points on 3D rotation groups. Take 3D skeleton data as an example. A 3D skeleton is a 3D point cloud associated with a graph structure. We can represent a skeleton by a set of 3D rotations, using relative rotations between edges (*i.e.*, joint rotations). Mathematically, the 3D rotation from a vector $\mathbf{v}_1 \in \mathbb{R}^3$ to another vector $\mathbf{v}_2 \in \mathbb{R}^3$ could be described as a rotation around an axis \mathbf{u} with a angle θ :

$$\begin{aligned} \text{Rotation Axis: } \mathbf{u} &= \frac{\mathbf{v}_1 \times \mathbf{v}_2}{\|\mathbf{v}_1 \times \mathbf{v}_2\|_2}, \\ \text{Rotation Angle: } \theta &= \arccos\left(\frac{\langle \mathbf{v}_1, \mathbf{v}_2 \rangle}{\|\mathbf{v}_1\|_2 \|\mathbf{v}_2\|_2}\right). \end{aligned} \quad (1)$$

Here, $\mathbf{v}_1 \times \mathbf{v}_2$ stands for cross product and $\langle \mathbf{v}_1, \mathbf{v}_2 \rangle$ is for inner product. Given such 3D rotation data, we would like

to design a representation method yielding to a certain kind of rotation robustness. With the help of quaternion algebra, we propose a quaternion product unit to achieve this aim.

2.1. Quaternion and 3D rotation

Quaternion $q = s + \mathbf{i}x + \mathbf{j}y + \mathbf{k}z \in \mathbb{H}$ is a type of hypercomplex number with 1D real part s and 3D imaginary part (x, y, z) . The imaginary parts satisfy $\mathbf{i}^2 = \mathbf{j}^2 = \mathbf{k}^2 = \mathbf{ijk} = -1$, $\mathbf{ij} = -\mathbf{ji} = \mathbf{k}$, $\mathbf{ki} = -\mathbf{ik} = \mathbf{j}$, $\mathbf{jk} = -\mathbf{kj} = \mathbf{i}$. For convenience, we ignore the imaginary symbols and represent a quaternion as the combination of a scalar and a vector as $q = [s, \mathbf{v}] = [s, (x, y, z)]$. When equipped with Hamilton product as multiplication, noted as \otimes , and standard vector addition, the quaternion forms an algebra. In particular, the Hamilton product between two quaternions $q_1 = [s_1, \mathbf{v}_1]$ and $q_2 = [s_2, \mathbf{v}_2]$ is defined as

$$\begin{aligned} q_1 \otimes q_2 &= [s_1 s_2 - \langle \mathbf{v}_1, \mathbf{v}_2 \rangle, \mathbf{v}_1 \times \mathbf{v}_2 + s_1 \mathbf{v}_2 + s_2 \mathbf{v}_1] \\ &= M_L(q_1)q_2 = M_R(q_2)q_1, \end{aligned} \quad (2)$$

where for a quaternion $q = [s, (x, y, z)]$, we have

$$M_L(q) = \begin{bmatrix} s & -x & -y & -z \\ x & s & -z & y \\ y & z & s & -x \\ z & -y & x & s \end{bmatrix}, M_R(q) = \begin{bmatrix} s & -x & -y & -z \\ x & s & z & -y \\ y & -z & s & x \\ z & y & -x & s \end{bmatrix}. \quad (3)$$

The second row in Eq. (2) uses matrix-vector multiplications, in which we treat one quaternion as a 4D vector and compute a matrix based on the other quaternion. Note that the Hamilton product is non-commutative — multiplying a quaternion on the left or the right gives different matrices.

Quaternion algebra provides us with a new representation of 3D rotations. Specifically, suppose that we rotate a 3D vector \mathbf{v}_1 to another 3D vector \mathbf{v}_2 , and the 3D rotation is with an axis \mathbf{u} and an angle θ shown in Eq. (1).

We can represent the 3D rotation as a unit quaternion $q = [s, \mathbf{v}] = [\cos(\frac{\theta}{2}), \sin(\frac{\theta}{2})\mathbf{u}]$, where $\|\mathbf{u}\|_2 = 1$ and $s^2 + \|\mathbf{v}\|_2^2 = \cos^2(\frac{\theta}{2}) + \sin^2(\frac{\theta}{2}) = 1$. Representing the two 3D vectors as two pure quaternions, *i.e.*, $[0, \mathbf{v}_1]$ and $[0, \mathbf{v}_2]$, we can achieve the rotation from \mathbf{v}_1 to \mathbf{v}_2 via the Hamilton products of the corresponding quaternions:

$$[0, \mathbf{v}_2] = q \otimes [0, \mathbf{v}_1] \otimes q^*, \quad (4)$$

where $q^* = [s, -\mathbf{v}]$ stands for the conjugation of q . Additionally, the combination of rotation matrices can also be represented by the Hamilton products of unit quaternions. For example, given two unit quaternions q_1 and q_2 , which correspond to two rotations, $(q_2 \otimes q_1) \otimes [0, \mathbf{v}_1] \otimes (q_1^* \otimes q_2^*)$ means rotating \mathbf{v}_1 sequentially through the two rotations.

Note that unit quaternion is a double cover of $\mathbb{S}\mathbb{O}(3)$ since q and $-q$ represents the same 3D rotation (in opposite directions). As shown in Eq. (1), we use inner product and arccos to generate unit quaternion and reduce this ambiguity by choosing the quaternion with positive real part. Refer to [5] for more details on unit quaternion and 3D rotation.

2.2. Quaternion product units

In standard deep learning models, each of their neurons can be represented as a weighted summation unit, *i.e.*, $y = \sigma(\sum_{i=1}^N w_i x_i + b)$, where $\{w_i\}_{i=1}^N$ and b are learnable parameters, $\{x_i\}_{i=1}^N$ are inputs, and $\sigma(\cdot)$ is a nonlinear activation function. As aforementioned, when the input $x_i \in \mathbb{S}\mathbb{O}(3)$, such a unit cannot keep the output y on $\mathbb{S}\mathbb{O}(3)$ as well. To design a computational unit guaranteeing the closure of $\mathbb{S}\mathbb{O}(3)$, we propose an alternate of this unit based on the quaternion algebra introduced above.

Specifically, given N unit quaternions $\{q_i = [s_i, \mathbf{v}_i]\}_{i=1}^N$ that represent 3D rotations, we can define a weighted chain of Hamilton products as

$$y = \bigotimes_{i=1}^N q_i^{w_i} = q_1^{w_1} \otimes q_2^{w_2} \otimes \dots \otimes q_N^{w_N}, \quad (5)$$

where the power of a quaternion $q = [s, \mathbf{v}]$ with a scalar w is defined as [5]:

$$q^w := \begin{cases} [\cos(w \arccos(s)), \frac{\mathbf{v}}{\|\mathbf{v}\|_2} \sin(w \arccos(s))], & \mathbf{v} \neq \mathbf{0}, \\ [1, \mathbf{0}], & \text{otherwise.} \end{cases}$$

Note that the power of a quaternion only scales the rotation angle and does not change the rotation axis.

Here, we replace the weighted summation with a weighted chain of Hamilton products, which makes $\mathbb{S}\mathbb{O}(3)$ closed under the operation. Based on the operation defined in Eq. (5), we proposed our quaternion product unit:

$$\text{QPU}(\{q_i\}_{i=1}^N; \{w_i\}_{i=1}^N, b) = \bigotimes_{i=1}^N \text{qpow}(q_i; w_i, b), \quad (6)$$

where for $q_i = [s_i, \mathbf{v}_i]$

$$\begin{aligned} & \text{qpow}(q_i; w_i, b) \\ &= [\cos(w_i(\arccos(s_i) + b)), \frac{\mathbf{v}_i}{\|\mathbf{v}_i\|_2} \sin(w_i(\arccos(s_i) + b))] \end{aligned}$$

represents the weighting function on the rotation angles with a weight w_i and a bias b .

Compared with Eq. (5), we add a bias to b to shift the origin. The output of arccos contains an infinite gradient at ± 1 . We solve this problem by clamping the input scalar part s between $-1 + \epsilon$ and $1 - \epsilon$, where ϵ is a small number.

2.3. Rotation-invariance and equivariance of QPU

The following proposition demonstrates the advantage of our QPU on achieving rotation robustness.

Proposition 1. *The output of the QPU is a quaternion containing a rotation-invariant real part and a rotation-equivariant imaginary part.*

Proof. This proposition is followed directly by the property of the Hamilton product. Given two quaternions $[s_1, \mathbf{v}_1]$ and $[s_2, \mathbf{v}_2]$, we apply a Hamilton product, *i.e.*, $[s_1, \mathbf{v}_1] \otimes [s_2, \mathbf{v}_1] = [s_o, \mathbf{v}_o]$. Applying a rotation R on the vector parts, we have

$$\begin{aligned} & [s_1, R(\mathbf{v}_1)] \otimes [s_2, R(\mathbf{v}_2)] \\ &= [s_1 s_2 - \langle R(\mathbf{v}_1), R(\mathbf{v}_2) \rangle, \\ & \quad R(\mathbf{v}_1) \times R(\mathbf{v}_2) + s_1 R(\mathbf{v}_2) + s_2 R(\mathbf{v}_1)]. \end{aligned} \quad (7)$$

Because $\langle R(\mathbf{v}_1), R(\mathbf{v}_2) \rangle = \langle \mathbf{v}_1, \mathbf{v}_2 \rangle$ and $R(\mathbf{v}_1) \times R(\mathbf{v}_2) = R(\mathbf{v}_1 \times \mathbf{v}_2)$, we have

$$\begin{aligned} & [s_1, R(\mathbf{v}_1)] \otimes [s_2, R(\mathbf{v}_2)] \\ &= [s_1 s_2 - \langle \mathbf{v}_1, \mathbf{v}_2 \rangle, R(\mathbf{v}_1 \times \mathbf{v}_2 + s_1 \mathbf{v}_2 + s_2 \mathbf{v}_1)] \\ &= [s_o, R(\mathbf{v}_o)]. \end{aligned} \quad (8)$$

Thus, the Hamilton product of two quaternions gives a rotation-invariant real part and a rotation-equivariant imaginary part. The same property holds for the chain of weighted Hamilton products. \square

The proof above indicates that the intrinsic property and the group law of $\mathbb{S}\mathbb{O}(3)$ naturally provided rotation-invariance and rotation-equivariance. Without forced tricks or hand-crafted representation strategies, the principled design of QPU makes it flexible for a wide range of applications with different requirements on rotation robustness. Figure 2 further visualizes the property of our QPU. Given a QPU, we feed it with a human skeleton and its rotated version, respectively. We compare their outputs and find that the imaginary parts of their outputs inherit the rotation discrepancy between them while the real parts of their outputs are the same with each other.

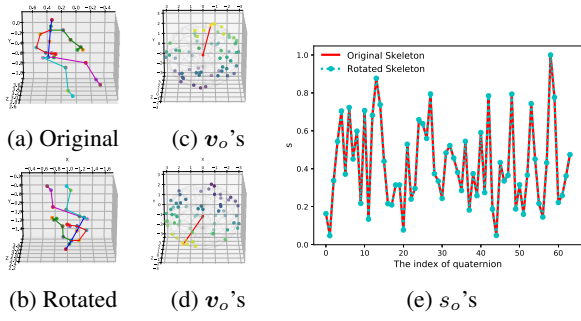


Figure 2: An illustration of the rotation-invariance and the rotation-equivariance of QPU. For each skeleton, we represent the relative rotations between its joints as unit quaternions and treat these quaternions as the input of a QPU. To visualize the rotation discrepancy, for a v_o in (c) and the corresponding v_o in (d), we connect each of them (normalized to a sphere) with the origin via a red line.

2.4. Backpropagation of QPU

The forward pass of QPU involves N individual weighting functions and a chain of Hamilton products. While it is easy to compute the gradients of the weighting functions [5], the gradient of the chain of Hamilton products is not straightforward.

Given quaternions $\{q_i\}_{i=1}^N$, we first compute the differential of the k -th quaternion q_k with respect to chain of Hamilton product $\bigotimes_{i=1}^N q_i$, denoted as $d(\bigotimes_i q_i)_{q_k}(\cdot)$, omitting N for simplicity. As Hamilton product is bilinear, we have

$$\begin{aligned} d(\bigotimes_i q_i)_{q_k}(h) &= (\bigotimes_{j<k} q_j) \otimes h \otimes (\bigotimes_{j>k} q_j) \\ &= M_L(\bigotimes_{j<k} q_j) M_R(\bigotimes_{j>k} q_j) h \quad (9) \\ &= M_L(B_k) M_R(A_k) h. \end{aligned}$$

where h represents a small variation of the k -th input quaternion, $B_k = \bigotimes_{j<k} q_j$ and $A_k = \bigotimes_{j>k} q_j$ are the chains of Hamilton product of the quaternions before and after the k -th quaternion, respectively.

We then compute the differential of loss L for the k -th quaternion. Given the gradient of the scalar loss L with respect to the output quaternion $\frac{\partial L}{\partial(\bigotimes_i q_i)}$, *e.g.*, computed with autograd [15] of Pytorch, we have

$$\begin{aligned} dL_{q_k}(h) &= \left\langle \frac{\partial L}{\partial(\bigotimes_i q_i)}, d(\bigotimes_i q_i)_{q_k}(h) \right\rangle \\ &= \left\langle \frac{\partial L}{\partial(\bigotimes_i q_i)}, M_L(B_k) M_R(A_k) h \right\rangle \quad (10) \\ &= \left\langle M_R^T(A_k) M_L^T(B_k) \frac{\partial L}{\partial(\bigotimes_i q_i)}, h \right\rangle, \end{aligned}$$

where $\langle \cdot, \cdot \rangle$ stands for the inner product between vectors and

M^T stands for the transpose of matrix M . Thus the gradient of loss L for q_k is given as

$$\frac{\partial L}{\partial q_k} = M_R^T(A_k) M_L^T(B_k) \frac{\partial L}{\partial(\bigotimes_i q_i)}. \quad (11)$$

To compute B_k and A_k , we first compute the cumulative Hamilton product $C = (c_1, c_2, \dots, c_N)$, where $c_k = \bigotimes_{j \leq k} q_j$. Then $B_k = c_k \otimes q_k^*$ and $A_k = c_k^* \otimes c_N$. We note that the gradient depends on a cumulative Hamilton product of all quaternions other than the k -th quaternion. This computation process is different from that of standard weighted summation units, in which the backpropagation through addition only involves the differential variable instead of other inputs. In contrast, it is interesting to see the gradient is a joint result of all other inputs to QPU except the items of the given differential variable.

3. Implementations and Further Analysis

3.1. QPU-based neural networks

As shown in Figure 1b, multiple QPUs receiving the same input quaternions form a fully-connected (FC) layer. The parameters of the QPU-based FC layer include the weights and the bias used in the QPUs. We initialize them using the Xavier uniform initialization [8]. Different from real-valued weighted summation, the QPU itself is nonlinear. Hence there is no need to introduce additional nonlinear activation functions. As aforementioned, the output of QPU is still a unit quaternion so that we can connect multiple QPU-based FC layers sequentially. Accordingly, the stack of multiple QPU-based FC layers establishes a quaternion multi-layer perceptron (QMLP) model for 3D rotation data. Note that the Hamilton product is not commutative, *i.e.*, $q_1 \otimes q_2 \neq q_2 \otimes q_1$ in general, the stacking of two QPU-based FC layers is not equivalent to one QPU-based FC layer.

Besides the QPU-based FC layer and the corresponding QMLP model, our QPU can also be used to implement a quaternion-based graph convolution layer. Specifically, given a graph with an adjacency matrix $A = [a_{ij}] \in \mathbb{R}^{N \times N}$, the aggregation of its node embeddings is achieved by the multiplication between its adjacency matrix and the embeddings [11]. When the node embeddings correspond to 3D rotations and are represented as unit quaternions, we can implement a new aggregation layer using the chain of Hamilton products shown in Eq. (5), *i.e.*, given input quaternions $\{q_i\}_{i=1}^N$ and the adjacency matrix A , the i -th output $q_i^{agg} = \bigotimes_{j=1}^N q_j^{a_{ij}}$. Stacking such a QPU-based aggregation layer with a QPU-based FC layer, we achieve a quaternion-based graph convolution layer accordingly.

3.2. Compatibility with real-valued models

Besides building pure quaternion-valued models (*i.e.*, the QMLP mentioned above), we can plug our QPU-based

layers into existing real-valued models easily. Suppose that we have one QPU-based layer, which takes N unit quaternions as inputs and derives M unit quaternions accordingly. When the QPU-based layer receives rotations from a real-valued model, we merely need to reformulate the inputs as unit quaternions. When a real-valued layer follows the QPU-based layer, we can treat the output of the QPU-based layer as a real-valued matrix with size $M \times 4$ and directly feed the output to the subsequent real-valued layer. Moreover, when rotation-invariant (rotation-equivariant) features are particularly required, we can feed only the real part (the imaginary part) to the real-valued layer accordingly.

Besides the straightforward strategy above, we further propose an angle-axis map to connect QPU-based layers with real-valued ones. Specifically, given a unit quaternion $q = [s, \mathbf{v}]$, the proposed mapping function is defined as:

$$\text{AngleAxisMap}(q) = [\arccos(s), \frac{\mathbf{v}}{\|\mathbf{v}\|_2}]. \quad (12)$$

We use this function to map the output of a QPU, which lies on a non-Euclidean manifold [10], to the Euclidean space. Using this function to connect QPU-based layers with real-valued ones makes the learning of the downstream real-valued models efficient, which helps us improve learning results. Our angle-axis mapping function is better than the logarithmic map used in [28, 27, 9, 24]. The logarithmic map $\log(q) = [0, \arccos(s) \frac{\mathbf{v}}{\|\mathbf{v}\|_2}]$ mixes the angular and axial information, while our map keeps this disentanglement, which is essential to preserve the rotation-invariance and rotation-equivariance of the output.

3.3. Computational Complexity and Accelerations

As shown in Section 2.4, our QPU supports gradient-based update, so both QPU-based layers and the models combining QPU-based layers with real-valued ones can be trained using backpropagation. Since each QPU handles a quaternion as one single element, QPU is efficient in terms of parameter numbers. A real-valued FC with N inputs and M outputs requires NM parameters, while a QPU with the same input and output dimensions requires only $\frac{1}{16}NM$ parameters.

Due to the usage of Hamilton products, however, QPU requires more computational time than the real-valued weighted summation unit. Take a QPU-based FC layer with N input quaternions and M output quaternions as an example. In the forward pass, the weighting step requires NM multiplications and NM sine and cosine computations. Each Hamilton product requires 16 multiplications and 12 additions. So the chain of $N - 1$ Hamilton products requires $16(N - 1)M$ multiplications and $12(N - 1)M$ additions. In total, a QPU-based FC layer requires $(17N - 16)M$ multiplications, $(13N - 12)M$ additions and NM cosine and sine computations. As a comparison, the real-valued FC

layer with the same input and output size (*i.e.*, $4N$ inputs and $4M$ outputs in terms of real numbers) requires $16NM$ multiplications, $16(N - 1)M$ additions.

In the backward pass, the gradient computation of a chain of Hamilton products requires the cumulative Hamilton product of the power weighted input quaternions. In particular, when computing the gradient of a chain of Hamilton products for an input quaternion, we need to do two Hamilton products and one matrix-matrix multiplication. As a result, the computational time is at least doubled compared with the forward pass and tripled if we recompute the cumulative Hamilton product. To deal with this challenge, we have two strategies: (a) store the result of cumulative Hamilton product in the forward pass, (b) recompute the same quantity during the backward pass. The strategy (a) saves computational time but requires to save a potentially large feature map (*i.e.*, N times the size of the feature map). The strategy (b) requires more computation time but no additional memory space. We tested both options in our experiments. When the QPU-based layer is with small feature maps (*i.e.*, at the beginning of a neural network with fewer input channels), both these two strategies work well. When we stack multiple QPU-based layers and apply the model to large feature maps, we need powerful GPUs with large memory spaces to implement the first strategy.

According to the above analysis, the computational bottleneck of our QPU is the chain of Hamilton products. Fortunately, we can accelerate this step by breaking down the computation like a tree. Specifically, as Hamilton product satisfies the combination law, *i.e.*, $(q_1 \otimes q_2) \otimes q_3 = q_1 \otimes (q_2 \otimes q_3)$, we can multiply quaternions at odd positions with quaternions at even positions in a parallel way and compute the whole chain by repeating this step recursively. For a chain of $N - 1$ Hamilton products, this parallel strategy reduces the time complexity from $\mathcal{O}(N)$ to $\mathcal{O}(\log(N))$. We implemented this strategy in the forward pass and witnessed a significant acceleration. In summary, although the backpropagation of our QPU-based model requires much more computational resources (*i.e.*, time, or memory) than that of real-valued models, its forward pass merely requires slightly more computations than that of real-valued models and can be accelerated efficiently via a simple parallel strategy.

4. Related Work

3D deep learning Deep learning has been widely used in the tasks relevant to 3D data. A typical application is skeleton-based action recognition [33, 23, 21, 9]. Most existing methods often use 3D positions of joints as their inputs. They explore the graph structure of the skeleton data and achieve significant gains in performance. For example, the AGC-LSTM in [23] introduces graph convolution to an LSTM layer and enhances the performance with an atten-

tion mechanism. The DGNN in [21] uses both node feature and edge feature to perform deep learning on a directed graph. In [12], deep learning on symmetric positive definite (SPD) matrix was used to perform skeleton-based hand gesture recognition. Recently, several works are proposed to perform action recognition based on 3D rotations between bones. Performing human action recognition by representing human skeletons in a Lie group (rotation and translation) is first proposed in [27] and further explored in [28]. The LieNet in [9] utilizes similar data representation in a deep learning framework on 3D rotation manifold, where input rotation matrices were transformed progressively by being multiplied with learnable rotation matrices. While the LieNet uses pooling to aggregate rotation features, our proposed QPU is a more general learning unit, where Hamilton products could explore the interaction between rotating features. Moreover, the weights in the LieNet are rotation matrices, whose optimization is constrained on the Lie group, while our QPU has unconstrained real-valued weights and can be optimized using standard backpropagation.

Besides the methods focusing on 3D skeletons, many neural networks are designed for 3D point clouds. The spherical CNNs [4, 6] project a 3D signal onto a sphere and perform convolutions in the frequency domain using spherical harmonics, which achieve rotation-invariant features. The ClusterNet [3] transforms point cloud into a rotation-invariant representation as well. The 3D Steerable CNNs [29] and the TensorField Networks [26] incorporate group representation theory in 3D deep learning and learn rotation-equivariant features by computing equivariant kernel basis on fields. Note that these models can only achieve either rotation-invariance or rotation-equivariance. To our knowledge, our QPU makes the first attempt to achieve these two important properties in a unified framework.

Quaternion-based Learning Quaternion is widely used in computer graphics and control theory to represent 3D rotation, which only requires four parameters to describe a rotation matrix. Recently, many efforts have been made to introduce quaternion-based models to the applications of computer vision and machine learning. Quaternion wavelet transforms (QWT) [1, 36] and quaternion sparse coding [34] have been successfully applied for image processing. For skeleton-based action recognition, the quaternion lifting schema [24] used the SLERP [22] to extract hierarchical information in rotation data to perform gait recognition. The QuaterNet [16] performs human skeleton action prediction by predicting the relative rotation of each joint in the next step with unit quaternion representation.

More recently, quaternion neural networks (QNNs) have produced significant advances [38, 14, 13, 25]. The QNNs replace multiplications in standard networks with Hamilton products, which can better preserve the interrelationship be-

tween channels and reduce parameters. The QCNN in [38] proposed to represent rotation in color space with quaternion. The QRNN in [14] extended RNN to a quaternion version and applied it to speech tasks. Different from these existing models, the proposed QPU roots its motivation in the way of interaction on the 3D rotation group and uses Hamilton product to merge power weighted quaternions.

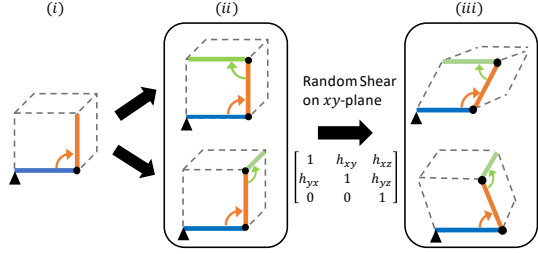
5. Experiments

To demonstrate the effectiveness of our QPU, we design the QMLP model mentioned above and test it on both synthetic and real-world datasets. In particular, we focus on the 3D skeleton classification tasks like human action recognition and hand action recognition. For our QPU-based models, we represent each skeleton as the relative rotations between its connected bones (*a.k.a.*, joint rotations). Taking these rotations as inputs, we train our models to classify the skeletons and compare them with state-of-the-art rotation representation methods. Besides testing the pure QPU-based models, we also use QPU-based layers to replace real-valued layers in existing models [23, 21] and verify the compatibility of our model accordingly. The code is available at https://github.com/IICNELAB/qpu_code.

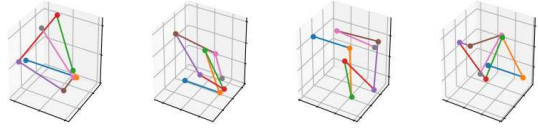
5.1. Synthetic dataset — CubeEdge

We first design a synthetic dataset called “CubeEdge” and propose a simple task to demonstrate the rotation robustness of our QPU-based model. The dataset consists of skeletons composed of consecutive edges taken from a 3D cube (a chain of edges). These skeletons are categorized into classes according to their shapes, *i.e.*, the topology of consecutive edges. For each skeleton, we use the rotations between consecutive edges as its feature.

For each skeleton in the dataset, we generate it via the following three steps, as shown in Figure 3a: *i)* We initialize the first two adjacent edges deterministically. *ii)* We ensure that the consecutive edges cannot be the same edge, so given the ending vertex of the previous edge, we generate the next edge randomly along either of the two valid directions. Accordingly, each added edge doubles the number of possible shapes. We repeat this step several times to generate a skeleton. *iii)* We apply random shear transformations to the skeletons and augment the dataset, and split the dataset for training and testing. We generate each skeleton with seven edges so that our dataset consists of 32 different shapes (classes) of skeletons. We use 2,000 samples for training and testing, respectively. For the skeletons in the testing set, we add Gaussian noise $\mathcal{N}(0, \sigma^2)$ to their vertices before shearing, where the standard deviation σ controls the level of noise. Figure 3b gives four testing samples. Additional random rotation is applied to testing set when validating the robustness to rotations.



(a) Generation process of CubeEdge. The root vertex is marked with a triangle



(b) Samples from testing set of CubeEdge with 7 edges

Figure 3: Illustrations of the CubeEdge dataset. (a) The steps to generate the skeletons. *i*) fix the two first edges; *ii*) pick the next edge from two possible candidates (doubling the number of classes); *iii*) apply random shear on *xy*-plane. (b) Samples from the testing set.

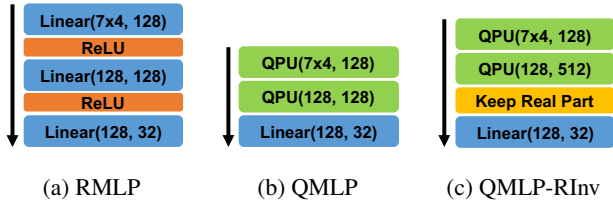
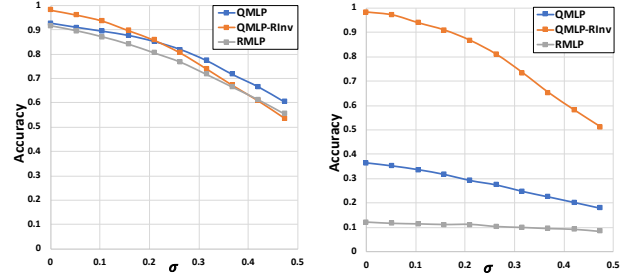


Figure 4: The architectures of the models for synthetic data.

To classify the synthetic skeletons, we apply our **QMLP** model and its variant **QMLP-RInv** and compared them with a standard real-valued MLP (**RMLP**). For fairness, all three models are implemented with three layers, and the size of output feature maps are the same. The RMLP is composed of three fully-connected layers connected by two ReLU activation functions. Our QMLP is composed of two stacked QPU layers, followed by a fully-connected layer. Our QMLP-RInv is composed of two stacked QPU layers with only the real part of the output retained, and a fully-connected layer is followed. The architectures of these three models are shown in Figure 4.

After training these three models, we test them on the following two scenarios: *i*) comparing them on the testing set directly; and *ii*) adding random rotations to the testing skeletons and then comparing the models on the randomly-rotated testing set. The first scenario aims at evaluating the feature extraction ability of different models, while the second one aims at validating their rotation-invariance. For each scenario, we compare these three models on their classification accuracy. Their results with respect to dif-



(a) Without random rotations (b) With random rotations

Figure 5: Comparisons on testing accuracy.

ferent levels of noise (*i.e.*, σ) are shown in Figure 5. In both scenarios, our QMLP and QMLP-RInv outperform the RMLP consistently, which demonstrates the superiority of our QPU model. Furthermore, we find that the QMLP-RInv is robust to the random rotations imposed on the testing data — it achieves nearly the same accuracy in both scenarios and works better than QMLP in the scenario with random rotations. This phenomenon verifies our claims: *i*) the real part of our QPU’s output is rotation-invariant indeed; *ii*) the disentangled representation of rotation-invariant feature and rotation-equivariant feature could inherently make the network robust to rotations.

5.2. Real-world data

Besides testing on synthetic data, we also consider two real-world skeleton datasets: the **NTU** dataset [20] for human action recognition and the **FPHA** dataset [7] for hand action recognition. The NTU dataset provides 56,578 human skeleton sequences performed by 40 actors belonging to 60 classes. The cross-view protocol is used in our experiments. We use 37,646 sequences for training and the remaining 18,932 sequences for testing. FPHA dataset provides hand skeletons recorded with magnetic sensors and computed by inverse kinematics, which consists of 1,175 action videos belonging to 45 categories and performed by 6 actors. Following the setting in [7], we use 600 sequences for training and 575 sequences for testing.

For these two datasets, we implement three real-valued models as our baselines.

RMLP-LSTM is composed of a two-layer MLP and a one-layer LSTM. The MLP merges each frame into a feature vector and is shared between frames. Features of each frame are then fed into the LSTM layer. We average the outputs of the LSTM at different steps and feed it into a classifier.

AGC-LSTM [23] first uses a FC layer on each joint to increase the feature dimension. The features are fed into a 3-layer graph convolutional network.

DGNN [21] is composed of 10 directed graph network

(DGN) blocks that merge both node features and edge features. Temporal information is extracted using temporal convolution between the DGN blocks.

To demonstrate the usefulness of our QPU, we substitute some of their layers with our QPU-based FC layers, and propose the following three QPU-based models:

QMLP-LSTM: The two FC layers of the RMLP-LSTM are replaced with two QPU-based FC layers.

QAGC-LSTM: The input FC layer of the AGC-LSTM replaced with a QPU-based FC layer. In the original AGC-LSTM, the first FC layer only receives a feature from one joint. As a result, the input would be a single quaternion, and the QPU layer would be unable to capture interactions between inputs. We thus augment the input by stacking the feature of the parent joint so that the input channel to QPU is two quaternions.

QDGNN: The FC layers in the first DGN block of QDGNN are replaced with QPU-based FC layers. The original DGNN uses 3D coordinates of joints as node features and vectors from parent joint to child joint as edge features. In our QDGNN, the node features are the unit quaternions corresponding to the joint rotations and the edge features are the differences between joint rotations, *i.e.*, $q_1 \otimes q_2^*$ computes the edge features for joint rotations q_1 and q_2 . Moreover, the original DGNN aggregates node features and edge features through the normalized adjacent matrix. In our QDGNN, we implement this aggregation as the QPU-based aggregation layer proposed in Section 3.1.

Similar to the experiments on synthetic data, we consider for the three QPU-based models their variants that only keep the real part of each QPU’s output, denoted as **QMLP-LSTM-RInv**, **QAGC-LSTM-RInv**, and **QDGNN-RInv**, respectively. For these variants, we multiply the output channels of the last QPU-based FC layer by 4 and keep only the real parts of its outputs. In all real-world experiments, we use the angle-axis map (Eq. (12)) to connect the QPU-based layers and real-valued neural networks. The number of parameters of the QPU-based models is almost the same with that of the corresponding real-valued models.

We compare our QPU-based models with corresponding baselines under two configurations: *i*) both training data and test data have no additional rotations (NR); *ii*) the training data are without additional rotations while the testing ones are with arbitrary rotations (AR). The comparison results on NTU and FPHA are shown in Table 1, respectively. We can find that in the first configuration, the performance of our QPU-based models is at least comparable to the baselines. In the second and more challenging setting, which requires rotation-invariance, our QPU-based models that use only outputs’ real parts as features retain high accuracy consistently in most situations, while the performance of the baselines degrades a lot. Additionally, we quantitatively analyze the impact of our angle-axis map on testing

Table 1: Comparisons on testing accuracy (%).

Model	#Param. (Million)	NTU		FPHA	
		NR	AR	NR	AR
RMLP-LSTM	0.691	72.92	24.67	67.13	17.39
QMLP-LSTM	0.609	69.72	26.60	76.17	24.00
QMLP-LSTM-RInv	0.621	75.84	75.84	68.00	67.83
AGC-LSTM	23.371	90.50	55.17	77.22	24.70
QAGC-LSTM	23.368	87.18	39.43	79.13	24.70
QAGC-LSTM-RInv	23.369	89.68	89.92	72.35	71.30
DGNN	4.078	87.45	23.30	80.35	24.35
QDGNN	4.075	86.55	41.06	82.26	27.13
QDGNN-RInv	4.076	83.88	83.88	76.35	76.35

Table 2: The impact of Angle-Axis map on accuracy (%).

Model	With Eq. (12)	Without Eq. (12)
QMLP-LSTM	76.17	73.04
QAGC-LSTM	79.13	75.65

accuracy. In Table 2, the results on the FPHA dataset verify our claim in Section 3.2: for our QPU-based models, their performance boosts a lot when we connect their QPU-based layers with the following real-valued layers through the angle-axis map.

6. Conclusion

In this work, we proposed a novel quaternion product unit for deep learning on 3D rotation groups. This model can be used as a new module to construct quaternion-based neural networks, which presents encouraging generalization ability and flexible rotation robustness. Moreover, our implementation makes this model compatible with existing real-valued models, achieving end-to-end training through backpropagation. Besides skeleton classification, we plan to extend our QPU to more applications and deep learning. In particular, we have done a preliminary experiment on applying our QPU to point cloud classification task. We designed a quaternion representation for the neighbor point set of a centroid which first converts 3D coordinates of neighbor points into quaternion-based 3D rotations then cyclically sort them according to their rotation order around the vector from the origin to the centroid. We designed our models based on Pointnet++ [19, 31] by replacing the first Set Abstraction layer with a QMLP module. We tested our model on ModelNet40 [32] and our rotation-invariant model achieved 80.1% test accuracy. Please refer to the supplementary file for more details of our quaternion-based point cloud representation, network architectures and experimental setups.

Acknowledgement The authors would like to thank David Filliat for constructive discussions.

References

- [1] Eduardo Bayro-Corrochano. The theory and use of the quaternion wavelet transform. *Journal of Mathematical Imaging and Vision*, 24(1):19–35, 2006. 6
- [2] Victoria Bloom, Dimitrios Makris, and Vasileios Argyriou. G3d: A gaming action dataset and real time action recognition evaluation framework. In *CVPRW*, 2012. 1
- [3] Chao Chen, Guanbin Li, Ruijia Xu, Tianshui Chen, Meng Wang, and Liang Lin. Clusternet: Deep hierarchical cluster network with rigorously rotation-invariant representation for point cloud analysis. In *CVPR*, 2019. 1, 6
- [4] Taco S. Cohen, Geiger Mario, Köhlerand Jonas, and Max Welling. Spherical cnns. In *ICLR*, 2016. 1, 6
- [5] Erik B Dam, Martin Koch, and Martin Lillholm. Quaternions, interpolation and animation. In *Technical Report DIKU-TR-98/5, Department of Computer Science, University of Copenhagen*, 1998. 3, 4
- [6] Carlos Esteves, Christine Allen-Blanchette, Ameesh Makhia, and Kostas Daniilidis. Learning so(3) equivariant representations with spherical cnns. In *ECCV*, 2018. 1, 6
- [7] Guillermo Garcia-Hernando, Shanxin Yuan, Seungryul Baek, and Tae-Kyun Kim. First-person hand action benchmark with rgb-d videos and 3d hand pose annotations. In *CVPR*, 2018. 7
- [8] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *AISTATS*, 2010. 4
- [9] Zhiwu Huang, Chengde Wan, Thomas Probst, and Luc Van Gool. Deep learning on lie groups for skeleton-based action recognition. In *CVPR*, 2017. 5, 6
- [10] Du Q Huynh. Metrics for 3d rotations: Comparison and analysis. *Journal of Mathematical Imaging and Vision*, 35(2):155–164, 2009. 5
- [11] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017. 4
- [12] Xuan Son Nguyen, Luc Brun, Olivier Lézoray, and Sébastien Bogleux. A neural network based on spd manifold learning for skeleton-based hand gesture recognition. In *CVPR*, 2019. 6
- [13] Titouan Parcollet, Mohamed Morchid, and Georges Linarès. Quaternion convolutional neural networks for heterogeneous image processing. In *ICASSP*, 2019. 6
- [14] Titouan Parcollet, Mirco Ravanelli, Mohamed Morchid, Georges Linarès, Chiheb Trabelsi, Renato De Mori, and Yoshua Bengio. Quaternion recurrent neural networks. In *ICLR*, 2019. 6
- [15] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *NeurIPS*, 2017. 4
- [16] Dario Pavllo, David Grangier, and Michael Auli. Quaternet: A quaternion-based recurrent model for human motion. In *BMVC*, 2018. 1, 6
- [17] Charles R Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J Guibas. Frustum pointnets for 3d object detection from rgb-d data. In *CVPR*, 2018. 1
- [18] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017. 1
- [19] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *NeurIPS*, 2017. 8
- [20] Amir Shihroudy, Jun Liu, Tian-Tsong Ng, and Gang Wang. Ntu rgb+ d: A large scale dataset for 3d human activity analysis. In *CVPR*, 2016. 7
- [21] Lei Shi, Yifan Zhang, Jian Cheng, and Hanqing Lu. Skeleton-based action recognition with directed graph neural networks. In *CVPR*, 2019. 5, 6, 7
- [22] Ken Shoemake. Animating rotation with quaternion curves. In *SIGGRAPH*, 1985. 6
- [23] Chenyang Si, Wentao Chen, Wei Wang, Liang Wang, and Tieniu Tan. An attention enhanced graph convolutional lstm network for skeleton-based action recognition. In *CVPR*, 2019. 5, 6, 7
- [24] Agnieszka Szczesna, Adam Świtoński, Janusz Słupik, Hafed Zghidi, Henryk Josiński, and Konrad Wojciechowski. Quaternion lifting scheme applied to the classification of motion data. *Information Sciences*, 2018. 5, 6
- [25] Yi Tay, Aston Zhang, Luu Anh Tuan, Jinfeng Rao, Shuai Zhang, Shuohang Wang, Jie Fu, and Siu Cheung Hui. Lightweight and efficient neural natural language processing with quaternion networks. In *ACL*, 2019. 6
- [26] Nathaniel Thomas, Tess Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley. Tensor field networks: Rotation-and translation-equivariant neural networks for 3d point clouds. *arXiv preprint arXiv:1802.08219*, 2018. 1, 6
- [27] Raviteja Vemulapalli, Felipe Arrate, and Rama Chellappa. Human action recognition by representing 3d skeletons as points in a lie group. In *CVPR*, 2014. 1, 5, 6
- [28] Raviteja Vemulapalli and Rama Chellappa. Rolling rotations for recognizing human actions from 3d skeletal data. In *CVPR*, 2016. 5, 6
- [29] Maurice Weiler, Mario Geiger, Max Welling, Wouter Boomsma, and Taco S. Cohen. 3d steerable cnns: Learning rotationally equivariant features in volumetric data. In *NeurIPS*, 2018. 1, 6
- [30] Thomas Whelan, Stefan Leutenegger, R Salas-Moreno, Ben Glocker, and Andrew Davison. Elasticfusion: Dense slam without a pose graph. *Robotics: Science and Systems*, 2015. 1
- [31] Erik Wijmans. Pointnet++ pytorch. https://github.com/erikwijmans/Pointnet2_PyTorch, 2018. 8
- [32] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *CVPR*, 2015. 8
- [33] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *AAAI*, 2018. 5
- [34] Licheng Yu, Yi Xu, Hongteng Xu, and Hao Zhang. Quaternion-based sparse representation of color image. In *ICME*, 2013. 6

- [35] Pengfei Zhang, Cuiling Lan, Junliang Xing, Wenjun Zeng, Jianru Xue, and Nanning Zheng. View adaptive neural networks for high performance skeleton-based human action recognition. In *ICCV*, 2017. 1
- [36] Jun Zhou, Yi Xu, and Xiaokang Yang. Quaternion wavelet phase based stereo matching for uncalibrated images. *Pattern Recognition Letters*, 28(12):1509–1522, 2007. 6
- [37] Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *CVPR*, 2018. 1
- [38] Xuanyu Zhu, Yi Xu, Hongteng Xu, and Changjian Chen. Quaternion convolutional neural networks. In *ECCV*, 2018. 6