

# Path Signatures on Lie Groups

**Darrick Lee**

*Department of Mathematics  
University of Pennsylvania*

LDARRICK@SAS.UPENN.EDU

**Robert Ghrist**

*Departments of Mathematics and Electrical & Systems Engineering  
University of Pennsylvania*

GHRIST@SEAS.UPENN.EDU

**Editor:**

## Abstract

Path signatures are powerful nonparametric tools for time series analysis, shown to form a universal and characteristic feature map for Euclidean valued time series data. We lift the theory of path signatures to the setting of Lie group valued time series, adapting these tools for time series with underlying geometric constraints. We prove that this generalized path signature is universal and characteristic. To demonstrate universality, we analyze the human action recognition problem in computer vision, using  $SO(3)$  representations for the time series, providing comparable performance to other shallow learning approaches, while offering an easily interpretable feature set. We also provide a two-sample hypothesis test for Lie group-valued random walks to illustrate its characteristic property. Finally we provide algorithms and a Julia implementation of these methods.

## 1. Introduction

Time series data is ubiquitous in modern data science, and may take values in a variety of forms. Perhaps the most common is a collection of simultaneous multivariate real-valued time series  $\{\gamma^i\}_{i=1}^N$ , where  $\gamma^i : [0, 1] \rightarrow \mathbb{R}$ . In this case, we may consider the entire collection  $\gamma = (\gamma^1, \dots, \gamma^N)$  as a path through Euclidean space,  $\gamma : [0, 1] \rightarrow \mathbb{R}^N$ . The *path signature* is a feature set that completely characterizes such paths, and has recently been applied to several tasks in machine learning (Chevyrev and Kormilitzin, 2016; Lyons, 2014). Recent work has provided the path signature with strong theoretical properties; namely that it is a universal and characteristic kernel for time series in Euclidean space  $\mathbb{R}^N$  (Chevyrev and Oberhauser, 2018).

However, in many scenarios, the data may have some geometric constraints, and may be better represented by elements of a (non-Euclidean) manifold. In this case, the time-varying data can be modelled as path on such a manifold, rather than on Euclidean space. *Lie groups* are smooth manifolds equipped with a compatible group structure. Paths (or time series) valued in Lie groups model a number of natural phenomena, including the following.

- The *special Euclidean group*  $SE(n)$  is the Lie group of all rigid body motions in  $\mathbb{R}^n$ . The group  $SE(3)$  is often used to model the position and pose of a rigid body, such as

a component of a robotic arm or an element of a drone swarm, with  $k$  such components or elements collectively giving rise to a path  $SE(3)^k$  (Selig, 2004).

- The *special orthogonal group*  $SO(n)$  is the Lie group of all rotations in  $\mathbb{R}^n$ ; this is a Lie subgroup of  $SE(n)$ . The Lie group  $SO(3)^k$  has recently been used to represent the pose of a human by recording the relative rotations of  $k$  pairs of body parts (Vemulapalli and Chellappa, 2016). Thus, human movement can be represented as a path in  $SO(3)^k$ . This representation has been used in the computer vision problem of human action recognition, and Lie group methods have achieved state-of-the-art results in this domain (Huang et al., 2017).
- The state of an oscillator may be described as an element of the *circle*  $S^1$ , and collective behavior of a network of oscillators can be describe by an element of the *n-torus*,  $T^n = (S^1)^n$ . The time evolution of oscillator networks can therefore be modelled as a path on  $T^n$  (Strogatz, 2000).
- The Euclidean space  $\mathbb{R}^N$  is the simplest example of a Lie group, where the group operation is addition. The classical path signature for Euclidean space can be viewed as a special case of path signatures on Lie groups.

In this paper, we extend path signatures to time series valued in Lie groups, and show that this extension is also a universal and characteristic kernel.

## 1.1 Contributions

We lift the theory of path signatures for time series valued in Euclidean space to the setting of time series valued in Lie groups, restricting ourselves to the class of piecewise regular paths on Lie groups.

**Definition 1** *Let  $G$  be a Lie group. A path  $\gamma : [a, b] \rightarrow G$  is regular if  $\gamma'_t$  is continuous and nonvanishing on the entire interval  $[a, b]$ . Such a path is piecewise regular if there exists a partition  $a = t_0 < t_1 < \dots < t_n = b$  such that  $\gamma$  is regular on each open subinterval  $(t_i, t_{i+1})$  for all  $i$ . The pathspace – the space of all piecewise regular paths on the unit interval,  $\gamma : [0, 1] \rightarrow G$  – will be denoted  $PG$ .*

Let  $G$  be a Lie group of dimension  $N$ , and let  $\mathfrak{g}$  be its Lie algebra (the tangent space at the identity). We denote the underlying vector space of  $\mathfrak{g}$  by  $\bar{\mathfrak{g}} \cong \mathbb{R}^N$ . The *path signature* is a function on paths,

$$S : PG \rightarrow T((\bar{\mathfrak{g}})),$$

valued in a formal power series of tensors,  $T((\bar{\mathfrak{g}}))$ , where we may view the coefficients as descriptors (or features) of the underlying path (or time series). Path signatures for general manifolds were originally defined by Chen (1958), but not in a manner conducive to data analysis. This paper gives a computationally clean derivation for path signatures on Lie groups tuned for use in data analysis.

Our generalization is designed to be analogous to the Euclidean case as much as possible, for ease of applicability. For example, the definition of the path signature for  $\gamma : [0, 1] \rightarrow G$

depends only on the derivative  $\gamma' : [0, 1] \rightarrow \mathfrak{g}$ . We exploit one of the key properties of Lie groups — that tangent vectors at a point correspond to elements of its Lie algebra  $\mathfrak{g}$ , a vector space. This will permit a signature construction making use of iterated integrals as per the Euclidean case.

In the Euclidean case  $G = \mathbb{R}^N$ , the Lie group is often conflated with its Lie algebra  $\mathfrak{t} = \mathbb{R}^N$ , and the fact that the integration is performed in the Lie algebra is often not made. By clarifying and emphasizing this point, the generalization to Lie groups illuminates understanding of the classical Euclidean case.

From a machine learning perspective, the basic properties of the path signature as a feature map provide several benefits.

- The signature is a feature set for a path as a whole, and can be used to compare time series with varying numbers of time points.
- Defined as iterated line integrals, the path signature is invariant under reparametrization, and thus only depends on the *order* in which events occur.
- The signature is left translation invariant, meaning the signatures of paths that differ by a constant element  $g \in G$  will be the same. This implies that the signature only depends on the dynamics of the time series and is unconcerned with the initial point.
- The antisymmetrization of the second degree signature tensor can be viewed as an indicator of lead-lag behavior in the time series. In the case of Lie groups, the interpretation will be considered in terms of left-invariant vector fields.

However, the most crucial property is that the path signature fully characterizes paths up to tree-like equivalence; that is, the map  $S$  is injective, up to quotienting  $PG$  out by an equivalence relation. This fact is originally due to Chen (1958) for the case of piecewise regular paths on Lie groups, and later generalized by Hambly and Lyons (2010) to the case of bounded variation paths in  $\mathbb{R}^n$ .

Our main contribution is to apply this injectivity result to prove that a normalized variant of the signature,  $S : PG \rightarrow T(\overline{\mathfrak{g}})$ , is a *universal* and *characteristic* feature map for time series in  $G$ , when we equip  $T(\overline{\mathfrak{g}})$  with the structure of a Hilbert space. This is proved in Section 4.2. This was originally shown for the Euclidean case by Chevyrev and Oberhauser (2018). Such feature maps can be used to two large classes of machine learning problems, in the context of kernel methods.

1. **(Studying functions on  $PG$ )** Solving a classification problem on  $PG$  can be reduced to finding a function  $f : PG \rightarrow \mathbb{R}$  such that the level set  $f = 0$  provides the decision boundary. The *universality* of the normalized signature map states that any continuous bounded function  $f : PG \rightarrow \mathbb{R}$  can be approximated using a linear functional  $f(\cdot) \approx \langle \ell, \tilde{S}(\cdot) \rangle$ . This allows us to reduce a nonlinear optimization problem into a linear one, greatly reducing the complexity.
2. **(Studying measures on  $PG$ )** Two-sample hypothesis testing on  $PG$  requires the computation of a set of statistics that is rich enough to distinguish any two probability

measures on  $PG$ . The *characteristicness* of the normalized signature map states that the kernel mean embedding (KME) is injective with respect to the normalized signature

$$\bar{\Phi} : \mathcal{M}(PG) \rightarrow T((\bar{\mathfrak{g}})), \quad \bar{\Phi}(\mu) = \mathbb{E}_\mu[S],$$

where  $\mathcal{M}(PG)$  denotes all finite regular Borel measures on  $PG$ , and  $S$  is appropriately normalized. This allows us to consider probability measures as elements of a linear space; furthermore, the norm induced by the Hilbert space structure coincides with the maximum mean discrepancy (MMD) between measures.

We perform two experiments that demonstrate the efficacy of the path signature for these two classes of problems. First, we consider the computer vision problem of human action recognition in Section 5.1. We show that the path signature method is much easier to use than shallow learning methods previously applied to this problem (Vemulapalli et al., 2014; Vemulapalli and Chellappa, 2016) while providing comparable results. Second, in Section 5.2, we consider a hypothesis testing problem for simulated random walks on the Lie group  $SO(3)$ . Here, we show that the Lie group valued path signature vastly outperforms the Euclidean path signature.

Along the way, we will establish extensions of other properties of the path signatures to Lie groups and discuss several concepts related to path signatures and data analysis on Lie groups more broadly. A summary of these contributions is given below.

1. We provide a detailed exposition of Lie group valued time series, and discuss a notion of scaling for such time series in Section 2.2. Scaling of data is sometimes required when the data needs to be normalized, and we discuss how scaling affects the path signature in Section 3.1. We also discuss the continuous interpretation of discrete time series on Lie groups in Section 2.3.
2. For  $G$  an  $N$ -dimensional Lie group, we give a signature-preserving bijection between  $PG$  and  $P\mathbb{R}^N$  in Section 3.2, which provides a Euclidean representation of Lie group valued time series. This bijection allows the exportation of Euclidean data analysis tools to Lie group valued data. With the metric introduced in Section 3.3, this bijection is an isometry.
3. It is well known that the Euclidean path signature is equivariant with respect to linear transformations (Friz and Victoir, 2010). We show that path signatures are equivariant under Lie group homomorphisms in general. Namely, given a homomorphism of Lie groups  $F : G_1 \rightarrow G_2$ , where  $\mathfrak{g}_1$  and  $\mathfrak{g}_2$  are the respective Lie algebras, we define the action of this homomorphism on the tensor algebra  $F_* : T((\bar{\mathfrak{g}}_1)) \rightarrow T((\bar{\mathfrak{g}}_2))$ , and show in Section 3.4 that

$$S(F\gamma) = F_*S(\gamma)$$

for all  $\gamma \in PG$ .

4. An important feature of the path signature is the interpretability of lower level signature terms. We discuss the extension of the lead-lag interpretation of second level signature terms for Euclidean paths, as well as a topological interpretation of the first level signature terms for abelian Lie groups in Section 3.6.
5. Path transformations, such as appending the time parameter or using a sliding window, are often used as a preprocessing step for Euclidean path signatures (Chevyrev and Kormilitzin, 2016). We discuss these transformations in the context of breaking reparametrization or left-translation invariance in Section 3.8. Empirical studies (Fermanian, 2019) have shown that the sliding window transformation (also called the lead-lag transformation) provides good classification results, despite the lack of a theoretical explanation. We propose one explanation, which is that the sliding window transformation breaks left-translation invariance, and we provide empirical evidence in the experiments in Section 5.1.
6. We provide both algorithmic details and a Julia package for the computation of path signatures valued in Lie groups, which can be found at <https://github.com/ldarrick/PathSignatures>. For details, see Appendix A.

## 1.2 Previous and related work

The concept of path signatures is relatively new in data science and machine learning (Lyons, 2014; Chevyrev and Kormilitzin, 2016; Giusti and Lee, 2020), but has deep roots in topology and geometry. Chen originally defined the path signature for piecewise regular paths on manifolds and proved several basic properties in a sequence of papers (Chen, 1954, 1957, 1958). He later studied the geometry and topology of path spaces and loop spaces by constructing a rational cochain model of these spaces, in which path signatures constitute 0-cochains (Chen, 1977).

Lyons (1998) developed the concept of the path signature in a different direction, using the path signature as a construction to lift bounded variation paths on  $\mathbb{R}^N$  to paths of power series of tensors  $T((\mathbb{R}^N))$ . This initiated the study of *rough paths*, which can be thought of as a generalization of the path signature to highly irregular paths. This theory was then used to study stochastic processes and stochastic differential equations (Lyons and Qian, 2007; Lyons et al., 2007; Friz and Victoir, 2010).

Within machine learning, path signatures have been used to study real-valued time series data in a variety of settings. Examples can be found in the study of financial time series (Gyurkó et al., 2013; Lyons et al., 2014), handwritten character recognition (Yang et al., 2016), human action recognition using position data (Yang et al., 2019), identifying psychological or neurological disorders (Moore et al., 2019; Zimmerman et al., 2018; Arribas et al., 2017) and featurizing the output of persistent homology in topological data analysis (Chevyrev et al., 2020).

The theoretical aspects of path signatures in the context of kernel methods were developed in Kiraly and Oberhauser (2019) and Chevyrev and Oberhauser (2018). The present paper is largely inspired by these two papers. The concept of using the path signature as a kernel for time series was first proposed in Kiraly and Oberhauser (2019), and efficient algorithms for computing the kernel were developed. The path signature for Euclidean space

was shown to be a universal and characteristic feature map in Chevyrev and Oberhauser (2018). This exploits the recently formalized duality between universal and characteristic kernels in Simon-Gabriel and Schölkopf (2018).

It is well known that Euclidean path signatures are translation invariant, and we will show that Lie group path signatures are left translation invariant. Diehl and Reizenstein (2019) has considered the related problem of determining the Euclidean path signature terms which are invariant under some matrix Lie group action.

We begin in Section 2 by reviewing basic facts on Lie groups and Lie algebras, and provide an exposition on continuous and discrete time series on Lie groups. We then define the path signature for Lie groups in Section 3, and discuss the bijection between  $PG$  and  $P\mathbb{R}^N$ , the equivariance of the path signature, detecting lead-lag behavior in time series, and path transformations. In Section 4, we provide a brief overview of kernel methods and prove our main result, which shows that the path signature kernel is universal and characteristic. Finally, in Section 5, we apply the path signature on Lie groups to a human action classification problem and a hypothesis testing problem involving random walks on  $SO(3)$ .

### 1.3 Notation

Throughout this paper, we will denote the time parameter for a path  $\gamma : [0, 1] \rightarrow G$  using a subscript  $t$ , meaning  $\gamma_t := \gamma(t)$ . Derivatives are shown using the prime notation, as in  $\gamma'_t := \frac{d\gamma}{dt}(t)$ . If we have a path in Euclidean space  $\alpha : [0, 1] \rightarrow \mathbb{R}^N$ , we will use superscripts to represent the components, such as  $\alpha = (\alpha^1, \alpha^2, \dots, \alpha^N)$ . If  $G$  is a Lie group, we will use  $\mathfrak{g}$  to denote its Lie algebra and use  $\bar{\mathfrak{g}}$  to be the underlying vector space of  $\mathfrak{g}$  (forgetting the Lie bracket structure).

Continuous paths will often be denoted using the lowercase Greek symbols  $\alpha, \beta, \gamma$ , and the space of all piecewise regular paths in  $G$  is denoted  $PG$ . For  $T \in \mathbb{N}$ , we let  $[T] = \{1, \dots, T\}$  denote the finite set of integers up to  $T$ . Discrete time series will be distinguished using the hat notation  $\hat{\gamma} : [T] \rightarrow G$ , and the space of all discrete time series in  $G$  will be denoted  $\hat{P}G$ .

There are also several parameters that will be used consistently throughout the paper. Unless otherwise specified, we reserve the following symbols for the given meaning.

- $N$  is the dimension of the Lie group  $G$  that paths take values in;
- $T + 1$  is the length of a discrete time series (so that the discrete derivative will be of length  $T$ );
- $M$  is the level of the truncated signature.

## 2. Lie groups, paths, and time series

We begin this section by recalling several basic facts about Lie groups Alexandrino and Bettiol (2015), followed by paths on Lie groups and the interpretation of sampled time series on Lie groups, stressing the differences from sampled time series on Euclidean space.

## 2.1 A review of Lie groups

Recall that a *Lie group*  $G$  is a smooth manifold with a group structure such that the multiplication and inversion maps are both smooth. Let  $g_0 \in G$ . The *left translation map* by  $g_0$ , written as  $L_{g_0} : G \rightarrow G$ , is defined to be  $L_{g_0}(g) = g_0g$ . The *right translation map*  $R_{g_0} : G \rightarrow G$  is defined analogously. This induces a mapping on tangent spaces  $L_{g_0*} : T_gG \rightarrow T_{g_0g}G$ . A vector field  $X$  on  $G$  is called *left-invariant* if

$$L_{g_0*}X(g) = X(L_{g_0}g) = X(g_0g)$$

for all  $g_0, g \in G$ . This implies that all left-invariant vector fields  $X$  are defined by their value at the identity  $e \in G$ ,

$$X(g) = L_{g*}X(e),$$

and thus, we obtain a one-to-one correspondence between left-invariant vector fields and the tangent space at the identity, which we denote by  $\mathfrak{g} := T_eG$ . Vector fields act on smooth functions  $f : G \rightarrow \mathbb{R}$ , and we define an operation of left-invariant vector fields  $X$  and  $Y$  by

$$[X, Y](f) := X(Y(f)) - Y(X(f)),$$

where  $[X, Y]$  is also left-invariant. This provides  $\mathfrak{g}$  with the structure of a Lie algebra, where the Lie bracket  $[\cdot, \cdot] : \mathfrak{g} \times \mathfrak{g} \rightarrow \mathfrak{g}$  is a bilinear mapping such that for all  $X, Y, Z \in \mathfrak{g}$

$$\begin{aligned} [X, Y] &= -[Y, X] \\ [X, [Y, Z]] + [Z, [X, Y]] + [Y, [Z, X]] &= 0. \end{aligned}$$

Similarly, left translation induces a map  $L_{g_0}^* : T_{g_0g}^*G \rightarrow T_g^*G$  on cotangent spaces. A 1-form  $\omega \in T^*G$  is called *left-invariant* if

$$L_{g_0}^*\omega(g_0g) = L_{g_0}^*\omega(L_{g_0}g) = \omega(g)$$

for all  $g, g_0 \in G$ . Again, we obtain a correspondence between left-invariant 1-forms and the cotangent space at the identity via the property

$$\omega(g) = L_{g^{-1}}^*\omega(e).$$

Thus, we may identify the left-invariant 1-forms by the dual of the Lie algebra,  $\mathfrak{g}^*$ .

**Remark 2** *We can think of tangent vectors (derivatives) of a path  $\gamma : [0, T] \rightarrow G$  as elements of the Lie algebra  $\mathfrak{g}$  in two ways. First, for a tangent vector  $v \in T_gG$ , we can compute the pushforward the tangent vector along the left multiplication map  $L_{g^{-1}}v \in T_eG = \mathfrak{g}$ . Second, a basis of the Lie algebra provides a global frame for  $G$ , meaning, it provides a basis for  $T_gG$  for all  $g$ . By considering  $v \in T_gG$  in terms of this basis, we may also think of  $v$  as an element of  $\mathfrak{g}$ .*

*In summary, the structure of the Lie group allows us to consider tangent vectors at any point on  $G$  using a single vector space: a fact repeatedly used throughout this paper.*

Given a left-invariant vector field  $X \in \mathfrak{g}$ , there exists a unique 1-parameter subgroup  $\rho_X : \mathbb{R} \rightarrow G$  such that  $\rho_X(0) = e$  and  $\rho_X'(0) = X(e)$ . This is defined by the integral curve of  $X$  which passes through the identity at  $t = 0$ .

**Definition 3** *The Lie exponential map of  $G$  is defined as*

$$\exp : \mathfrak{g} \rightarrow G, \quad \exp(X) := \rho_X(1),$$

where  $\rho_X$  is the 1-parameter subgroup defined above.

This exponential map provides a way to move between a Lie group and its Lie algebra.

**Proposition 4** *The exponential map  $\exp : \mathfrak{g} \rightarrow G$  is smooth and  $d(\exp)_0 = \text{id}$ . Thus,  $\exp$  is a diffeomorphism between an open neighborhood of the origin  $0 \in \mathfrak{g}$  and an open neighborhood of the identity  $e \in G$ .*

Thus, if elements are near the origin, we can define an inverse map.

**Definition 5** *Suppose  $U \subset \mathfrak{g}$  is a neighborhood of the origin such that the exponential map is a diffeomorphism. Let  $V = \exp(U)$ . The logarithm map on  $V$  is defined to be*

$$\log : V \rightarrow \mathfrak{g}, \quad \log(g) := \exp^{-1}(g).$$

A homomorphism of Lie groups  $F : G \rightarrow H$  is a smooth map which is also a group homomorphism, and a homomorphism of Lie algebras  $\phi : \mathfrak{g} \rightarrow \mathfrak{h}$  is a linear map that preserves the Lie bracket  $F([X, Y]) = [F(X), F(Y)]$  for all  $X, Y \in \mathfrak{g}$ . A Lie group homomorphism  $F : G \rightarrow H$  induces a Lie algebra homomorphism  $F_* : \mathfrak{g} \rightarrow \mathfrak{h}$  between the respective Lie algebras by the induced map between the tangent spaces at the identity  $F_* : T_e G \rightarrow T_e H$ .

**Example 1** *The special orthogonal group  $SO(3)$  — orientation-preserving rotations of  $\mathbb{R}^3$  — will be the running example used throughout this paper. This is a matrix Lie group and can be explicitly described as the space of all  $3 \times 3$  orthogonal matrices ( $AA^\top = A^\top A = I$ ) with determinant  $+1$ . The Lie algebra of  $SO(3)$  is  $\mathfrak{so}(3)$ , which consists of all  $3 \times 3$  skew-symmetric matrices ( $B = -B^\top$ ). An explicit basis for  $\mathfrak{so}(3)$  is*

$$e_1 = \begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad e_2 = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{pmatrix}, \quad e_3 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{pmatrix}.$$

We will denote the duals of these basis vectors to be  $\omega_i = e_i \in \mathfrak{g}^*$ . For all matrix Lie groups, the Lie exponential and logarithm are simply the matrix exponential and logarithm. Suppose  $\theta \in \mathbb{R}$ . The exponential map in these three basis directions gives us

$$\begin{aligned} \exp(\theta e_1) &= \begin{pmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{pmatrix}, \\ \exp(\theta e_2) &= \begin{pmatrix} \cos \theta & 0 & \sin \theta \\ 0 & 1 & 0 \\ -\sin \theta & 0 & \cos \theta \end{pmatrix}, \\ \exp(\theta e_3) &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \theta & -\sin \theta \\ 0 & \sin \theta & \cos \theta \end{pmatrix}. \end{aligned}$$



These are exactly the rotation matrices about the  $z$ ,  $y$ , and  $x$  axes respectively. Therefore, we may think of the basis vectors  $e_i$  of the Lie algebra as infinitesimal rotations in the respective directions. In particular, given a path  $\gamma \in P(SO(3))$ , the value  $\omega_i(\gamma'_t)$  corresponds to the infinitesimal rotation of  $\gamma$  at time  $t$  in the direction of  $e_i$ . If we integrate this over the domain of the path,

$$\int_0^1 \omega_i(\gamma'_t) dt,$$

we obtain the cumulative rotation of  $\gamma$  in the direction of  $e_i$  over the unit interval. This interpretation will be important to keep in mind when we define the path signature in Section 3.

Finally, we briefly discuss the Riemannian structure of Lie groups. Recall that a *Riemannian metric* on a smooth manifold  $M$  is the assignment of an inner product  $\langle \cdot, \cdot \rangle_p$  to the tangent space  $T_p M$  for every point  $p \in M$ , which varies smoothly. Specifically, this means that if  $X, Y$  are smooth vector fields defined on a neighborhood of  $p$ , then the map  $p \mapsto \langle X_p, Y_p \rangle_p$  is smooth. On a Lie group, we often want a Riemannian metric that is compatible with the algebraic structure of  $G$ . A Riemannian metric is *left-invariant* if

$$\langle X, Y \rangle_g = \langle L_{h*} X, L_{h*} Y \rangle_{hg}$$

for all  $g, h \in G$  and  $X, Y \in T_g G$ , and a right-invariant Riemannian metric is defined similarly. Such left-invariant metrics can simply be defined on the tangent space at the identity.

**Proposition 6** *There is a one-to-one correspondence between left-invariant metrics on a Lie group  $G$  and inner products on its Lie algebra  $\mathfrak{g}$ .*

Namely, evaluating the inner product  $\langle \cdot, \cdot \rangle_g$  simply corresponds to viewing the tangent vectors as elements of the identity, and then evaluating the chosen inner product on  $\mathfrak{g}$ . We will assume that all Riemannian metrics under discussion are left-invariant, and simply call them Riemannian metrics.

## 2.2 Paths on Lie groups

A Riemannian metric  $\langle \cdot, \cdot \rangle$ , where we now omit the subscript  $g$  since it is left-invariant, provides a notion of length for piecewise regular paths on  $G$ . Suppose  $\gamma \in PG$ . Then the *length* of  $\gamma$  is defined to be

$$\ell(\gamma) := \int_0^1 \sqrt{\langle \gamma'_t, \gamma'_t \rangle} dt.$$

This allows us to define a metric on the Lie group. If  $g_1, g_2 \in G$ , then the distance between  $g$  and  $h$  is defined to be the infimum length of paths connecting  $g_1$  and  $g_2$ ,

$$d(g_1, g_2) := \inf \{ \ell(\gamma) : \gamma \in PG, \gamma_0 = g_1, \gamma_1 = g_2 \}.$$

Note that since the Riemannian metric is left invariant, this metric is also left invariant,

$$d(hg_1, hg_2) = d(g_1, g_2),$$

for all  $h \in G$ . The more familiar notion of length in the path signature literature is the 1-variation of a path.

**Definition 7** *Suppose  $(X, d_X)$  is a metric space and let  $\gamma \in PX$ . The 1-variation of  $\gamma$  on  $[0, 1]$  is defined as*

$$|\gamma|_{1-var} = \sup_{(t_i)} \sum_i d_X(\gamma_{t_i}, \gamma_{t_{i+1}}), \quad (1)$$

where the sum is taken over all partitions  $0 = t_1 \leq \dots \leq t_n = 1$  of  $[0, 1]$ .

Using the metric induced by the Riemannian metric, we may consider the 1-variation length of paths in  $G$ . Under the piecewise regular hypothesis, these two lengths are equivalent.

**Lemma 8 (Burtscher (2015))** *Let  $\gamma \in PG$ . We have  $\ell(\gamma) = |\gamma|_{1-var}$ .*

At this point, in the case of paths on Euclidean space, we may use the 1-variation to define a metric on  $P\mathbb{R}_0^N$ , which are the paths which start at the origin. Given a Lie group  $G$  with a left-invariant Riemannian metric, we could follow the same procedure to obtain a metric space structure on  $PG_e$ . However, this is not the metric space structure on  $PG$  that is the most compatible with the path signature. We will defer this discussion until Section 3.3.

The path space  $P\mathbb{R}^N$  is endowed with a vector space structure since  $\mathbb{R}^N$  itself is a vector space. Similarly, we can endow  $PG$  with a group structure by pointwise multiplication, where the identity is the constant path at the identity, and the inverse to a path  $\gamma \in PG$  is the pointwise inverse. However, we are missing a notion of scaling for paths in  $PG$ , and such an operation is important to have in machine learning, since algorithms may require normalization of data. Such a scaling is obtained by proving a correspondence between paths in  $G$  and paths in  $\mathfrak{g}$ , and then transferring the scaling operation from  $\mathfrak{g}$  to  $G$ .

This is done by considering paths on  $G$  from the point of view of differential equations. We have the following existence and uniqueness theorem for first order ordinary differential equations. Let  $\overline{P}\mathfrak{g}$  denote the space of piecewise continuous paths  $\gamma : [0, 1] \rightarrow \mathfrak{g}$  which are right continuous, meaning  $\lim_{t \downarrow t_0} \gamma_t = \gamma_{t_0}$ .

**Theorem 9** *Let  $f \in \overline{P}\mathfrak{g}$ , so that  $f : [0, 1] \rightarrow \mathfrak{g}$  is piecewise continuous and right continuous, where we consider elements of  $\mathfrak{g}$  as left-invariant vector fields. Then, the solution of the first order ODE*

$$\gamma'_t = f_t(\gamma_t), \quad \gamma_0 = g \quad (2)$$

*exists and is unique.*

Note that in this theorem, we consider a function  $\gamma : [0, 1] \rightarrow G$  to be a solution of this ODE if the differential equation holds at all points except the points of discontinuity of  $f$ . This implies that we can represent piecewise regular paths in  $G$  as paths in the Lie algebra  $\mathfrak{g}$ , along with its initial point. Let  $PG_g \subset PG$  be defined as

$$PG_g = \{\gamma \in PG : \gamma_0 = g\}.$$

**Corollary 10** *Suppose  $G$  is a Lie group and  $\mathfrak{g}$  its Lie algebra. The map  $\Psi_g : \overline{P\mathfrak{g}} \rightarrow PG_g$ , which takes  $f \in \overline{P\mathfrak{g}}$  to the solution of the ODE in Equation 2 with initial condition  $\gamma_0 = g$ , is a bijection.*

**Proof** Firstly, the map  $\Psi_g$  is well defined by the existence and uniqueness theorem above. The inverse to  $\Psi_g$  can be defined by taking the derivative at every differentiable point. Suppose  $\gamma \in PG_g$ , and let  $d(\gamma) \subset [0, 1]$  denote the set of points such that  $\gamma$  is differentiable. Note that  $[0, 1] - d(\gamma)$  is a finite set since  $\gamma$  is piecewise regular. Now, define  $\Psi_g^{-1}(\gamma)(t) = \gamma'_t$  for all  $t \in d(\gamma)$ , and at the nondifferentiable points by right continuity

$$\Psi_g^{-1}(\gamma)(t) = \lim_{s \downarrow t} \gamma'_s.$$

This map is well defined:  $\Psi_g^{-1}(\gamma)(t)$  is continuous for every  $t \in [0, 1] - d(\gamma)$ , and right continuous by definition. ■

We can view  $\overline{P\mathfrak{g}}$  as a Lie algebra, with pointwise vector space operations, and pointwise Lie bracket. Because the group structure of  $PG$  and the Lie algebra structure of  $\overline{P\mathfrak{g}}$  are defined pointwise, the map  $\Psi_g$  is compatible with Lie algebra morphisms induced by Lie group morphisms. Namely, if  $F : G \rightarrow H$  is a Lie group morphism, we obtain a group homomorphism  $F : PG \rightarrow PH$  by applying the map pointwise. Analogously, if  $F_* : \mathfrak{g} \rightarrow \mathfrak{h}$  is the induced Lie algebra morphism, we obtain a Lie algebra morphism  $F_* : \overline{P\mathfrak{g}} \rightarrow \overline{P\mathfrak{h}}$ . The following lemma is immediate since the group structure on  $PG$  and the Lie algebra structure on  $\overline{P\mathfrak{g}}$  are defined pointwise.

**Lemma 11** *Suppose  $F : G \rightarrow H$  is a morphism of Lie groups, and  $F_* : \mathfrak{g} \rightarrow \mathfrak{h}$  is the induced morphism of Lie algebras. Then the following diagram commutes*

$$\begin{array}{ccc} \overline{P\mathfrak{g}} & \xrightarrow{F_*} & \overline{P\mathfrak{h}} \\ \Psi_g^{-1} \uparrow & & \downarrow \Psi_{F(g)} \\ PG_g & \xrightarrow{F} & PH_{F(g)}. \end{array}$$

The map  $\Psi_g$  allows us to view paths on Lie groups as paths in a linear space, while retaining all first order differential information. We can use the fact that many operations for paths on  $\mathbb{R}^N$  are defined via operations on the Lie algebra, and thus generalize these operations to Lie groups.

For a path  $\alpha \in P\mathbb{R}^N$  and  $\lambda \geq 0$ , denote the *vector space scaling* operation as

$$(\lambda\alpha)_t := \lambda\alpha_t.$$

However, another way of viewing the scaling operation for paths that begin at the origin is by scaling in the Lie algebra. Suppose  $\lambda \geq 0$ , and denote the vector space scaling in a Lie algebra  $\mathfrak{g}$  by  $c_\lambda : \mathfrak{g} \rightarrow \mathfrak{g}$ .

**Lemma 12** *Let  $\alpha \in P\mathbb{R}_0^N$ . Then*

$$\lambda\alpha = \Psi_0 \circ c_\lambda \circ \Psi_0^{-1}(\alpha).$$

**Proof** In  $\mathbb{R}^N$ , the map  $\Psi_0$  is simply integration in  $\mathbb{R}^N$ , and  $\Psi_0^{-1}$  is differentiation. Thus, we have

$$\begin{aligned} (\Psi_0 \circ c_\lambda \circ \Psi_0^{-1}(\alpha))_t &= \int_0^t \lambda \alpha'_s ds \\ &= \lambda \int_0^t \alpha'_s ds \\ &= (\lambda \alpha)_t. \end{aligned}$$

■

We use this fact as motivation to define scaling on Lie groups.

**Definition 13** *Suppose  $G$  is a Lie group and  $\mathfrak{g}$  its Lie algebra. Let  $\gamma \in PG$  and  $\lambda \geq 0$ . We define the Lie algebra scaling of  $\gamma$  by  $\lambda$  to be*

$$\lambda \cdot \gamma := \Psi_{\gamma_0} \circ c_\lambda \circ \Psi_{\gamma_0}^{-1}(\gamma). \quad (3)$$

**Remark 14** *We highlight three important differences between vector space scaling for paths in  $\mathbb{R}^N$  and Lie algebra scaling for paths in an arbitrary Lie group  $G$ , and provide a reason for each.*

1. *Returning to the setting of paths in  $\mathbb{R}^N$ , the two notions of scaling differ slightly when the path does not start at the origin. If we have  $\alpha \in P\mathbb{R}^N$  such that  $\alpha_0 = x$ , then  $(\lambda \alpha)_0 = \lambda x$ , while  $(\lambda \cdot \alpha)_0 = x$ . However, if we align the initial points, the paths coincide,*

$$(\lambda \alpha) - \lambda x = (\lambda \cdot \alpha) - x.$$

*This difference is due to the fact that arbitrary Lie groups do not have a natural scaling operation. However, if our Lie group was equipped with a suitable scaling operation, such as a Carnot group (Le Donne, 2017), then we would be able to do define a scaling operation that coincides with the vector space scaling in  $P\mathbb{R}^N$ .*

2. *We have only defined scaling by a nonnegative number. Definition 13 could be extended to all real numbers  $\lambda$  without any changes, but the interpretation of negative scaling is more difficult in arbitrary Lie groups. For a path  $\alpha \in P\mathbb{R}_0^N$ , scaling by  $\lambda = -1$  simply produces the pointwise inverse of a path. However, this is not the case in a general Lie group. For example, let  $X, Y \in \mathfrak{g}$  and consider the piecewise path*

$$\gamma_t = \begin{cases} e^{2tX} & : t \in [0, \frac{1}{2}) \\ e^X e^{(2t-1)Y} & : t \in [\frac{1}{2}, 1]. \end{cases}$$

*Here, we have  $\gamma_1 = e^X e^Y$  and  $(-1 \cdot \gamma)_1 = e^{-X} e^{-Y}$ , which are not inverses in general. Thus, we see that the obstruction to this interpretation is the noncommutativity of arbitrary Lie groups. However, in the setting of abelian Lie groups, such an interpretation would hold.*

3. By definition, the vector space scaling in  $P\mathbb{R}^N$  obeys the distributive law:  $\lambda(\alpha + \beta) = (\lambda\alpha) + (\lambda\beta)$  for  $\alpha, \beta \in P\mathbb{R}^N$  and  $\lambda \in \mathbb{R}$ . In other words, the vector space scaling is a pointwise Lie group homomorphism for  $\mathbb{R}^N$ . However,  $c_\lambda : \mathfrak{g} \rightarrow \mathfrak{g}$  is not a morphism of Lie algebras in general since  $c_\lambda([X, Y]) = \lambda[X, Y] \neq \lambda^2[X, Y] = [c_\lambda X, c_\lambda Y]$ . Thus, it cannot be the induced map of an underlying Lie group homomorphism for  $G$ , so the Lie algebra scaling for  $G$  is not distributive,  $\lambda \cdot (\alpha\beta) \neq (\lambda \cdot \alpha)(\lambda \cdot \beta)$ , in general. In the case of an abelian Lie group  $H$ , the associated Lie algebra  $\mathfrak{h}$  is abelian so that  $[X, Y] = 0$  for all  $X, Y \in \mathfrak{h}$ , and thus Lie algebra scaling can be viewed as a pointwise Lie group morphism.

Due to these remarks, we must keep in mind that the scaling operation for paths in Lie groups is not compatible with the algebraic structure of  $G$ .

### 2.3 Discrete time series on $G$

In this subsection, we will consider the interpretation of discrete time series on an arbitrary Lie group  $G$ , and also discuss derivative computations for these discrete time series. We will continue the theme of comparison with the corresponding notions in  $\mathbb{R}^N$ .

**Remark 15** Here, we will assume that discrete time series are uniformly sampled at integer times. This does not result in any loss of generality due to the reparametrization invariance of the path signature, given in Proposition 22.

Let  $T \in \mathbb{N}$  and  $\hat{x} : [T + 1] \rightarrow \mathbb{R}^N$  be a discrete time series in  $\mathbb{R}^N$  of length  $T + 1$ . There is a natural interpretation of  $\hat{x}$  as a continuous time series  $x : [T + 1] \rightarrow \mathbb{R}^N$  by linear interpolation between points. Namely, it is the interpolation with a constant derivative between the discrete points defined in  $\hat{x}$ . This is the interpretation that we implicitly take when we compute derivatives of discrete time series by finite differences  $\hat{x}'_i = \hat{x}_{i+1} - \hat{x}_i$  to get the discrete derivative  $\hat{x}' : [T] \rightarrow \mathbb{R}^N$ . Additionally, we can think about the continuous path  $x$  as a geodesic interpolation of the discrete path  $\hat{x}$ .

However, the interpretation is more subtle in the case of arbitrary Lie groups. Suppose we have a discrete time series in  $G$ , which we denote by  $\hat{\gamma} : [T + 1] \rightarrow G$ . We wish to obtain an interpolation such that the derivative, when viewed in the Lie algebra  $\mathfrak{g}$ , is constant between adjacent points. This can be achieved by taking the logarithm of the difference between adjacent points. We define the *discrete derivative*  $\hat{\gamma}' : [T] \rightarrow \mathfrak{g}$  of a discrete Lie group valued path by

$$\hat{\gamma}'_i := \log(\hat{\gamma}_i^{-1} \hat{\gamma}_{i+1}) \in \mathfrak{g}. \quad (4)$$

Then, we can define the continuous interpolation  $\gamma : [0, T + 1] \rightarrow G$  using the exponential map such that for  $t \in [i, i + 1)$ , the interpolation is

$$\gamma_t := \hat{\gamma}_i \exp((t - i)\hat{\gamma}'_i).$$

We note that this construction reduces to linear interpolation in the case of  $G = \mathbb{R}^n$ . This is due to the fact that for the additive Lie group  $\mathbb{R}^N$ , the exponential and logarithm map are both the identity and are both globally defined. Additionally, the group operation

is addition, so we should interpret all of the products as sums. However, there are two essential differences between the case of arbitrary Lie groups and Euclidean space.

Firstly, for an arbitrary Lie group  $G$ , the logarithm map is only defined in a neighborhood of the identity. The two reasons the logarithm may not be defined in a larger neighborhood are the loss of injectivity and the loss of surjectivity of the exponential map. On any compact Lie group, the exponential map will not be injective at any point. In this case, we can define the logarithm to be the value closest to the origin, but non-injectivity may still occur. For example, the point antipodal to the identity in  $S^1$  has no unique logarithm since there are two paths of equal distance to the identity. However, if we perturb the target point in either direction, there exists a unique shortest path. This implies that by undersampling the underlying time series, we may infer incorrect information. The case of  $S^1$  is exactly the situation encountered in the Nyquist sampling theorem.

The exponential map is not always surjective, with the simplest examples being non-connected Lie groups. However, connected Lie groups such as  $SL(2, \mathbb{R})$  can still have non-surjective exponential maps. In these cases, discrete derivatives may not exist, and finer sampling is required so that the difference between adjacent points  $\tilde{\gamma}_i^{-1}\tilde{\gamma}_{i+1}$  is closer to the identity and has a well-defined logarithm. However, for compact Lie groups such as  $SO(3)$ , the Lie exponential map is surjective.

Secondly, the interpolation defined here may not be a geodesic connecting the two points. Suppose  $h$  is a Riemannian metric on  $G$ . In general, geodesics do not coincide with the one-parameter subgroups of  $G$ . In other words, in these cases, the Riemannian exponential map and the Lie exponential map are *not* the same. However, for bi-invariant metrics, they coincide.

**Theorem 16** *The Lie exponential map and the Riemannian exponential map at the identity agree on Lie groups with bi-invariant metrics.*

Thus, for all Lie groups equipped with bi-invariant metrics, we may continue to interpret the interpolation as a geodesic interpolation. In fact, this holds for all compact Lie groups.

**Proposition 17** *Every compact Lie group admits a bi-invariant metric.*

From this discussion, we find that for a compact Lie group  $G$ , the interpretation of discrete time series on  $G$  is similar to the case of  $\mathbb{R}^N$ , with the main difference being the non-injectivity of the exponential map.

### 3. Path signatures on Lie groups

This subsection, based on the exposition of path signatures on Euclidean space given in Giusti and Lee (2020), begins by defining the path signature for Lie groups. We show several basic properties which are well-known for path signatures on Euclidean space, culminating in the definition of tree-like equivalence for paths and the property that the signature is an injective group homomorphism. This material was originally developed by Chen (1954, 1957, 1958) and is not novel.

We then prove a signature preserving bijection between paths on an  $N$ -dimensional Lie group  $G$  and paths on  $\mathbb{R}^N$ , which highlights the extent to which the theory naturally

extends to the case of Lie groups. This result provides a Euclidean representation of Lie group valued time series, and can thus be used to apply classical Euclidean data analysis techniques to Lie group valued time series. To the authors' knowledge, this result has not previously appeared in literature.

We then consider the extension of the equivariance property of path signatures. This is followed by an interpretation of the second-level signature terms as indicators of lead-lag behavior between the directions corresponding to our choice of basis vectors for the Lie algebra  $\mathfrak{g}$ . Finally, we close this section by discussing computational aspects of the path signature for discrete time series, as well as symmetry breaking path transformations which can be used as a preprocessing step.

In this section, we use  $(e_1, \dots, e_N)$  to denote an ordered basis of  $\mathfrak{g}$  and use  $(\omega_1, \dots, \omega_N)$  to denote the dual basis of  $\mathfrak{g}^*$  such that  $\omega_i(e_j) = \delta_{i,j}$ , where  $\delta_{i,j}$  is the Kronecker delta.

### 3.1 Path signature as a group homomorphism

Let  $G$  be an  $N$ -dimensional Lie group. Recall that  $PG$  denotes the space of piecewise regular paths  $\gamma : [0, 1] \rightarrow G$ .

**Definition 18** *Let  $\gamma \in PG$ . Suppose  $\omega_1, \dots, \omega_N \in \mathfrak{g}^*$  form a basis of  $\mathfrak{g}^*$ . For  $i \in [N]$ , define a path  $S^i(\gamma)_t : [0, 1] \rightarrow \mathbb{R}$  as*

$$S^i(\gamma)_t := \int_0^t \omega_i(\gamma'_s) ds.$$

*Next, let  $I = (i_1, \dots, i_m)$  be a multi-index, where  $i_j \in [N]$ . Higher order paths  $S^I(\gamma)_t : [0, 1] \rightarrow \mathbb{R}$  are inductively defined as*

$$S^I(\gamma)_t := \int_0^t S^{(i_1, \dots, i_{m-1})}(\gamma)_s \omega_{i_m}(\gamma'_s) ds. \quad (5)$$

*The path signature of  $\gamma$  with respect to  $I$  is defined to be  $S^I(\gamma) := S^I(\gamma)_1$ .*

We can also present the definition in a non-inductive way. Let  $\Delta^m$  be the standard  $m$ -simplex

$$\Delta^m = \{(t_1, \dots, t_m) : 0 \leq t_1 < t_2 < \dots < t_m \leq 1\}.$$

By collapsing the inductive definition, we can write the path signature of  $\gamma$  with respect to  $I = (i_1, \dots, i_m)$  as

$$S^I(\gamma) = \int_{\Delta^m} \omega_{i_1}(\gamma'_{t_1}) \dots \omega_{i_m}(\gamma'_{t_m}) dt_1 \dots dt_m. \quad (6)$$

We can amalgamate the path signatures with respect to every multi-index  $I$  into an element of a tensor algebra.

**Definition 19** Suppose  $V$  is a real vector space of dimension  $N$ . The tensor algebra with respect to  $V$  is defined to be

$$T((V)) = \prod_{m \geq 0} V^{\otimes m}.$$

Suppose  $(e_1, \dots, e_N)$  is an ordered basis for  $V$ . Suppose  $\mathbf{s}, \mathbf{t} \in T((V))$ . Let  $\mathbf{t}_m \in V^{\otimes m}$  be the degree  $m$  part of  $\mathbf{t}$  and if  $I = (i_1, \dots, i_m)$  is a multi-index with  $i_j \in [N]$ , then  $\mathbf{t}^I$  is the coefficient of  $e_{i_1} \otimes \dots \otimes e_{i_m}$  in  $\mathbf{t}$ . Addition and scalar multiplication is defined element-wise:

- $(\mathbf{s} + \mathbf{t})^I = \mathbf{s}^I + \mathbf{t}^I,$
- $(\lambda \mathbf{t})^I = \lambda \mathbf{t}^I,$

and multiplication is defined by tensor multiplication

- $(\mathbf{s} \otimes \mathbf{t})^I = \sum_{j=0}^m \mathbf{s}^{(i_1, \dots, i_j)} \mathbf{t}^{(i_{j+1}, \dots, i_m)}.$

Let  $\bar{\mathfrak{g}}$  be the underlying vector space of the Lie algebra  $\mathfrak{g}$ . Let  $e_1, \dots, e_N$  be a basis for  $\bar{\mathfrak{g}}$ . We define the *path signature* of  $\Gamma \in PG$  to be

$$S(\gamma) := 1 + \sum_{m \geq 1} \sum_{|I|=m} S^I(\gamma) e_{i_1} \otimes \dots \otimes e_{i_m} \in T((\bar{\mathfrak{g}})). \quad (7)$$

**Remark 20** For path signatures defined on Euclidean space  $\mathbb{R}^N$ , we often choose the standard 1-forms  $(dx_1, \dots, dx_N)$  to be the basis of  $\mathfrak{r}$ , the Lie algebra of  $\mathbb{R}^N$ . Suppose  $\alpha \in P\mathbb{R}^N$ . We can also write our path component-wise as  $\alpha = (\alpha^1, \dots, \alpha^N)$ , where each  $\alpha^i : [0, 1] \rightarrow \mathbb{R}$ . Then, evaluation of the standard 1-forms is simply  $dx_i(\alpha'_t) = (\alpha^i)'_t$ . Thus, in the Euclidean case, the definition of the path signature reduces to

$$S^I(\alpha) = \int_{\Delta^m} (\alpha^{i_1})'_{t_1} \dots (\alpha^{i_m})'_{t_m} dt_1 \dots dt_m. \quad (8)$$

Let  $\gamma \in PG$  and  $g \in G$ . The *left translation* of  $\gamma$  by  $g$  is defined to be the path  $(g\gamma)_t := g(\gamma_t)$ , where we left translate the path  $\gamma$  by  $g$  pointwise (one can analogously define the *right translation* of a path). Similar to the case of Euclidean space, path signatures are left translation invariant and reparametrization invariant.

**Proposition 21 (Left translation invariance)** Let  $\gamma \in PG$  and  $g \in G$ . Then  $S(g\gamma) = S(\gamma)$ .

**Proof** It suffices to show that  $S^I(g\gamma) = S^I(\gamma)$  for all multi-indices  $I$ . Note that we have

$$(g\gamma)'_t = L_{g*} \gamma'_t.$$

Specifically, this implies that  $\gamma'_t$  and  $g\gamma'_t$  are represented by the same element in the Lie algebra  $\mathfrak{g}$ . Therefore for any  $\omega \in \mathfrak{g}^*$ , we have  $\omega(g\gamma'_t) = \omega(\gamma'_t)$ . Thus,  $S^I(g\gamma) = S^I(\gamma)$  for all  $I$ . ■



**Proposition 22 (Reparametrization invariance)** *Let  $\gamma : [a, b] \rightarrow G$  be a piecewise regular path, and let  $\phi : [c, d] \rightarrow [a, b]$  be a strictly increasing function. Then  $S(\gamma \circ \phi) = S(\gamma)$ .*

**Proof** This is the Change of Variables Theorem. Reparametrization invariance of the first level of the signature is given as

$$S^i(\gamma \circ \phi) = \int_c^d \omega_i((\gamma \circ \phi)'_t) dt = \int_c^d \omega_i(\gamma'_{\phi_t}) \phi'_t dt = \int_a^b \omega_i(\gamma'_\tau) d\tau = S^i(\gamma).$$

Invariance for higher order terms is shown by induction using the same argument. ■

In particular this proposition justifies our choice of only considering paths parametrized by  $[0, 1]$ , as any other path can be reparametrized into this domain. Next, we would like to understand how scaling of paths in  $G$  given in Definition 13 affects the path signature. Note that the vector space scaling in  $\bar{\mathfrak{g}}$  induces a dilation map in  $T((\bar{\mathfrak{g}}))$ . Explicitly, we define the map  $\delta_\lambda : T((\bar{\mathfrak{g}})) \rightarrow T((\bar{\mathfrak{g}}))$  as

$$\delta_\lambda \mathbf{t} := (\mathbf{t}_0, \lambda \mathbf{t}_1, \lambda^2 \mathbf{t}_2, \dots). \quad (9)$$

**Proposition 23** *Let  $\gamma \in PG$  and  $\lambda \geq 0$ . Then  $S(\lambda \cdot \gamma) = \delta_\lambda S(\gamma)$ .*

**Proof** Consider the multi-index  $I = (i_1, \dots, i_k)$ . Then,

$$\begin{aligned} S^I(\lambda \cdot \gamma) &= \int_{\Delta^k} \omega_{i_1}(\lambda \gamma'_{t_1}) \dots \omega_{i_k}(\lambda \gamma'_{t_k}) dt_1 \dots dt_k \\ &= \lambda^k \int_{\Delta^k} \omega_{i_1}(\gamma'_{t_1}) \dots \omega_{i_k}(\gamma'_{t_k}) dt_1 \dots dt_k \\ &= \lambda^k S^I(\gamma). \end{aligned}$$

■

We have seen that the group structure on  $G$  allows us to define a group structure on  $PG$  by pointwise multiplication. The group structure on  $G$  allows us to obtain another group structure on a quotient of  $PG$  where the group operation is given by concatenation. Let  $\alpha, \beta \in PG$ . The *concatenation* of  $\alpha$  and  $\beta$  is defined to be

$$(\alpha * \beta)_t = \begin{cases} \alpha_{2t} & : t \in [0, \frac{1}{2}] \\ \alpha_1(\beta_0)^{-1} \beta_{2t-1} & : t \in [\frac{1}{2}, 1]. \end{cases}$$

The inverse of a path  $\gamma$  is defined to be the same path, but in the reverse direction

$$(\gamma^{-1})_t = \gamma_{1-t}.$$

Concatenation or inversion of piecewise regular paths is still piecewise regular. In order to obtain an identity element, we must quotient out by an equivalence relation.

**Definition 24** *A path  $\gamma \in PG$  is called reducible if there exist paths  $\alpha, \beta, \zeta \in PG$  such that  $\gamma = \alpha * \zeta * \zeta^{-1} * \beta$ , up to reparametrization. The path  $\alpha * \beta$  is called a reduction of  $\gamma$ . We define the reduction of  $\zeta * \zeta^{-1}$  to be  $c_e$ , the constant path at the identity  $e \in G$ . A path  $\gamma$  is irreducible if no reduction exists. An irreducible path  $\tilde{\gamma}$  obtained by finitely many iterative reductions of a path  $\gamma$  is called an irreducible reduction of  $\gamma$ .*




Figure 1: (Left) An example of a reducible path  $\alpha * \zeta * \zeta^{-1} * \beta$ . (Right) The irreducible reduction of the path on the left.

**Theorem 25 (Chen (1958))** *Every piecewise regular path  $\gamma \in PG$  has a unique irreducible reduction up to reparametrization.*

This result allows us to define the notion of tree-like equivalence.

**Definition 26** *A path  $\gamma \in PG$  is a tree-like path if its irreducible reduction is  $c_e$ , the constant path at the identity. Two paths  $\alpha, \beta$  are tree-like equivalent,  $\alpha \sim_t \beta$ , if  $\alpha * \beta^{-1}$  is a tree-like path.*

**Remark 27** *The definition of tree-like equivalence includes translations. Indeed, suppose  $\gamma \in PG$  and  $g \in G$ . Define  $g\gamma$  and  $\gamma g$  to be the left and right translations of the path  $\gamma$  by  $g$ . Then  $\gamma \sim_t g\gamma$  since  $\gamma * (g\gamma)^{-1} = \gamma * \gamma^{-1}$  by the definition of the concatenation operator. The same holds for right translations.*

*Additionally, tree-like equivalence also includes reparametrization since the definition of reductions are reparametrization invariant.*

**Proposition 28** *Tree-like equivalence is an equivalence relation.*

**Proof** Let  $\gamma, \gamma_1, \gamma_2, \gamma_3 \in PG$ . Note that the subscript here denotes distinct paths, and does not denote the time parameter. By definition the reduction of  $\gamma * \gamma^{-1}$  is the constant path, so  $\gamma \sim_t \gamma$ .

Next, if  $\gamma = \alpha * \zeta * \zeta^{-1} * \beta$ , for paths  $\alpha, \beta, \zeta \in PG$ , then  $\gamma^{-1} = \beta^{-1} * \zeta * \zeta^{-1} * \alpha^{-1}$ . Thus, a path is reducible if and only if its inverse is reducible. Additionally, the reduction  $\beta^{-1} * \alpha^{-1}$  of  $\gamma^{-1}$  is the inverse of the reduction  $\alpha * \beta$  of  $\gamma$ . Now, suppose  $\gamma_1 \sim_t \gamma_2$  so that  $\gamma_1 * \gamma_2^{-1}$  is tree-like. By the above argument,  $\gamma_2 * \gamma_1^{-1}$  is also tree-like, so  $\gamma_2 \sim_t \gamma_1$ .

Finally, the concatenation  $\alpha * \beta$  of two tree-like paths is also tree-like, by performing all the reductions of  $\alpha$  and then performing all the reductions on  $\beta$ . Suppose  $\gamma_1 \sim_t \gamma_2$  and  $\gamma_2 \sim_t \gamma_3$ . Then,  $\gamma_1 * \gamma_3^{-1}$  is a reduction of  $(\gamma_1 * \gamma_2^{-1}) * (\gamma_2 * \gamma_3^{-1})$ , and the latter path is tree-like since it is a concatenation of two tree-like paths. By the uniqueness of irreducible reductions,  $\gamma_1 * \gamma_3^{-1}$  is tree-like. Thus,  $\gamma_1 \sim_t \gamma_3$ . ■

We can now define  $\widetilde{PG} := PG / \sim_t$  to be the space of tree-like equivalence classes of piecewise regular paths in  $G$ . We define the identity element to be  $[c_e] \in \widetilde{PG}$ , the equivalence class of the constant path at the identity. Compatibility of concatenation and inversion are implicit in the above proof, and the group axioms are easily checked. Thus, we have shown the following.

**Proposition 29** *The quotient  $\widetilde{PG}$  is a group.*

We can now state Chen's injectivity theorem.

**Theorem 30 (Chen (1958))** *Suppose  $G$  is a real Lie group. Let  $\alpha, \beta \in PG$ . Then  $S(\alpha) = S(\beta)$  if and only if  $\alpha$  and  $\beta$  are tree-like equivalent.*

Chen also showed that the signature is a group homomorphism. Namely, suppose  $\alpha, \beta \in PG$ . Chen's identity (Chen, 1954) states that

$$S(\alpha * \beta) = S(\alpha) \otimes S(\beta). \quad (10)$$

Putting the previous results together, we obtain the following characterization.

**Proposition 31** *The path signature map  $S : \widetilde{PG} \rightarrow T((\mathfrak{g}))$  is an injective group homomorphism.*

We will also require an internal multiplicative structure on the path signature coefficients which is an immediate generalization of the Euclidean path signature.

**Definition 32** *Let  $k$  and  $l$  be non-negative integers. A  $(k, l)$ -shuffle is a permutation  $\sigma$  of the set  $\{1, 2, \dots, k+l\}$  such that*

$$\sigma^{-1}(1) < \sigma^{-1}(2) < \dots < \sigma^{-1}(k)$$

and

$$\sigma^{-1}(k+1) < \sigma^{-1}(k+2) < \dots < \sigma^{-1}(k+l).$$

We denote by  $Sh(k, l)$  the set of  $(k, l)$ -shuffles. Given two finite ordered multi-indices  $I = (i_1, \dots, i_k)$  and  $J = (j_1, \dots, j_l)$ , let  $R = (r_1, \dots, r_k, r_{k+1}, \dots, r_{k+l}) = (i_1, \dots, i_k, j_1, \dots, j_l)$  be the concatenated multi-index. The shuffle product of  $I$  and  $J$  is defined to be the multiset

$$I \sqcup J = \{(r_{\sigma(1)}, \dots, r_{\sigma(k+l)}) : \sigma \in Sh(k, l)\}.$$

As an example, suppose  $I = (1, 2)$  and  $J = (2, 3)$ . Then

$$I \sqcup J = \{(1, 2, 2, 3), (1, 2, 3, 2), (2, 1, 2, 3), (2, 1, 3, 2), (2, 3, 1, 2)\}.$$

**Theorem 33** *Let  $I$  and  $J$  be multi-indices in  $[N]$ , of lengths  $k$  and  $l$  respectively, and suppose  $\gamma \in PG$ . Then*

$$S^I(\gamma)S^J(\gamma) = \sum_{K \in I \sqcup J} S^K(\gamma). \quad (11)$$

**Proof** Let  $R = (r_1, \dots, r_k, r_{k+1}, \dots, r_{k+l}) = (i_1, \dots, i_k, j_1, \dots, j_l)$ . Writing out the signature on the left side of the equation using Equation 6, we get

$$\begin{aligned} & \int_{\Delta^k} \omega_{i_1}(\gamma'_{t_1}) \dots \omega_{i_k}(\gamma'_{t_k}) dt_1 \dots dt_k \int_{\Delta^l} \omega_{j_1}(\gamma'_{t_1}) \dots \omega_{j_l}(\gamma'_{t_l}) dt_1 \dots dt_l \\ &= \int_{\Delta^k \times \Delta^l} \omega_{r_1}(\gamma'_{t_1}) \dots \omega_{r_{k+l}}(\gamma'_{t_{k+l}}) dt_1 \dots dt_{k+l}, \end{aligned}$$

and the sum on the right side is

$$\sum_{\sigma \in Sh(k,l)} \int_{\Delta^{k+l}} \omega_{\sigma(r_1)}(\gamma'_{t_1}) \dots \omega_{\sigma(r_{k+l})}(\gamma'_{t_{k+l}}) dt_1 \dots dt_{k+l}.$$

The equivalence of the two formulas is given by the standard decomposition of  $\Delta^k \times \Delta^l$  into  $(k+l)$ -simplices,

$$\begin{aligned} \Delta^k \times \Delta^l &= \{(t_1, \dots, t_{k+l}) : 0 < t_1 < \dots < t_k < 1, 0 < t_{k+1} < \dots < t_{k+l} < 1\} \\ &= \bigsqcup_{\sigma \in Sh(k,l)} \{(t_{\sigma(1)}, \dots, t_{\sigma(k+l)}) : 0 < t_1 < \dots < t_{k+l} < 1\}. \end{aligned}$$

■

### 3.2 Relationship between paths in $G$ and $\mathbb{R}^N$

In this subsection, we define a signature-preserving bijection between piecewise regular paths in  $G$  and piecewise regular paths in  $\mathbb{R}^N$  which start at the identity and origin respectively.

The idea behind the following proposition is that the path signature computation only requires the first derivative of paths. The Lie bracket is unused in the computation of path signatures, so we can simply consider the Lie algebras of Lie group as vector spaces. Thus, we can identify the underlying vector space of the Lie algebra  $\mathfrak{g}$  with the underlying vector space of the Lie algebra  $\mathfrak{r}$  of  $\mathbb{R}^N$ . We then use the correspondence  $\Psi_g : \overline{P}\mathfrak{g} \rightarrow PG_g$  between piecewise continuous paths  $\overline{P}\mathfrak{g}$  and piecewise regular paths  $PG$  given in Corollary 10, to map paths on  $G$  to paths on  $\mathbb{R}^N$ .

In the following proposition, we abuse notation and consider elements of the Lie algebras  $\mathfrak{r}$  of  $\mathbb{R}^N$  and  $\mathfrak{g}$  of  $G$  as both the tangent space at the identity, and the vector space of left-invariant vector fields. Similarly, we consider elements of the dual of the Lie algebra  $\mathfrak{r}^*$  and  $\mathfrak{g}^*$  as both the cotangent space at the identity, and the vector space of left-invariant 1-forms.

Because we will be using two different path signature functions, we will denote by  $S_{\mathbb{R}} : P\mathbb{R}^N \rightarrow T((\mathbb{R}^N))$  the path signature for  $\mathbb{R}^N$  with respect to the ordered basis of standard 1-forms  $(dx_1, \dots, dx_N)$  of  $\mathfrak{r}^*$ . We denote  $S_G : PG \rightarrow T((\mathbb{R}^N))$  to be the path signature for  $G$  with respect to a given ordered basis  $(\omega_1, \dots, \omega_N)$  of  $\mathfrak{g}^*$ .

**Proposition 34** *Suppose  $G$  is an  $N$ -dimensional Lie group with Lie algebra  $\mathfrak{g}$ . Let  $\phi : \bar{\mathfrak{r}} \rightarrow \bar{\mathfrak{g}}$  be an isomorphism of vector spaces, and  $\phi^* : \mathfrak{g}^* \rightarrow \mathfrak{r}^*$  be its dual. Let  $(dx_1, \dots, dx_N)$*

denote the standard 1-forms of  $\mathbb{R}^N$ , and define  $\omega_i = (\phi^*)^{-1}(dx_i)$ . Let  $S_G : PG \rightarrow T((\mathbb{R}^N))$  be the path signature map for  $G$  with respect to the ordered basis  $(\omega_1, \dots, \omega_N)$  of  $\mathfrak{g}^*$ . Then, there exists a bijection  $\Phi : P\mathbb{R}_0^N \rightarrow PG_e$  such that  $S_{\mathbb{R}}(\gamma) = S_G(\Phi(\gamma))$  for all  $\gamma \in P\mathbb{R}^N$ .

**Proof** The construction of the map  $\Phi$  is derived from Corollary 10. Let  $\Psi_{\mathbb{R}} : \overline{P}\mathfrak{r} \rightarrow P\mathbb{R}_0^N$  and  $\Psi_G : \overline{P}\mathfrak{g} \rightarrow PG_e$  be the bijections from Corollary 10 for  $\mathbb{R}^N$  and  $G$  respectively. Now, define  $\Phi$  by

$$\Phi : P\mathbb{R}_0^N \xrightarrow{\Psi_{\mathbb{R}}^{-1}} \overline{P}\mathfrak{r} \xrightarrow{\phi} \overline{P}\mathfrak{g} \xrightarrow{\Psi_G} PG_e.$$

The idea is that we start with a path  $\gamma \in P\mathbb{R}_0^N$ , and apply the following maps:

1.  $\Psi_{\mathbb{R}}^{-1}$  : take the derivative  $\gamma'$  to obtain a path in  $\mathfrak{r}$
2.  $\phi$  : identify the underlying vector space of  $\mathfrak{r}$  with  $\mathfrak{g}$
3.  $\Psi_G$  : solve the differential equation (Equation 2) with the identity initial condition to obtain a path in  $G$ .




Figure 2: An example of the function  $\Phi$  when we take  $G = S^1 \times S^1$ . (Left) A path  $\gamma \in P\mathbb{R}_0^2$ . (Middle) The derivative of  $\gamma$  as a path in  $\mathfrak{r}$  or  $\mathfrak{g}$ . (Right) The corresponding path  $\Phi(\gamma)$ .

Because all three maps are bijective,  $\Phi$  is also bijective. To show that the signatures are invariant under this mapping, let  $I = (i_1, \dots, i_k)$  and  $\gamma \in P\mathbb{R}_0^N$ . The path signature of  $\gamma$  with respect to  $I$  is

$$S_{\mathbb{R}}^I(\gamma) = \int_{\Delta^k} dx_{i_1}(\gamma_{t_1}) \dots dx_{i_k}(\gamma_{t_k}) dt_1 \dots dt_k.$$

Note that the derivative of  $\Phi(\gamma)$  is given by  $\Phi(\gamma)'_t = \phi(\gamma'_t)$  and thus, the path signature of  $\Phi(\gamma)$  with respect to  $I$  is

$$\begin{aligned} S_G^I(\Phi(\gamma)) &= \int_{\Delta^k} \omega_{i_1}(\phi(\gamma'_{t_1})) \dots \omega_{i_k}(\phi(\gamma'_{t_k})) dt_1 \dots dt_k \\ &= \int_{\Delta^k} \phi^*(\omega_{i_1})(\gamma'_{t_1}) \dots \phi^*(\omega_{i_k})(\gamma'_{t_k}) dt_1 \dots dt_k \\ &= S_{\mathbb{R}}^I(\gamma). \end{aligned}$$

The final equality holds because the dual isomorphism  $\phi^*$  takes  $\omega_i$  to  $dx_i$ . Thus,  $S_{\mathbb{R}}^I(\gamma) = S_G^I(\Phi(\gamma))$  for all  $\Gamma \in P\mathbb{R}_e^N$  and all multi-indices  $I$ .  $\blacksquare$

### 3.3 Stability of the path signature

In this section, we will discuss the stability of path signatures, which is of crucial importance in machine learning applications. Because we will only be interested in truncated path signatures in applications, we will consider the truncated signature map  $S_M : PG \rightarrow T^{\leq M}(\bar{\mathfrak{g}})$ , where

$$T^{\leq M}(\bar{\mathfrak{g}}) = \bigoplus_{m=0}^M (\bar{\mathfrak{g}})^{\otimes m},$$

where  $S_M$  only retains information about the first  $M$  levels of the path signature. In addition, we define the projection map

$$\pi_m : T((\bar{\mathfrak{g}})) \rightarrow \bar{\mathfrak{g}}^{\otimes m}$$

to a particular tensor level. Such a map can also be defined on the truncated tensor algebra  $\pi_m : T^{\leq M}(\bar{\mathfrak{g}}) \rightarrow \bar{\mathfrak{g}}^{\otimes m}$ , and we denote all such maps in the same manner.

By stability of the path signature, we mean to say that the truncated signature map  $S_M : PG \rightarrow T^{\leq M}(\bar{\mathfrak{g}})$  is Lipschitz continuous. In order to discuss such a notion, we must provide both  $PG$  and  $T^{\leq M}(\bar{\mathfrak{g}})$  with metrics. We begin with the metric on  $T((\bar{\mathfrak{g}}))$ , which is required in Section 4.2 and is analogous to the metric on  $T^{\leq M}(\bar{\mathfrak{g}})$ .

Recall that a basis  $(e_1, \dots, e_N)$  of  $\bar{\mathfrak{g}}$  induces a natural inner product on  $\bar{\mathfrak{g}}$  by defining the basis to be orthonormal. This extends to an inner product structure on  $\bar{\mathfrak{g}}^{\otimes m}$ , and given  $\mathbf{t}_k \in \bar{\mathfrak{g}}^{\otimes m}$ , we will denote the norm by  $\|\mathbf{t}\|_m$ . In addition, this also extends to an inner product on  $T^{\leq M}(\bar{\mathfrak{g}})$ . Let  $\mathbf{s}, \mathbf{t} \in T^{\leq M}(\bar{\mathfrak{g}})$ . Such an inner product and norm are defined to be

$$\langle \mathbf{s}, \mathbf{t} \rangle = \sum_{m=0}^M \sum_{|I|=m} \mathbf{s}^I \mathbf{t}^I, \quad \|\mathbf{t}\| = \sqrt{\sum_{m=0}^M \sum_{|I|=m} (\mathbf{t}^I)^2}. \quad (12)$$

Then, we can use the norm to define a metric on both  $\bar{\mathfrak{g}}^{\otimes m}$  and  $T^{\leq M}(\bar{\mathfrak{g}})$ . Namely, given  $\mathbf{s}_m, \mathbf{t}_m \in \bar{\mathfrak{g}}^{\otimes m}$  and  $\mathbf{s}, \mathbf{t} \in T^{\leq M}(\bar{\mathfrak{g}})$ , we have

$$\begin{aligned} d_m(\mathbf{s}_m, \mathbf{t}_m) &= \|\mathbf{s}_m - \mathbf{t}_m\|_m \\ d(\mathbf{s}, \mathbf{t}) &= \|\mathbf{s} - \mathbf{t}\|. \end{aligned}$$

Note that this norm on the tensor algebra extends to  $T((\bar{\mathfrak{g}}))$ , where the inner product and norm for  $\mathbf{s}, \mathbf{t} \in T((\bar{\mathfrak{g}}))$  are defined as in (12) with  $M \rightarrow \infty$ . In this case, the inner product and norm may be infinite. However, image of the path signature lies in a subalgebra of  $T((\bar{\mathfrak{g}}))$  where the norm is finite. Namely, we define

$$T_1((\bar{\mathfrak{g}})) := \{\mathbf{t} \in T((\bar{\mathfrak{g}})) \quad : \quad \|\mathbf{t}\| < \infty, \mathbf{t}_0 = 1\}. \quad (13)$$

We now show that the signature of any path  $\gamma \in PG$  lies in this subspace.

**Lemma 35** *Let  $\gamma \in PG$ . Then  $\|S(\gamma)\| < \infty$ .*

**Proof** Without loss of generality, we suppose that  $\gamma$  is parametrized by length such that it is defined as  $\gamma : [0, L] \rightarrow G$ , where  $L$  is the length, and  $\|\gamma'_t\| = 1$  for all differentiable  $t$ ; this assumption is valid due to the reparametrization invariance of the signature. We will inductively bound each signature term. At the first level, we have

$$\begin{aligned} |S^i(\gamma)(t)| &\leq \int_0^t |\omega_i(\gamma'_s)| ds \\ &\leq t, \end{aligned}$$

using the fact that  $|\omega_i(\gamma'_t)| \leq \|\gamma'_t\| = 1$ . Assume that for any multi-index  $I = (i_1, \dots, i_{m-1})$  of length  $m - 1$ , we have

$$|S^I(\gamma)(t)| \leq \frac{t^{m-1}}{(m-1)!}.$$

Now consider the multi-index  $I = (i_1, \dots, i_m)$  of length  $m$ . Using the induction hypothesis, and the recursive definition of the signature, we have

$$\begin{aligned} |S^I(\gamma)(t)| &\leq \int_0^t |S^{(i_1, \dots, i_{m-1})}(s)| |\omega_{i_m}(\gamma'_s)| ds \\ &\leq \int_0^t \frac{s^{m-1}}{(m-1)!} ds = \frac{t^m}{m!}. \end{aligned}$$

Therefore, for any multi-index  $I$  of length  $m$ , we have

$$|S^I(\gamma)| \leq \frac{L^m}{m!}, \tag{14}$$

and the norm of  $S(\gamma)$  is bounded by

$$\begin{aligned} \|S(\gamma)\|^2 &= \sum_{m=0}^{\infty} \sum_{|I|=m} (S^I(\gamma))^2 \\ &\leq \sum_{m=0}^{\infty} \frac{N^m L^{2m}}{(m!)^2} < \infty \end{aligned}$$

where the last inequality uses the fact that there are  $N^m$  multi-indices of length  $m$ . ■

Next, we consider a metric structure on  $PG$ . We mentioned in Section 2.2 that a metric on  $P\mathbb{R}_0^N$  can be obtained by using the 1-variation of paths. Namely, given  $\alpha, \beta \in P\mathbb{R}_0^N$ , we may consider the distance between these two paths as  $|\alpha - \beta|_{1-var}$ , where the difference is performed pointwise. Such an approach could be used for  $PG$  in theory. Now suppose  $\alpha, \beta \in PG_e$ , we can use the metric on  $G$  to define the distance to be  $|\beta^{-1}\alpha|_{1-var}$ , where the inversion and multiplication are both performed pointwise. However, this notion of distance is not well-suited for the path signature.

The main reason for this is that the computation of  $|\beta^{-1}\alpha|_{1-var}$  depends fundamentally on the adjoint action of the Lie group on the Lie algebra, which is governed by the Lie bracket. Namely, the adjoint action is trivial if and only if the Lie bracket is zero. However, the path signature ignores the Lie bracket structure, so the prospect of Lipschitz continuity of the signature with respect to this metric is problematic.

We therefore consider a different metric. Note that our path signature computations have consistently been performed on the underlying vector space of the Lie algebra  $\bar{\mathfrak{g}}$ , so it seems natural to directly define a metric using the derivatives  $\alpha', \beta' \in \hat{P}\mathfrak{g}$ . One such notion of a distance would be the  $L^1$  distance between these derivatives

$$\|\alpha' - \beta'\|_{L^1} = \int_0^1 \|\alpha'_t - \beta'_t\|_{\mathfrak{g}} dt$$

which in particular does not use the Lie bracket structure. In fact this  $L^1$  distance is exactly the 1-variation of the corresponding paths  $\Phi^{-1}(\alpha), \Phi^{-1}(\beta) \in P\mathbb{R}_0^N$ , given by the bijection in Proposition 34. Thus, we can define the metric on  $PG_e$  to be

$$d_{\mathbb{R}}(\alpha, \beta) := |\Phi^{-1}(\alpha) - \Phi^{-1}(\beta)|_{1-var}.$$

Note that equipped with this metric, the map  $\Phi$  is trivially an isometry.

**Lemma 36** *Suppose  $\Phi : P\mathbb{R}_0^N \rightarrow PG_e$  is the map defined in Proposition 34. Suppose  $P\mathbb{R}_0^N$  is equipped with the 1-variation metric, and  $PG_e$  is equipped with the metric  $d_{\mathbb{R}}$ . Then,  $\Phi$  is an isometry.*

Using this isometry, stability for Lie group path signatures is a direct corollary of stability for Euclidean path signatures.

**Proposition 37 (Friz and Victoir (2010))** *Let  $\alpha, \beta \in P\mathbb{R}_0$ , and let*

$$L \geq \max\{|\alpha|_{1-var}, |\beta|_{1-var}\}.$$

*Then, for all  $k \geq 1$ , there exists some  $C_k > 0$  such that*

$$\left\| \pi_k(S(\alpha) - S(\beta)) \right\|_k \leq C_k L^{k-1} |\alpha - \beta|_{1-var}.$$

**Corollary 38** *Let  $\alpha, \beta \in PG_e$ , and let  $L \geq \max\{|\alpha|_{1-var}, |\beta|_{1-var}\}$ . Then, for all  $k \geq 1$ , there exists some  $C_k > 0$  such that*

$$\left\| \pi_k(S(\alpha) - S(\beta)) \right\|_k \leq C_k L^{k-1} d_{\mathbb{R}}(\alpha, \beta).$$

### 3.4 Equivariance of the path signature

At this point, a natural question to consider is how do path signatures behave under Lie group morphisms? Let  $G_1$  and  $G_2$  be Lie groups of dimensions  $N_1$  and  $N_2$  respectively. Given a Lie group morphism  $F : G_1 \rightarrow G_2$ , we have an induced Lie algebra morphism  $F_* : \mathfrak{g}_1 \rightarrow \mathfrak{g}_2$  between the corresponding Lie algebras. In particular, all Lie algebra morphisms are linear transformations, so if we forget the Lie bracket, this results in a map  $F_* : \bar{\mathfrak{g}}_1 \rightarrow \bar{\mathfrak{g}}_2$



between the underlying vector spaces. Because linear transformations induce maps on tensor products of the space  $F_*^{\otimes m} : \bar{\mathfrak{g}}_1^{\otimes m} \rightarrow \bar{\mathfrak{g}}_2^{\otimes m}$ , we also get an induced map of algebras between tensor algebras

$$F_* : T((\bar{\mathfrak{g}}_1)) \rightarrow T((\bar{\mathfrak{g}}_2)).$$

If  $(e_1, \dots, e_{N_1})$  is an ordered basis for  $\mathfrak{g}_1$  and  $(f_1, \dots, f_{N_2})$  is an ordered basis for  $\mathfrak{g}_2$ , then we can write  $F_* : \mathfrak{g}_1 \rightarrow \mathfrak{g}_2$  as an  $N_2 \times N_1$  matrix in terms of these bases, which we call  $M$ . We can describe the action of  $F_*$  in the tensor algebra using this matrix. Let  $\mathbf{t} \in T((\bar{\mathfrak{g}}_1))$ . In general, the action on the order  $m$  elements  $\mathbf{t}_m \in \bar{\mathfrak{g}}^{\otimes m}$  is a tensor-matrix multiplication, as described in Pfeffer et al. (2019), in which all  $m$  sides of the tensor  $\mathbf{t}_m$  are multiplied by the matrix  $M$ . This can be written out as

$$F_* \mathbf{t} = \sum_{m=0}^{\infty} \sum_{|I|=m} \mathbf{t}^I (M e_{i_1}) \otimes (M e_{i_2}) \otimes \dots \otimes (M e_{i_k}).$$

The low order tensors can be written out in usual matrix notation. Consider  $\mathbf{t}_1$  as a column vector. The action on first order elements is matrix multiplication,

$$(F_* \mathbf{t})_1 = M \mathbf{t}_1.$$

Considering  $\mathbf{t}_2$  as a matrix, the action on the second order elements is conjugation by  $M$ ,

$$(F_* \mathbf{t})_2 = M \mathbf{t}_2 M^\top.$$

For higher orders, we can no longer use matrix notation, so we explicitly define the action for a given index. Let  $J = (j_1, \dots, j_n)$  be a multi-index where  $j_k \in [N_2]$ . Then, the element of  $F_* \mathbf{t}$  corresponding to the multi-index  $J$  is

$$(F_* \mathbf{t})^J = \sum_{i_1=1}^{N_1} \sum_{i_2=1}^{N_1} \dots \sum_{i_n=1}^{N_1} \mathbf{t}^{(i_1, \dots, i_n)} M_{j_1, i_1} M_{j_2, i_2} \dots M_{j_n, i_n}.$$

The following is a generalization of the equivariance of the path signature in Euclidean space, which is discussed in Friz and Victoir (2010) and Pfeffer et al. (2019). Here, suppose  $(\omega_1, \dots, \omega_{N_1})$  is the dual basis to  $(e_1, \dots, e_{N_1})$  and  $(\nu_1, \dots, \nu_{N_2})$  is the dual basis to  $(f_1, \dots, f_{N_2})$ .

**Proposition 39** *Let  $G_1$  and  $G_2$  be Lie groups, with Lie algebras  $\mathfrak{g}_1$  and  $\mathfrak{g}_2$  respectively. Suppose  $F : G_1 \rightarrow G_2$  is a Lie group morphism and  $\gamma \in P(G_1)$ . Then*

$$S(F\gamma) = F_* S(\gamma).$$

**Proof** The proof of this claim is simply due to the linearity of integrals and 1-forms. Consider the multi-index  $J = (j_1, \dots, j_m)$ . Then,

$$S^J(F\gamma) = \int_{\Delta^m} \nu_{j_1}(F_* \gamma'_{t_1}) \dots \nu_{j_m}(F_* \gamma'_{t_m}) dt_1 \dots dt_m.$$

Consider a single factor in the integrand. Using the basis  $(e_1, \dots, e_{N_1})$  for  $\mathfrak{g}$ , write the derivative  $\gamma'$  as

$$\gamma'_t = \sum_{i=1}^{N_1} \dot{c}_t^i e_i$$

where  $c^i : [0, 1] \rightarrow \mathbb{R}$  are the component paths. Then, since  $\nu_j(F_*\gamma'_t)$  denotes the  $j^{\text{th}}$  component of  $F_*\gamma'_t$ , we can write this as

$$\nu_j(F_*\gamma'_t) = \sum_{i=1}^{N_1} M_{j,i} \omega_i(\gamma'_t).$$

Substituting this back into the formula for  $S^J(F\gamma)$ , we get

$$\begin{aligned} S^J(F\gamma) &= \sum_{i_1=1}^{N_1} \dots \sum_{i_m=1}^{N_1} (M_{j_1, i_1} \dots M_{j_m, i_m}) \int_{\Delta^m} \omega_{i_1}(\gamma'_t) \dots \omega_{i_m}(\gamma'_t) dt_1 \dots dt_m \\ &= \sum_{i_1=1}^{N_1} \dots \sum_{i_m=1}^{N_1} (M_{j_1, i_1} \dots M_{j_m, i_m}) S^{(i_1, \dots, i_m)}(\gamma) \\ &= (F_*S(\gamma))^J. \end{aligned}$$

■

### 3.5 Lead-lag relationships

For path signatures defined on Euclidean space, a certain linear combination of second degree signature terms provides a reparametrization invariant indicator of lead-lag behavior in cyclic real-valued time series, as initially introduced in Baryshnikov and Schlafly (2016). In this subsection, we will extend this interpretation to time series valued in Lie groups.

A *cyclic* time series in  $G$  is one which is periodic up to a time-reparametrization. More precisely, a time series  $\gamma$  is cyclic if it factors through the circle,

$$\gamma : [0, 1] \xrightarrow{\phi} S^1 \rightarrow G,$$

with  $\phi$  monotone and winding around the circle at least twice, the winding condition enforcing nontrivial repetition.

Consider the interpretation for Euclidean paths in  $\mathbb{R}^2$ . Suppose  $\gamma = (\gamma^1, \gamma^2) \in P\mathbb{R}^2$  is a cyclic time series. We say that the component  $\gamma^1$  exhibits a cyclic leading behavior with respect to the component  $\gamma^2$  if the following two conditions hold:

1. when  $\gamma^1$  is large (small), then  $\gamma^2$  is increasing (decreasing), and
2. when  $\gamma^2$  is large (small), then  $\gamma^1$  is decreasing (increasing).

The first condition can be viewed as a reparametrization invariant definition of a time series  $\gamma^1$  leading another time series  $\gamma^2$ . The second condition is used because we are working with cyclic time series, so we also consider the negative influence of  $\gamma^2$  on  $\gamma^1$ . We may think of this phenomena as a feedback loop in which  $\gamma^1$  positively influences  $\gamma^2$  and  $\gamma^2$  negatively influences  $\gamma^1$ . The standard example of such behavior is  $\gamma_t = (\sin(t), -\cos(t))$ .

To quantify what we mean by large or small in the two conditions above, we translate the time series such that  $\gamma_0 = (0, 0)$  and interpret large (small) to mean positive (negative). Then, a measure for these two conditions are given by  $S^{1,2}(\gamma)$  and  $-S^{2,1}(\gamma)$  respectively,

$$S^{1,2}(\gamma) = \int_0^1 \gamma_t^1(\gamma^2)'_t dt, \quad S^{2,1}(\gamma) = \int_0^1 \gamma_t^2(\gamma^1)'_t dt.$$

Thus, a measure for cyclic leading behavior can be defined as

$$A^{1,2}(\gamma) = \frac{1}{2} (S^{1,2}(\gamma) - S^{2,1}(\gamma)) = \frac{1}{2} \int_0^1 \gamma_t^1(\gamma^2)'_t - \gamma_t^2(\gamma^1)'_t dt.$$

Because the signature is translation invariant, the translation to the origin described above does not affect this measure. Moreover, if we consider a time series  $\gamma \in P\mathbb{R}^N$ , then we can consider all pairwise cyclic leading behavior between components. We can place all of this information into a matrix called the *lead matrix*,  $A(\gamma)$ , which has entries

$$A^{i,j}(\gamma) = \frac{1}{2} (S^{i,j}(\gamma) - S^{j,i}(\gamma)). \quad (15)$$

The entries  $A^{i,j}(\gamma)$  have a geometric interpretation in terms of the *signed area* of the path, as per Baryshnikov and Schlafly (2016).

An example of the second level signatures and the signed area is shown in the figure below.




Figure 3: Second level signature computations  $S^{1,2}$  (left),  $S^{2,1}$  (middle), and the signed area  $A^{1,2}$  (right). Blue represents positive area, while red represents negative area.

Returning to the setting of Lie groups, we can define the lead matrix of a path  $\gamma \in PG$  in the same manner, but the interpretation must be slightly modified. Writing out the

integral for  $S^{i,j}(\gamma)$  given a basis  $(\omega_1, \dots, \omega_N)$  of  $\mathfrak{g}^*$ , we get

$$S^{i,j}(\gamma) = \int_0^1 \left( \int_0^t \omega_i(\gamma'_s) ds \right) \omega_j(\gamma'_t) dt.$$

The inner integral is simply  $S^i(\gamma)_t$  and represents the cumulative variation of the path in the direction of  $e_i$  (the dual of  $\omega_i$ ), which is the analogue of the displacement in Euclidean space. Thus, for a cyclic time series  $\gamma \in PG$ , we say that the  $e^i$  direction exhibits cyclic leading behavior with respect to the  $e^j$  direction if the following holds:

1. **(positive influence)** when  $S^i(\gamma)_t$  is positive (negative), then  $\omega_i(\gamma'_t)$  is positive (negative), and
2. **(negative influence)** when  $S^j(\gamma)_t$  is positive (negative), then  $\omega_j(\gamma'_t)$  is negative (positive).

Thus the lead matrix, as defined in Equation 15, can be interpreted as a measure of this cyclic leading behavior for Lie group time series. An example of this interpretation is given in Section 5.1.

However, the geometric interpretation in terms of signed area is no longer available. This is because any area computation on Lie groups will require second-order differential information about the paths, but path signatures are only defined using first order differential information. This suggests that an interpretation in terms of areas on the Lie group will not be possible. However, by using Proposition 34, the value  $A^{i,j}(\gamma)$  can still be interpreted as the signed area of the corresponding path  $\Phi^{-1}(\gamma)$ , where  $\Phi$  is the bijection given in Proposition 34.

### 3.6 Topological considerations

In this section, we will consider the topological interpretation of the first level signature terms for some Lie groups. Namely, the first level signature term  $S^i$  is homotopy invariant if the differential 1-form corresponding to the basis vector  $\omega_i \in \mathfrak{g}^*$  is a closed form. For simplicity, we will assume that all paths are piecewise smooth in this section. Note in particular that the continuous interpretation of discrete time series is piecewise smooth, so the discussion in this section holds for analysis of these discrete time series.

Recall the following definition of a homotopy between paths.

**Definition 40** *Suppose  $\alpha, \beta : [0, 1] \rightarrow G$  are homotopic relative to endpoints if  $\alpha_0 = \beta_0$ ,  $\alpha_1 = \beta_1$  and there exists a continuous function  $h : [0, 1]^2 \rightarrow G$ , called a homotopy, such that*

$$h(0, t) = \alpha_t, \quad h(1, t) = \beta_t, \quad h(s, 0) = \alpha_0 = \beta_0, \quad h(s, 1) = \alpha_1 = \beta_1.$$

*We use the notation  $\alpha \simeq \beta$  if the paths  $\alpha$  and  $\beta$  are homotopic relative to endpoints.*

Loosely speaking, two paths are homotopic relative to endpoints if their endpoints coincide, and there exists a continuous deformation from one path to the other. Namely,

homotopy relative to endpoints forms an equivalence relation in  $PG$ . We say that a map  $f : PG \rightarrow \mathbb{R}$  is *homotopy invariant* if  $f(\alpha) = f(\beta)$  whenever  $\alpha \simeq \beta$ .

Recall that a differential form  $\omega$  is *closed* if its exterior derivative is trivial,  $d\omega = 0$ . By Stokes' theorem, the first level signature terms for closed forms are homotopy invariant. Indeed, let  $\alpha, \beta \in PG$ , and let  $h : [0, 1]^2 \rightarrow G$  be a homotopy between  $\alpha$  and  $\beta$ . By Stokes' theorem, we have

$$\int_{\partial h} \omega = \int_h d\omega = 0,$$

but the boundary of  $h$  is exactly  $\alpha * \beta^{-1}$ . Thus, we have

$$\int_{\alpha} \omega - \int_{\beta} \omega = \int_{\partial h} \omega = 0.$$

For left-invariant forms on Lie groups, there is a simple way to determine whether the form is closed. We begin with the invariant formula for the exterior derivative (Lee, 2003). Let  $\omega$  be a 1-form on  $G$ , and  $X, Y$  are vector fields on  $G$ , then

$$d\omega(X, Y) = X(\omega(Y)) - Y(\omega(X)) - \omega([X, Y]).$$

If in particular, if  $\omega \in \mathfrak{g}^*$  is a left-invariant 1-form and  $X, Y \in \mathfrak{g}$  are left-invariant vector fields, then this formula reduces to

$$d\omega(X, Y) = \omega([X, Y])$$

since  $\omega(X)$  and  $\omega(Y)$  are constant functions. Thus, the left invariant form  $\omega$  is closed if and only if  $\omega([X, Y]) = 0$  for all  $X, Y \in \mathfrak{g}$ . In particular, this implies that all left-invariant 1-forms are closed on abelian Lie groups such as  $\mathbb{R}^N$  and  $T^N$ , since  $[X, Y] = 0$  for all  $X, Y \in \mathfrak{g}$ . However, there are no closed left invariant 1-forms on  $SO(3)$  since a nontrivial  $\omega \in \mathfrak{so}(3)^*$  must be nonzero for at least some  $Z \in \mathfrak{so}(3)$ . However, for all  $Z \in \mathfrak{so}(3)$ , there exist  $X, Y \in \mathfrak{so}(3)$  such that  $Z = [X, Y]$ . In fact, this argument extends to all semisimple Lie groups, and thus there are no closed left-invariant 1-forms on any semisimple Lie group.

### 3.7 Discretization of the path signature

We have focused our discussion of the path signature so far on the continuous setting in order to discuss the theoretical framework. However, applications are studied in the discrete setting. In this section, we provide the explicit computation of path signatures for discrete time series, and discuss a useful discrete approximation.

Let  $v = (v^1, \dots, v^N) \in \mathfrak{g}$ , where we have written out the components of  $v$  in terms of some choice of basis on  $\mathfrak{g}$ . Consider the continuous time series  $\gamma_t = \exp(vt)$ , where  $\exp$  is the Lie exponential. Note that the derivative  $\gamma'_t = v$  is constant. In this case, the path

signature of  $\gamma$  is straightforward to compute. Given  $I = (i_1, \dots, i_m)$ , the path signature is

$$\begin{aligned} S^I(\gamma) &= \int_{\Delta^m} \omega_{i_1}(\gamma'_{t_1}) \dots \omega_{i_m}(\gamma'_{t_m}) dt_1 \dots dt_m \\ &= \int_{\Delta^m} v_{i_1} \dots v_{i_m} dt_1 \dots dt_m \\ &= \frac{v_{i_1} \dots v_{i_m}}{m!}. \end{aligned}$$

The entire path signature can be written concisely as the tensor exponential.

**Definition 41** *Let  $V$  be a real vector space. The tensor exponential  $\exp_{\otimes} : V \rightarrow T((V))$  is defined to be*

$$(\exp_{\otimes}(v))_m = \frac{v^{\otimes m}}{m!}.$$

Then, we may write the path signature of  $\gamma = \exp(vt)$  to be

$$S(\gamma) = \exp_{\otimes}(v).$$

Now, suppose we have a discrete time series  $\hat{\gamma} : [T + 1] \rightarrow G$ . Recall from Equation 4 that we may compute the discrete derivative by

$$\hat{\gamma}'_t = \log(\hat{\gamma}_t^{-1} \hat{\gamma}_{t+1})$$

to obtain a discrete time series  $\hat{\gamma}' : [T] \rightarrow \mathfrak{g}$ . As discussed in Section 2.3, we interpret these discrete time series as continuous paths by interpolating using the exponential. Thus, the continuous interpretation of the discrete time series can be thought of as a concatenation of several exponential paths

$$\gamma = \exp(\hat{\gamma}'_1 t) * \exp(\hat{\gamma}'_2 t) * \dots * \exp(\hat{\gamma}'_T t).$$

Therefore, by the above computation of the path signature of an exponential path and Chen's identity, we define the continuous path signature of the discrete time series to be

$$S(\hat{\gamma}) = \exp_{\otimes}(\hat{\gamma}'_1) \otimes \exp_{\otimes}(\hat{\gamma}'_2) \otimes \dots \otimes \exp_{\otimes}(\hat{\gamma}'_T). \quad (16)$$

By using tensor operations, this formula provides an effective implementation for the computation of the path signature.

An alternative approach is to compute an approximation of the path signature for discrete time series.

**Definition 42** *Let  $\hat{\gamma} : [T + 1] \rightarrow G$ . Suppose  $\omega_1, \dots, \omega_N \in \mathfrak{g}^*$  form a basis of  $\mathfrak{g}^*$ . Let  $I = (i_1, \dots, i_m)$  be a multi-index, where  $i_j \in [N]$ . Let  $\hat{\gamma}' : [T] \rightarrow \mathfrak{g}$  be the discrete derivative of  $\hat{\gamma}$ . We define the discrete  $m$ -simplex with length  $T$  to be*

$$\hat{\Delta}_T^m = \{(t_1, \dots, t_m) \in [T]^m \quad : \quad 0 \leq t_1 < t_2 < \dots < t_m \leq 1\}.$$

The discrete path signature of  $\hat{\gamma}$  with respect to  $I$  is defined to be

$$\hat{S}^I(\hat{\gamma}) := \sum_{(t_1, \dots, t_m) \in \hat{\Delta}_T^m} \omega_{i_1}(\hat{\gamma}'_{t_1}) \dots \omega_{i_m}(\hat{\gamma}'_{t_m}),$$

where  $t_1, \dots, t_m \in [T]$ . The discrete path signature can be viewed as a map

$$\hat{S} : \hat{P}G \rightarrow T((\bar{\mathfrak{g}})).$$

The discrete path signature can be viewed as an approximation to the continuous path signature. Let  $\gamma \in PG$  be a continuous path. Given a partition  $\pi = (0 = t_1 < t_2 < \dots < t_{T+1} = 1)$ , the discretization of  $\gamma$  with respect to  $\pi$ , denoted  $\hat{\gamma}^{(\pi)} : [T+1] \rightarrow G$ , is defined to be

$$\hat{\gamma}_i^{(\pi)} := \gamma_{t_i}.$$

The following proposition in Kiraly and Oberhauser (2019) shows that the discrete signature indeed approximates the continuous path signature.

**Proposition 43 (Corollary 4.3, Kiraly and Oberhauser (2019))** *Let  $\gamma \in P\mathbb{R}^N$ , and define a partition  $\pi = (0 = t_1 < t_2 < \dots < t_{T+1} = 1)$ . Then,*

$$\|\hat{S}(\hat{\gamma}^{(\pi)}) - S(\gamma)\| \leq |\gamma|_{1-var} e^{|\gamma|_{1-var}} \max_{i=1, \dots, T} |\gamma|_{[t_i, t_{i+1}]} |_{1-var}.$$

By applying the map  $\Phi : P\mathbb{R}^N \rightarrow PG$  from Proposition 34, and using the fact that it is an isometry, we immediately get the following corollary for Lie group valued paths.

**Corollary 44** *Let  $\gamma \in PG$ , and define a partition  $\pi = (0 = t_1 < t_2 < \dots < t_{T+1} = 1)$ . Then,*

$$\|\hat{S}(\hat{\gamma}^{(\pi)}) - S(\gamma)\| \leq |\gamma|_{1-var} e^{|\gamma|_{1-var}} \max_{i=1, \dots, T} |\gamma|_{[t_i, t_{i+1}]} |_{1-var}.$$

We consider the discrete path signature since it is more amenable to computation. In particular, Kiraly and Oberhauser (2019) derived efficient algorithms to compute the discrete path signature kernel, and we extend these algorithms to Lie groups in Section 4.3. In addition, the discrete path signature is of independent interest and the algebraic properties of the discrete path signature are studied in Diehl et al. (2020), where it is called the iterated sums signature.

### 3.8 Path transformations

We turn to discussion of several transformations which add information to paths.

### 3.8.1 TIME TRANSFORMATION

The *time transformation* is a simple method to remove reparametrization invariance (common in the path signatures literature such as in Chevyrev and Kormilitzin (2016)), defined by appending the time parameter to the path:

$$\begin{aligned} T_{\text{Time}} : PG_e &\rightarrow \tilde{P}(G \times \mathbb{R}) \\ \gamma_t &\mapsto (\gamma_t, t). \end{aligned} \tag{17}$$

**Lemma 45** *The map  $T_{\text{Time}} : PG_e \rightarrow \tilde{P}(G \times \mathbb{R})$  is injective.*

**Proof** Consider  $\alpha, \beta \in PG_e$ . The time parameter of  $T_{\text{Time}}(\alpha)$  is monotone increasing, implying the path  $T_{\text{Time}}(\alpha)$  is irreducible. Thus, since we also assume that  $\alpha_0 = \beta_0 = e$ , the paths  $T_{\text{Time}}(\alpha)$  and  $T_{\text{Time}}(\beta)$  are tree-like equivalent if and only if they differ by a reparametrization. However, all paths in the image of  $T_{\text{Time}}$  have the same parametrization in the time coordinate. Thus,  $T_{\text{Time}}(\alpha)$  is tree-like equivalent to  $T_{\text{Time}}(\beta)$  if and only if  $\alpha = \beta$ . ■

Because the path signature is injective with respect to tree-like equivalence classes of paths, this lemma implies that by first embedding a path  $\gamma \in PG_e$  into  $\tilde{P}(G \times \mathbb{R})$ , we obtain a parametrization-dependent feature map for  $PG_e$  using the path signature. More generally, this also removes tree-like invariance, so in fact we obtain an injective feature map for  $PG_e$ .

### 3.8.2 IDENTITY INITIALIZED TRANSFORMATION

The *identity initialized (IdInit) transformation* is a simple method to remove translation invariance: to the authors' knowledge, this has not been explicitly discussed in the context of path signatures. Suppose  $g \in G$  and we define  $\ell_g \in PG$  to be the exponential path from the identity  $e \in G$  to the point  $g \in G$ . Then the IdInit transformation is defined to be

$$\begin{aligned} T_{\text{IdInit}} : PG &\rightarrow PG_e \\ \gamma &\mapsto \ell_{\gamma_0} * \gamma. \end{aligned} \tag{18}$$

For discrete time series, this simply amounts to appending the identity element to the beginning of the time series. The following lemma is clear by definition.

**Lemma 46** *The map  $T_{\text{IdInit}} : PG \rightarrow PG_e$  is injective.*

The implication of this lemma is that by embedding  $PG$  into  $PG_e$ , we can obtain a translation-dependent feature map for  $PG$  using the path signature. Moreover, by combining the time and identity start transformations, the path signature provides an injective feature map for  $PG$

**Corollary 47** *The composition  $T_{\text{Time}} \circ T_{\text{IdInit}} : PG \rightarrow \tilde{P}(G \times \mathbb{R})$  is injective.*



### 3.8.3 SLIDING WINDOW TRANSFORMATION

The sliding window transformation (often called the lead-lag transformation in the path signature literature) has been used for several applications of the path signatures, such as Gyurkó et al. (2013) and Yang et al. (2019), and has been found to produce good results. An empirical study by Fermanian (2019) has shown that the sliding window transformation performs well on classification tasks. The exact definition may vary slightly between these papers, and we give the definition from Yang et al. (2019) and Fermanian (2019).

Let  $\tau \in \mathbb{R}^+$ , and  $m \in \mathbb{N}$ . Given a path  $\gamma \in PG$ , which we recall is defined on the unit interval  $\gamma : [0, 1] \rightarrow G$ , we can extend its definition to all of  $\mathbb{R}$  as

$$\gamma_t = \begin{cases} \gamma_0 & : t < 0 \\ \gamma_t & : t \in [0, 1] \\ \gamma_1 & : t > 1. \end{cases}$$

Then, we define the sliding window transformation with  $m$  lags to be

$$\begin{aligned} T_{\text{SWin},m} : PG &\rightarrow PG^{m+1} \\ \gamma_t &\mapsto (\gamma_t, \gamma_{t-\tau}, \gamma_{t-2\tau}, \dots, \gamma_{t-m\tau}). \end{aligned} \tag{19}$$

For discrete time series, we assume that the data is temporally uniformly sampled, and we choose  $\tau$  to be the time in between samples. Then, if we consider a discrete time series to be  $\tilde{\gamma} : \{0, \dots, n\} \rightarrow G$ , then the sliding window transformation will be

$$T_{\text{SWin},m}(\tilde{\gamma})_i := (\tilde{\gamma}_i, \tilde{\gamma}_{i-1}, \dots, \tilde{\gamma}_{i-m}).$$

In the context of Euclidean path signatures, Fermanian (2019) empirically shows that the sliding window embedding often performs well on classification tasks, though there is no theoretical explanation. We note that due to the choice of padding the start of the delayed time series with the identity, this transformation breaks translation invariance. We suggest that breaking the translation symmetry is one reason the sliding window transformation performs well in practice. This is discussed in Remark 64 in Section 5.1.

## 4. The path signature kernel for Lie groups

In this section, we show that a normalized variant of the path signature can be used to define a universal and characteristic kernel for Lie group valued time series. We begin by giving an overview of kernel methods and the types of problems that kernel methods can solve. Next, we extend the results of Chevyrev and Oberhauser (2018) to show that path signatures for Lie group valued paths are universal and characteristic. Finally, we show that the algorithms introduced in Kiraly and Oberhauser (2019) for efficient computation of the signature kernel can also be used for Lie group valued paths.

### 4.1 Background on universal and characteristic kernels

Suppose  $\mathcal{X}$  is a topological space which represents the space of data we would like to consider; we will call this the *input space*. Many tasks in machine learning can be separated into two broad classes of problems.

1. Those which involve making inferences about functions  $f \in \mathcal{F}$ , where  $\mathcal{F} \subset \mathbb{R}^{\mathcal{X}}$  is the function class which we are considering. Performing binary classification reduces to learning a function  $f \in \mathcal{F}$  such that the level set  $f = 0$  represents the *decision boundary* between the two classes of data.
2. Those which involve making inferences about probability measures  $\mu \in \mathcal{P}(\mathcal{X})$ , where  $\mathcal{P}(\mathcal{X})$  denotes the space of Borel probability measures on  $\mathcal{X}$ . For example, in two sample hypothesis testing, we begin with samples  $\{x_1, \dots, x_n\}$  and  $\{y_1, \dots, y_m\}$  taken from probability distributions  $p$  and  $q$  on  $\mathcal{X}$  respectively. Testing the null hypothesis that  $p = q$  then corresponds to learning about the underlying measures of  $p$  and  $q$ .

The general philosophy behind kernel methods is to map the input space  $\mathcal{X}$  into a reproducing kernel Hilbert space (RKHS)  $\mathcal{H}$  using a *feature map*

$$\Phi : \mathcal{X} \rightarrow \mathcal{H}_\kappa,$$

where the corresponding kernel is

$$\kappa : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}, \quad \kappa(x, x') = \langle \Phi(x), \Phi(x') \rangle.$$

Problems involving learning nonlinear functions  $f \in \mathbb{R}^{\mathcal{X}}$  given some input data  $\{x_i\}$ , where  $x_i \in \mathcal{X}$ , can be reformulated as problems involving learning an element  $g \in \mathcal{H}$  (which can be thought of as a function  $g \in \mathbb{R}^{\mathcal{X}}$ ) given the data  $\{\Phi(x_i)\}$ . Additionally, the norm induced by the Hilbert space provides a metric between points  $x, y \in \mathcal{X}$  as  $\|x - y\|$ . In essence, this translates a *nonlinear* learning problem into a *linear* learning problem. This allows the application of linear methods, which are much simpler and better developed in many cases.

Measures  $\mu$  on  $\mathcal{X}$  can be mapped into the RKHS via the *kernel mean embedding* (KME),

$$\bar{\Phi} : \mathcal{M}(\mathcal{X}) \rightarrow \mathcal{H}_\kappa, \quad \bar{\Phi}(\mu) := \int_{\mathcal{X}} \Phi(x) d\mu(x) = \mathbb{E}_\mu[\Phi]. \quad (20)$$

A priori, this map is not necessarily well-defined, so we will usually require restrictions on the feature map or kernel such that the integral exists.

**Lemma 48 (Sriperumbudur et al. (2010))** *Let  $\mu \in \mathcal{M}(\mathcal{X})$ . If the kernel  $\kappa : \mathcal{X} \rightarrow \mathbb{R}$  is measurable and the integral  $\int_{\mathcal{X}} \sqrt{\kappa(x, x)} d\mu(x) < \infty$ , then  $\bar{\Phi}(\mu) \in \mathcal{H}_\kappa$ .*

Specifically, the bounded integral condition is satisfied if we know that  $\kappa(x, x) = \|\Phi(x)\|^2 < C$  for all  $x \in \mathcal{X}$  for some fixed constant  $C$ . In other words, if the image of the feature map  $\Phi$  is contained in a bounded subset of  $\mathcal{H}_\kappa$ , then the KME is well defined.

Similar to the previous case, we can use the norm on  $\mathcal{H}_\kappa$  to define a notion of distance on  $\mathcal{P}(\mathcal{X})$ . Although this only provides a pseudometric since  $\bar{\Phi}$  may not be injective, it coincides with a well known measure of discrepancy between probability measures.

**Definition 49** *Let  $\mathcal{F} \subset \mathbb{R}^{\mathcal{X}}$  be a class of functions and  $\mu, \nu \in \mathcal{P}(\mathcal{X})$ . The maximum mean discrepancy (MMD) of  $\mu$  and  $\nu$  with respect to  $\mathcal{F}$  is*

$$\text{MMD}[\mathcal{F}, \mu, \nu] := \sup_{f \in \mathcal{F}} (\mathbb{E}_\mu[f] - \mathbb{E}_\nu[f]). \quad (21)$$

When we take the function class  $\mathcal{F}$  to be the unit ball in the RKHS  $\mathcal{H}_\kappa$ , the MMD can be written as the distance between the mean embeddings, with respect to the norm on  $\mathcal{H}$ .

**Lemma 50 (Borgwardt et al. (2006))** *Suppose the KME map  $\bar{\Phi}$  is well-defined and suppose  $\mu, \nu \in \mathcal{M}(\mathcal{X})$ . Let  $\mathcal{F} = \{f \in \mathcal{H} \subset \mathbb{R}^{\mathcal{X}} : \|f\| \leq 1\}$ . Then,*

$$\text{MMD}^2[\mathcal{F}, \mu, \nu] = \|\bar{\Phi}(\mu) - \bar{\Phi}(\nu)\|^2. \quad (22)$$

This simplifies the study of probability measures by considering them as elements of a linear space, and also provides a straightforward method to compute an unbiased finite sample estimate of the MMD in terms of the kernel.

**Lemma 51 (Gretton et al. (2012))** *Suppose the KME map  $\bar{\Phi}$  is well-defined and suppose  $\mu, \nu \in \mathcal{M}(\mathcal{X})$ . Let  $\mathcal{F} = \{f \in \mathcal{H} \subset \mathbb{R}^{\mathcal{X}} : \|f\| \leq 1\}$ . Let  $X = (x_1, \dots, x_n)$  and  $Y = (y_1, \dots, y_m)$  be i.i.d. samples from  $\mu$  and  $\nu$  respectively. An unbiased estimate of  $\text{MMD}^2[\mathcal{F}, \mu, \nu]$  is given as the MMD of the empirical distributions of  $X$  and  $Y$ ,*

$$\begin{aligned} \text{MMD}_u^2[\mathcal{F}, X, Y] &= \frac{1}{n(n-1)} \sum_{i=1}^n \sum_{j \neq i}^n \kappa(x_i, x_j) + \frac{1}{m(m-1)} \sum_{i=1}^m \sum_{j \neq i}^m \kappa(y_i, y_j) \\ &\quad - \frac{2}{nm} \sum_{i=1}^n \sum_{j=1}^m \kappa(x_i, y_j). \end{aligned} \quad (23)$$

From this discussion, kernels provide a unified way to study both nonlinear functions and probability measures using the linear space  $\mathcal{H}$ . However, there are deficiencies in both scenarios.

1. In the case of nonlinear functions, we usually begin by choosing our function class  $\mathcal{F} \subset \mathbb{R}^{\mathcal{X}}$ . How do we know that any function  $f : \mathcal{X} \rightarrow \mathbb{R}$  can be represented arbitrarily closely by an element in  $\ell \in \mathcal{H}$  such that  $f(x) \approx \langle \ell, \Phi(x) \rangle$  for all  $x \in \mathcal{X}$ ?
2. In the case of probability measures, it is often crucial that the MMD is in fact a metric instead of just a pseudometric. How do we know that the feature map  $\Phi$  is rich enough to distinguish all probability measures  $\mu \in \mathcal{M}(\mathcal{X})$ ?

The answer is given by the definitions of universal and characteristic kernels. This will require us to extend the definition of the KME to Schwarz distributions rather than just measures. We provide a quick exposition of the definitions here, and refer the reader to a more thorough treatment in Simon-Gabriel and Schölkopf (2018). As usual, let  $\mathcal{F} \subset \mathbb{R}^{\mathcal{X}}$  be a function class, and let  $\mathcal{F}'$  denote its topological dual of all continuous linear functionals. The definition of the KME for distributions is analogous to the case of measures

$$\bar{\Phi} : \mathcal{F}' \rightarrow \mathcal{H}_\kappa, \quad \bar{\Phi}(D) := \int_{\mathcal{X}} \Phi(x) dD(x), \quad (24)$$

where the integral here is the *weak-* or *Pettis-* integral (Simon-Gabriel and Schölkopf, 2018). Similar to the case of measures, this map is a priori not well-defined. However, we have a simple criterion for the existence of these weak integrals.

**Lemma 52 (Simon-Gabriel and Schölkopf (2018))** *Let  $\mathcal{F} \subset \mathbb{R}^{\mathcal{X}}$ . If the map  $\iota : \mathcal{H}_\kappa \rightarrow \mathbb{R}^{\mathcal{X}}$  defined by  $\iota(\ell) := \langle \ell, \Phi(\cdot) \rangle$  has image contained in  $\mathcal{F}$ , then the integral  $\bar{\Phi}(D)$  exists for all  $D \in \mathcal{F}'$ . Thus, the KME map  $\bar{\Phi} : \mathcal{F}' \rightarrow \mathcal{H}_\kappa$  is well-defined.*

We can now state the definition of a universal and characteristic feature map.

**Definition 53** *Fix an input space  $\mathcal{X}$ , and a function space  $\mathcal{F} \subset \mathbb{R}^{\mathcal{X}}$ . Consider a feature map*

$$\Phi : \mathcal{X} \rightarrow \mathcal{H}_\kappa$$

*into an RKHS  $\mathcal{H}_\kappa$  with respect to a kernel  $k$ . Suppose that  $\langle \ell, \Phi(\cdot) \rangle \in \mathcal{F}$  for all  $\ell \in \mathcal{H}_\kappa$ . We say that  $\Phi$  is*

1. universal to  $\mathcal{F}$  if the map

$$\iota : \mathcal{H}_\kappa \rightarrow \mathcal{F}, \quad \ell \mapsto \langle \ell, \Phi(\cdot) \rangle$$

*has a dense image in  $\mathcal{F}$ ; and*

2. characteristic to a subset  $\mathcal{P} \subseteq \mathcal{F}'$  if the KME map

$$\bar{\Phi} : \mathcal{P} \rightarrow \mathcal{H}_\kappa, \quad D \mapsto \int_{\mathcal{X}} \Phi(x) dD(x)$$

*is injective.*

Note that we have assumed that the image  $\iota(\mathcal{H}_\kappa) \subset \mathcal{F}$  so by Lemma 52, the KME map is well defined. The property of universality allows us to approximate any function  $f \in \mathcal{F}$  using linear functionals  $\langle \ell, \Phi(\cdot) \rangle$  for  $\ell \in \mathcal{H}_\kappa$ . The dual of a class of functions  $\mathcal{F}$  is generally much larger than the set of probability measures on  $\mathcal{X}$ . If  $\mathcal{M}(\mathcal{X}) \subset \mathcal{F}'$ , then a characteristic feature map is able to represent probability measures on  $\mathcal{X}$  with elements of  $\mathcal{H}$ . Moreover the MMD becomes a metric due to the injectivity of the KME.

We have the following equivalence between universality and characteristicness, as shown in Simon-Gabriel and Schölkopf (2018) and Chevyrev and Oberhauser (2018).

**Theorem 54** *Suppose that  $\mathcal{F}$  is a locally convex topological vector space. A feature map  $\Phi$  is universal to  $\mathcal{F}$  if and only if  $\Phi$  is characteristic to  $\mathcal{F}'$ .*

## 4.2 The path signature kernel

In this subsection, we will define the path signature kernel, and show that it is universal and characteristic. This was shown for the Euclidean case in Chevyrev and Oberhauser (2018), which states that these properties hold when studying paths evolving in a Hilbert space. Through our definition of the path signature for Lie groups, we provide a clarification: the space itself need not be a Hilbert space, but rather the space of tangent vectors must be a Hilbert space. In our case of Lie groups, this is the Lie algebra with a Riemannian metric. With this setup, the path signature kernel for Lie groups is universal and characteristic.

Recall that we have defined  $T_1((\bar{\mathfrak{g}}))$  in Equation 13 to be the subspace of  $T((\bar{\mathfrak{g}}))$  with constant value 1 and finite norm. We will view the path signature

$$S : \widetilde{PG} \rightarrow T_1((\bar{\mathfrak{g}}))$$

as a feature map, and recall that  $T_1((\bar{\mathfrak{g}}))$  is equipped with an inner product, and is in particular a Hilbert space. However, as discussed in the previous subsection, we will need to ensure that the signature map sends paths to a bounded subset of  $T_1((\bar{\mathfrak{g}}))$  in order for the KME to be defined. This will be done by using a tensor normalization, which was first discussed in Chevyrev and Oberhauser (2018).

**Definition 55** *A tensor normalization is a continuous injective map of the form*

$$\begin{aligned} \Lambda : T_1((\bar{\mathfrak{g}})) &\rightarrow \{\mathbf{t} \in T_1((\bar{\mathfrak{g}})) : \|\mathbf{t}\| \leq K\} \\ \mathbf{t} &\mapsto \delta_{\lambda(\mathbf{t})}\mathbf{t} \end{aligned}$$

where  $K > 0$  is a constant,  $\lambda : T_1((\bar{\mathfrak{g}})) \rightarrow (0, \infty)$  is a function, and  $\delta_\lambda$  is the tensor dilation from Definition 9.

We will discuss the construction and computational aspects of tensor normalization in Appendix B. For now, we will assume that a tensor normalization exists, and set  $\Lambda : T_1((\bar{\mathfrak{g}})) \rightarrow T_1((\bar{\mathfrak{g}}))$  to be a fixed tensor normalization. Now, we have the *normalized signature*

$$\Phi_S : \widetilde{PG} \rightarrow T_1((\bar{\mathfrak{g}})), \quad \Phi_S = \Lambda \circ S,$$

which is a continuous injective map from  $\widetilde{PG}$  into a bounded subset of  $T_1((\bar{\mathfrak{g}}))$ . Note that due to the scaling property of the path signature from Proposition 23, this is equivalent to

$$\Phi_S(\gamma) = S(\lambda(S(\gamma)) \cdot \gamma),$$

where we first scale the path in  $G$  by  $\lambda(S(\gamma))$ , using the Lie algebra scaling from Definition 13.

Following Chevyrev and Oberhauser (2018), we will show universality and then use the duality in Theorem 54 to show characteristicness with respect to probability measures. Using the theory discussed in the previous section, the objective is to find a function class  $\mathcal{F} \subset \mathbb{R}^{PG}$  and a topology on  $PG$  such that

1. the function class  $\mathcal{F}$  can be approximated by linear functionals  $\langle \ell, \Phi_S(\cdot) \rangle$ , and
2. the dual  $\mathcal{F}'$  contains probability measures on  $PG$ .

The difficulty with such a result is due to the fact that  $PG$  is not locally compact. However, the class of continuous bounded functions  $C_b(PG, \mathbb{R})$  has such properties when  $PG$  is endowed with the *strict topology*, originally defined in Giles (1971).

**Definition 56** *Let  $X$  be a topological space. We say that a function  $\psi : X \rightarrow \mathbb{R}$  vanishes at infinity if for each  $\epsilon > 0$ , there exists a compact set  $K \subset X$  such that  $\sup_{x \in X-K} |\psi(x)| < \epsilon$ . Denote by  $B_0(X, \mathbb{R})$  the set of functions that vanish at infinity. The strict topology on  $C_b(X, \mathbb{R})$  is the topology generated by the seminorms*

$$p_\psi(f) = \sup_{x \in X} |f(x)\psi(x)|, \quad \psi \in B_0(X, \mathbb{R}).$$

**Theorem 57** *Let  $X$  be a metrizable topological space.*

1. *The strict topology on  $C_b(X, \mathbb{R})$  is weaker than the uniform topology and stronger than the topology of uniform convergence on compact sets.*
2. *If  $\mathcal{F}_0$  is a subalgebra of  $C_b(X, \mathbb{R})$  such that for all  $x, y \in X$ , there exists some  $f \in \mathcal{F}_0$  such that  $f(x) \neq f(y)$  ( $\mathcal{F}_0$  separates points), and for all  $x \in X$ , there exists some  $f \in \mathcal{F}_0$  such that  $f(x) \neq 0$ , then  $\mathcal{F}_0$  is dense in  $C_b(X, \mathbb{R})$  under the strict topology.*
3. *The topological dual of  $C_b(X, \mathbb{R})$  equipped with the strict topology is the space of finite regular Borel measures on  $X$ .*

Specifically, note that the space of finite regular Borel measures on  $X$  includes all probability measures on  $X$ . Finally, we are ready to state the universality and characteristicness result.

**Theorem 58** *Let  $\Lambda : \overline{T}_1(\mathbb{R}^n) \rightarrow \overline{T}_1(\mathbb{R}^n)$  be a tensor normalization. The normalized signature*

$$\Phi : \widetilde{PG} \rightarrow T_1((\overline{\mathfrak{g}})), \quad \Phi = \Lambda \circ S,$$

1. *is a continuous injection from  $\widetilde{PG}$  into a bounded subset of  $\overline{T}_1(\mathbb{R}^n)$ ,*
2. *is universal to  $\mathcal{F} := C_b(\widetilde{PG}, \mathbb{R})$  equipped with the strict topology, and*
3. *is characteristic to the space of finite regular Borel measures on  $\widetilde{PG}$ .*

**Proof** The fact that  $\Phi$  is an injection follows from the injectivity of the path signature from Theorem 30 and the definition of the tensor normalization. Continuity follows from the stability property of Corollary 38. Next, we move on to universality. Define

$$L = 1 + \bigoplus_{m=1}^{\infty} (\overline{\mathfrak{g}})^{\otimes m}$$

to be a dense subspace of  $T_1((\overline{\mathfrak{g}}))$  (note that  $L$  only contains finite linear combinations of tensors, whereas  $T_1((\overline{\mathfrak{g}}))$  contains all power series of tensors) and define

$$\mathcal{F}_0 = \{ \langle \ell, \Phi(\cdot) \rangle : \widetilde{PG} \rightarrow \mathbb{R} : \ell \in L \}.$$

We aim to show that  $\mathcal{F}_0$  satisfies the hypotheses of the second point of Theorem 57. By the injectivity of  $\Phi$ , the class of functions  $\mathcal{F}_0$  separates points, and because the path signature is defined with constant term 1, the path signature is nonzero for all paths  $\gamma \in \widetilde{PG}$ . Finally, by the shuffle product identity from Theorem 33, the class of functions  $\mathcal{F}_0$  is closed under shuffle multiplication and is therefore a subalgebra of  $C_b(\widetilde{PG}, \mathbb{R})$ . Namely, let  $I$  and  $J$  be multi-indices, and  $e_I$  and  $e_J$  be the corresponding basis vectors in  $L$ . Then, we may define multiplication in  $\mathcal{F}_0$  by

$$\langle e_I, \Phi(\cdot) \rangle \langle e_J, \Phi(\cdot) \rangle = \left\langle \sum_{K \in I \sqcup J} e_K, \Phi(\cdot) \right\rangle,$$

which is closed. Thus,  $\Phi$  is universal with respect to  $\mathcal{F}$ . Finally, by the duality in Theorem 54 and the third point of Theorem 57, the function class  $\mathcal{F}$  is characteristic with respect to finite regular Borel measures on  $\widetilde{PG}$ .  $\blacksquare$

**Remark 59** *Although this theorem is stated for tree-like equivalence classes of paths in  $G$ , by precomposing with the time transformation or identity start transformation discussed in Section 3.8, we can also obtain universal and characteristic feature maps that are not reparametrization or translation invariant.*

### 4.3 The kernel trick

The *kernel trick* refers to an efficient method to compute kernels without having to compute explicit representations of elements in the feature space. Several efficient algorithms to compute the Euclidean path signature kernel are provided in Kiraly and Oberhauser (2019), who state that their algorithms hold for path signatures computed for Hilbert space-valued data (their methods hold in more generality). The algorithms depend only on an inner product structure in the space where the integrals are being computed, namely in  $\bar{\mathfrak{g}}$ , and thus also hold in our present context of Lie group-valued data. In this section, we will provide an explicit generalization of their main algorithm.

As in Section 3.3, we will restrict ourselves to the truncated signature

$$S_M : PG \rightarrow T^{\leq M}(\bar{\mathfrak{g}}).$$

Here, the inner product for  $T^{\leq M}(\bar{\mathfrak{g}})$  was given in Equation 12. The *signature kernel truncated at level  $M$*  is defined to be

$$K_M : PG \times PG \rightarrow \mathbb{R}, \quad (\alpha, \beta) \mapsto \langle S_M(\alpha), S_M(\beta) \rangle \quad (25)$$

We will begin by simplifying the computation of the kernel for continuous paths.

**Proposition 60** *The signature kernel can be computed as*

$$K_M(\alpha, \beta) = \sum_{m=0}^M \int_{(s,t) \in \Delta^m \times \Delta^m} \prod_{i=1}^m \langle \alpha'_{s_i}, \beta'_{t_i} \rangle_{\mathfrak{g}} ds dt, \quad (26)$$

where we view  $\alpha'_t$  and  $\beta'_t$  as elements of the Lie algebra  $\mathfrak{g}$ , and the inner product in the integrand  $\langle \cdot, \cdot \rangle_{\mathfrak{g}}$  is computed in the Lie algebra. Also, we denote  $s = (s_1, \dots, s_m)$  and  $t = (t_1, \dots, t_m)$  as elements of  $\Delta^m$ , and write  $ds := ds_1 \dots ds_m$  and  $dt := dt_1 \dots dt_m$ .

**Proof** Let's consider the inner product at a single level  $m$ . Recall that  $\pi_m : T((\bar{\mathfrak{g}})) \rightarrow \bar{\mathfrak{g}}^{\otimes m}$  is the projection on to the level  $m$  tensors, and  $\langle \cdot, \cdot \rangle_m$  refers to the inner product on  $\bar{\mathfrak{g}}^{\otimes m}$ .

Then,

$$\begin{aligned}
 & \langle \pi_m S(\alpha), \pi_m S(\beta) \rangle_m \\
 &= \sum_{I:|I|=m} S^I(\alpha) \cdot S^I(\beta) \\
 &= \sum_{I=(i_1, \dots, i_m)} \int_{\Delta^m} \omega_{i_1}(\alpha'_{s_1}) \dots \omega_{i_m}(\alpha'_{s_m}) ds \int_{\Delta^m} \omega_{i_1}(\beta'_{t_1}) \dots \omega_{i_m}(\beta'_{t_m}) dt \\
 &= \int_{(s,t) \in \Delta^m \times \Delta^m} \sum_{I=(i_1, \dots, i_m)} [\omega_{i_1}(\alpha'_{s_1}) \omega_{i_1}(\beta'_{t_1})] \dots [\omega_{i_m}(\alpha'_{s_m}) \omega_{i_m}(\beta'_{t_m})] ds dt \\
 &= \int_{(s,t) \in \Delta^m \times \Delta^m} \prod_{i=1}^m \langle \alpha'_{s_i}, \beta'_{t_i} \rangle_{\mathfrak{g}} ds dt.
 \end{aligned}$$

Then, adding up all of the levels, we get our desired result. ■

As noted by Kiraly and Oberhauser (2019), the expression in Equation 26 can be efficiently computed by a method that is similar to Horner's scheme for computing polynomial expressions. Suppose we wish to compute the expression  $p(x) = \sum_{i=0}^M x^i$ . By expanding this polynomial as

$$p(x) = 1 + x(1 + x(1 + \dots + x(1 + x)))$$

where the recursion occurs  $M$  times and computing the brackets from the inside to the outside, we can evaluate the expression using  $M$  additions and  $M$  multiplications. In contrast, the naive computation of  $p(x)$  would require  $M$  additions and  $\frac{M^2+M}{2}$  multiplications. Note that we may write out this recursion explicitly as follows. Let

$$q_1 = 1 + x, \quad q_m = 1 + xq_{m-1}.$$

Then, we may write  $p(x) = q_M$ . We can significantly reduce the number of operations required to compute the integrals in Equation 26 by adapting this procedure.

**Corollary 61** *Let*

$$Q_1(s, t) = 1 + \int_{\substack{s' \in [0, s] \\ t' \in [0, t]}} \langle \alpha'_{s'}, \beta'_{t'} \rangle_{\mathfrak{g}} ds' dt'. \quad (27)$$

*Then, we recursively define*

$$Q_m(s, t) = 1 + \int_{\substack{s' \in [0, s] \\ t' \in [0, t]}} Q_{m-1}(s', t') \langle \alpha'_{s'}, \beta'_{t'} \rangle_{\mathfrak{g}} ds' dt'. \quad (28)$$

*The signature kernel can be computed as*

$$K_M(\alpha, \beta) = Q_M(1, 1). \quad (29)$$



**Proof** We will consider the case of  $M = 2$ . Here, we have

$$\begin{aligned}
 Q_2(1, 1) &= 1 + \int_{\substack{s_1 \in [0,1] \\ t_1 \in [0,1]}} Q_1(s_1, t_1) \langle \alpha'_{s_1}, \beta'_{t_1} \rangle_{\mathfrak{g}} ds_1 dt_1 \\
 &= 1 + \int_{\substack{s_1 \in [0,1] \\ t_1 \in [0,1]}} \left( 1 + \int_{\substack{s_2 \in [0, s_1] \\ t_2 \in [0, t_1]}} \langle \alpha'_{s_2}, \beta'_{t_2} \rangle_{\mathfrak{g}} ds_2 dt_2 \right) \langle \alpha'_{s_1}, \beta'_{t_1} \rangle_{\mathfrak{g}} ds_1 dt_1 \\
 &= 1 + \int_{\substack{s_1 \in [0,1] \\ t_1 \in [0,1]}} \langle \alpha'_{s_1}, \beta'_{t_1} \rangle_{\mathfrak{g}} ds_1 dt_1 + \int_{\substack{s_1 \in [0,1] \\ t_1 \in [0,1]}} \int_{\substack{s_2 \in [0, s_1] \\ t_2 \in [0, t_1]}} \prod_{i=1}^2 \langle \alpha'_{s_i}, \beta'_{t_i} \rangle_{\mathfrak{g}} ds_1 ds_2 dt_1 dt_2 \\
 &= K_2(\alpha, \beta).
 \end{aligned}$$

The general proof proceeds in the same manner. By each successive unfolding of the definition of  $Q_m$ , we recover an additional summand in Equation 26.  $\blacksquare$

Next, we will consider the discrete formulation of this expression. For simplicity, we will consider discrete time series of the same length, though the following results also hold when the two discrete time series are of different lengths. Suppose  $\hat{\alpha}, \hat{\beta} : [T + 1] \rightarrow G$ , and let  $\hat{\alpha}', \hat{\beta}' : [T] \rightarrow \mathfrak{g}$  be their corresponding discrete derivatives. Following the notation in Section 3.7 for the discrete signature, we define the *discrete signature kernel truncated at level  $M$*  to be

$$\hat{K}_M(\hat{\alpha}, \hat{\beta}) := \langle \hat{S}_M(\hat{\alpha}), \hat{S}_M(\hat{\beta}) \rangle_M. \quad (30)$$

Notice we are using the discrete signature given in Definition 42. Then, the discrete analogues of Proposition 60 and Corollary 61 are as follows.

**Proposition 62** *The discrete signature kernel can be computed as*

$$\hat{K}_M(\hat{\alpha}, \hat{\beta}) = \sum_{m=0}^M \sum_{(s,t) \in \hat{\Delta}_T^m \times \hat{\Delta}_T^m} \prod_{i=1}^m \langle \alpha'_{s_i}, \beta'_{t_i} \rangle_{\mathfrak{g}}. \quad (31)$$

**Corollary 63** *Let*

$$\hat{Q}_1(s, t) = 1 + \sum_{s' \in [s], t' \in [t]} \langle \alpha'_{s'}, \beta'_{t'} \rangle_{\mathfrak{g}} \quad (32)$$

*and recursively define*

$$\hat{Q}_m(s, t) = 1 + \sum_{s' \in [s], t' \in [t]} \hat{Q}_{m-1}(s', t') \langle \alpha'_{s'}, \beta'_{t'} \rangle_{\mathfrak{g}}. \quad (33)$$

*The discrete signature kernel can be computed as*

$$\hat{K}_M(\hat{\alpha}, \hat{\beta}) = \hat{Q}_M(T, T). \quad (34)$$

The proofs for these discrete formulas proceed in exactly the same manner as their continuous counterparts. This final recursive formula for the kernel provides an efficient computation of the discrete signature. We set some notation before writing down the algorithm. Suppose  $A, B$  are  $T \times T$  arrays. We will use the notation  $A[i, j]$  to denote elements of the array, and we suppose that the arrays are 1-indexed. The notation for the pseudocode is explained in Appendix A.

For the algorithm, suppose we have discrete time series  $\alpha, \beta : [T + 1] \rightarrow G$ , and the corresponding derivatives  $\alpha', \beta' : [T] \rightarrow \mathfrak{g}$  are already computed, as per Appendix A. We assume that the Lie group  $G$  is  $N$  dimensional. In the pseudocode, we let  $\mathbf{a}, \mathbf{b}$  be the discrete derivatives  $\alpha', \beta'$  respectively.

---

**Algorithm 1:** Discretized Signature Kernel

---

**Input** :  $\mathbf{a}, \mathbf{b}$  ( $T \times N$  arrays): Two paths as discrete derivatives  
 $M$  (Int): The truncation level  
**Output:**  $R$  (Float): The kernel value  $K_M(\alpha, \beta)$ .

- 1 Compute the Gram matrix of the derivatives  
 $K \leftarrow \mathbf{a}\mathbf{b}^\top$ ;
- 2 Initialize ( $T \times T$ ) arrays  $A$  and  $Q$  ;
- 3 Initialize the first step of the recursion  
 $A \leftarrow K$  ;
- 4 **for**  $m=2..M$  **do**
- 5      $Q \leftarrow 1 + A[\boxplus, \boxplus]$  ;
- 6      $A \leftarrow K \cdot Q$  ;
- 7 **end**
- 8  $R \leftarrow 1 + A[\Sigma, \Sigma]$  ;
- 9 Return  $R$

---

As discussed by Kiraly and Oberhauser (2019), the runtime of this algorithm is  $O(T^2 \cdot M)$ . Now, consider the naive computation of the kernel where we first compute the truncated signatures of  $\hat{\alpha}$  and  $\hat{\beta}$ , and then compute the inner product. From the analysis in Appendix A, a signature computation requires  $O(TN^M)$  operations. The leading-order term in the computation of the inner product is computing the inner product of the respective level  $M$  tensors. This requires  $N^M$  operations. Thus, the complexity of the naive computation is  $O(TN^M)$ .

This suggests that if the length  $T$  of the time series is large and the truncation level is small, then a naive computation may be more efficient. However, if we wish to compute the kernel at a high truncation level, the recursive algorithm provided here scales significantly better. Furthermore, several variants of this algorithm for Euclidean path signatures are considered in Kiraly and Oberhauser (2019) such as by incorporating low-rank approximations. These algorithms can similarly be extended to the setting of Lie groups by applying them to the discrete derivatives.

## 5. Experiments

In this section, we provide two detailed experiments to demonstrate the universal and characteristic properties of the path signature. First, we consider the human action recognition

problem from computer vision, using a Lie group representation of the data. We find that the path signature method is simple to implement, achieves comparable classification performance to shallow learning methods, and provides an interpretable feature set. Second, we perform a kernel two-sample hypothesis test aiming to distinguish between two different random walks on  $SO(3)$ . Here, we find that the path signature for  $SO(3)$  significantly outperforms the same hypothesis test done using the Euclidean representation of  $SO(3)$ .

## 5.1 Human action recognition

In this subsection, we aim to utilize the universal property of the signature kernel to study a classification problem. Namely, we will be working in the domain of *human action recognition*, in which several recent works, such as Huang et al. (2017), Rhif et al. (2018) and Li et al. (2019), have used Lie group representations and deep learning to achieve state of the art results. We also note that Euclidean path signatures have also been used to study the Euclidean valued and joint-based representation of this problem in Yang et al. (2019).

Our goal in this section is to demonstrate the utility, simplicity and interpretability of the Lie group path signature. Thus, our focus will be on establishing baseline results using support vector machines and random forests, and our main comparison will be with the “shallow” learning methods described in Vemulapalli et al. (2014) and Vemulapalli and Chellappa (2016). We achieve comparable classification results using a random forest algorithm and the second level signature as a feature set. A major advantage of the second level signature is that we can interpret this feature set in terms of the underlying movement, and we provide a brief discussion of this analysis.

### 5.1.1 LIE GROUP REPRESENTATION

In this experiment, we aim to classify actions based on human skeletal motion data, which we represent as a  $SO(3)^k$  valued time series. This Lie group representation of human motion was introduced in Vemulapalli et al. (2014) and Vemulapalli and Chellappa (2016), which describes a pose as an element of  $SO(3)^k$ . The relative rotation of two body parts is an element of  $SO(3)$ , and by recording the relative rotation of  $k$  pairs of body parts, we can describe the full pose of a human.

A skeleton  $S = (V, E)$  can be described using a set of joints  $V = \{v^i\}_{i=1}^{n_V}$ , where  $v^i \in \mathbb{R}^3$ , and a set of body parts  $E = \{e^i\}_{i=1}^{n_E}$ , where each body part is a function  $e^i : \{0, 1\} \rightarrow V$ . We may consider  $e^i(0)$  and  $e^i(1)$  to respectively be the start and end point of the body part  $e^i$ . Additionally, we will denote the unit vector describing the direction from the start to the end of the body part by

$$\hat{e}^i = \frac{e^i(1) - e^i(0)}{\|e^i(1) - e^i(0)\|}.$$

We will consider all pairs of body parts  $(e^i, e^j)$  that share a joint, such that  $e^i(c_1) = e^j(c_2)$  for some  $c_1, c_2 \in \{0, 1\}$ , so  $k$  is the number of adjacent pairs of body parts. To obtain the rotation matrix for a chosen pair  $(e^i, e^j)$ , we rotate the global coordinate system (with minimum rotation) such that  $\hat{e}^i$  is the  $x$ -axis. Then, the rotation matrix  $R_{i,j} \in SO(3)$  is the minimum rotation from  $\hat{e}^i$  to  $\hat{e}^j$  in this coordinate system. By repeating this for all adjacent

pairs, we obtain an element of  $SO(3)^k$ , and further repeating this for all time steps, we can represent this motion as a time series in  $SO(3)^k$ .

### 5.1.2 DATA AND PREPROCESSING

The dataset we use is the G3D-Gaming dataset from Bloom et al. (2012). This contains 663 sequences of 20 different gaming motions performed by 10 subjects, and each subject performed each action at least two times. The 3D locations of 20 joints are provided for every frame. However, the number of frames for each recorded action varies widely from 3 frames to 330 frames. We use the code provided in Vemulapalli and Chellappa (2016) to generate the  $SO(3)$  representation.




Figure 4: Histogram of number of frames per recorded action in G3D dataset.

Note that in Vemulapalli et al. (2014) and Vemulapalli and Chellappa (2016), all pairs of body parts are used in the Lie group representation, which results in  $k = 342$  pairs for this dataset. In contrast, we only use all adjacent pairs of body parts, which we call the *primary pairs*, resulting in  $k = 18$  pairs for this dataset. We use significantly less data because the path signatures take into account the relationships between all input pairs, so we have information regarding non-primary pairs through the higher order signature terms. The numbering of the primary pairs are given in the following figure.




Figure 5: Numbering of the primary pairs of body parts.

We categorize these 18 primary pairs into 5 classes.

Body	1-2
L Arm	3-6
R Arm	7-10
L Leg	11-14
R Leg	15-18

Table 1: The body regions corresponding to groups of primary pairs.

Extensive preprocessing of this data is performed in Vemulapalli et al. (2014) and Vemulapalli and Chellappa (2016) in order to deal with several difficulties. During the training stage, the following steps were taken.

1. Each time series is resampled via interpolation so that all time series have a fixed length.
2. A nominal curve in  $SO(3)^k$  is generated for each action class.
3. To handle issues of rate variation and temporal misalignment, dynamic time warping (DTW) is used to warp each time series to its corresponding nominal curve.
4. A rolling and unwrapping procedure is computed with respect to its nominal curve to obtain a curve in the Lie algebra  $\mathfrak{so}(3)^k$ .
5. (Optional) A Fourier temporal pyramid (FTP) representation of the Lie algebra curve may be computed to further deal with temporal misalignment.

A classifier such as a support vector machine (SVM) is then trained for every action class (one vs. rest). In the testing stage, the data is preprocessed with respect to all nominal curves and the corresponding SVM is used for prediction. This amounts to a large pre-processing cost, especially for test samples, which must be preprocessed with respect to all action classes.

In contrast, we perform minimal preprocessing since we can compare time series with varying numbers of frames using path signatures, and the issues of rate variation and temporal misalignment are handled by reparametrization invariance.

### 5.1.3 RESULTS

Following Vemulapalli et al. (2014); Vemulapalli and Chellappa (2016), we use a cross-subject test setting, where we use half of the subjects for training, and the other half for testing. All of the reported classification results are averaged over ten different combinations of the train/test split. We perform the classification using a kernel SVM, as well as a random forest. For the kernel SVM, we report the results using the signature kernel truncated at level 6. We use the Julia implementation of `scikit-learn`, with the `SVC` implementation for support vector machines, which uses the one-against-one approach (Knerr et al., 1990) for multi-class classification. For the random forest, we compute the level 2 signature of the time series and treat it as a feature set. We use the random forest implementation in the Julia `DecisionTrees` package. We follow the suggested default random forest hyperparameters in Probst et al. (2019), and use 1,000 trees,  $\sqrt{n_f}$  features for each tree, where  $n_f$  is the total number of input features, a maximum depth of 100, and use 70% of the data to train each tree. A tensor renormalization is used for all path signature computations as described in Proposition 65, and using the function

$$\psi(x) = \begin{cases} x^2 & \text{if } x \leq \sqrt{M} \\ M + M^{1+a}(M^{-a} - x^{-a})/a & \text{if } x > \sqrt{M} \end{cases}$$

where  $M = 4$  and  $a = 1$  as in Chevyrev and Oberhauser (2018).

We apply each algorithm to several embeddings of the discrete time series. In addition to using the raw time series, we also use the identity initialized (IdInit) embedding to remove translation invariance, and the sliding window (SWin) embedding with 1 to 6 lags.

	Random Forest	SVM
Raw	69.99%	65.20%
IdInit	80.48%	77.86%
SW, 1 lag	80.32%	79.24%
SW, 2 lags	82.71%	80.27%
SW, 3 lags	83.77%	80.84%
SW, 4 lags	84.47%	81.08%
SW, 5 lags	84.86%	81.44%
SW, 6 lags	85.56%	81.44%

Table 2: Classification results using a random forest and SVM for different embeddings.

Based on our results, a random forest trained using level 2 signatures outperforms an SVM trained using level 6 signatures for all embeddings. This may be due to the fact that random forests are in general better suited for multi-class classification tasks, since the SVM approach is to split the multi-class problem into  $\binom{20}{2}$  binary classification tasks.

We note by using the random forest classifier, and introducing lags, we are able to achieve results which are comparable to the 87.95% accuracy of Vemulapalli and Chellappa (2016), given the fact that we are using significantly less input data, minimal preprocessing, and default hyperparameters on the random forest.

**Remark 64** *In Section 3.8.3, we mentioned that one possible explanation for the strong empirical performance of the sliding window embedding is the breaking of translation invariance. In these results, we can isolate the effect of breaking translation invariance by using the IdInit embedding. We see that for both the random forest and the SVM, the performance improves significantly after using the IdInit embedding when compared to the raw time series. The sliding window embedding with 1 lag also significantly improves the performance when compared to the raw time series, but we see that the increase is comparable to that of the IdInit embedding. When we increase the number of lags, the path signature is able to capture more information by integrating with respect to past values, and thus the performance continues to improve. Fermanian (2019) suggested the problem of explaining why the sliding window embedding performs well in practice. This empirical evidence indicates that one such reason is the breaking of the translation invariance of the path signature.*

#### 5.1.4 INTERPRETATION OF THE SECOND LEVEL SIGNATURE MATRIX

In addition to comparable classification accuracy, the second level signature matrix affords an interpretable feature set. This is one of the highlights of this method, as none of the cited classification methods are easily interpretable. Namely, in both the narrow and deep learning approaches, the input to the classification algorithm is the stacked (and in some cases preprocessed) Lie group valued time series, resulting in an extremely high dimensional vector. In this section, we will discuss the interpretation of the raw second level signature matrix, which was used for the random forest classification.

The second level path signature of a path  $\gamma \in PG$ , where  $G$  is  $N$ -dimensional, can be viewed as a  $(N \times N)$  matrix where the  $(i, j)$  entry is simply  $S^{(i,j)}(\gamma)$ . For these applications, our Lie group is  $G = SO(3)^{18}$ , which has  $N = 54$  dimensions. We will consider the absolute value of the averaged second level signature matrices for each action class of the G3D data set. Specifically, because we are focused only on the interpretation in this section, we will use all of the data when computing averages. Also, we take the absolute value to simplify the interpretation of these matrices. We begin by considering a single action class (walk) in order to describe how to read the matrix.

Here we have visualized the  $(54 \times 54)$  matrix as a heatmap, where the entries are arranged in accordance with standard matrix notation. The thin red grid lines separate the matrix into  $(3 \times 3)$  blocks, where each triplet of entries corresponds to a single primary pair. The numerical labels on the axes enumerate the blocks (or primary pairs), which correspond to the primary pairs as shown in Figure 5. The thicker red grid lines are used to separate the different categories of primary pairs, as outlined in Table 1, and labelled in the figure. The diagonal entries of the matrix measure the square of the cumulative rotation of a given

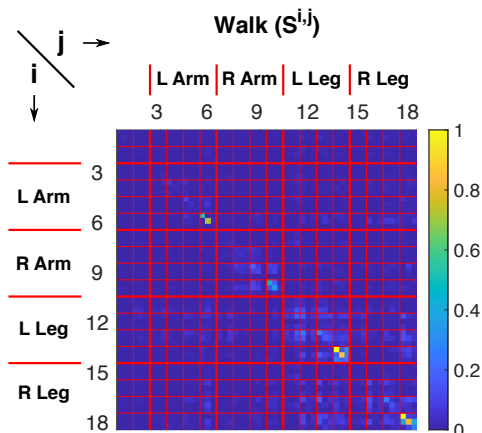


Figure 6: Averaged absolute second level signature matrix for the action class “walk.”

basis direction for a given primary pair. The off diagonal entries measure the positive or negative influence of the basis direction  $i$  on the basis direction  $j$ , as defined in Section 3.5. For a closer look at this example, we isolate the blocks corresponding to the left and right legs, and reduce the color threshold.

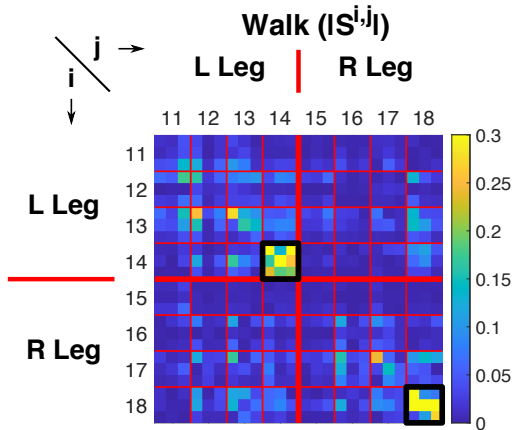


Figure 7: Averaged absolute second level signature matrix for the action class “walk.”

The blocks with the largest magnitude are the diagonal blocks 14 and 18, which correspond to the joints of the left and right foot (see Figure 5). In addition, many of the off-diagonal entries are nonzero, which corresponds to the action of walking. While walking, we alternate moving our left and right legs, and the signature matrix measures this as an influence of rotations about joints in one leg on the rotations about joints in the other leg.

In Figure 8, we have plotted the absolute value of the averaged second level signature matrix for all action classes in the data set. We omit the labelling to simplify the figure, but all labelling is the same as Figure 6. In particular, the colors range from 0 to 1.



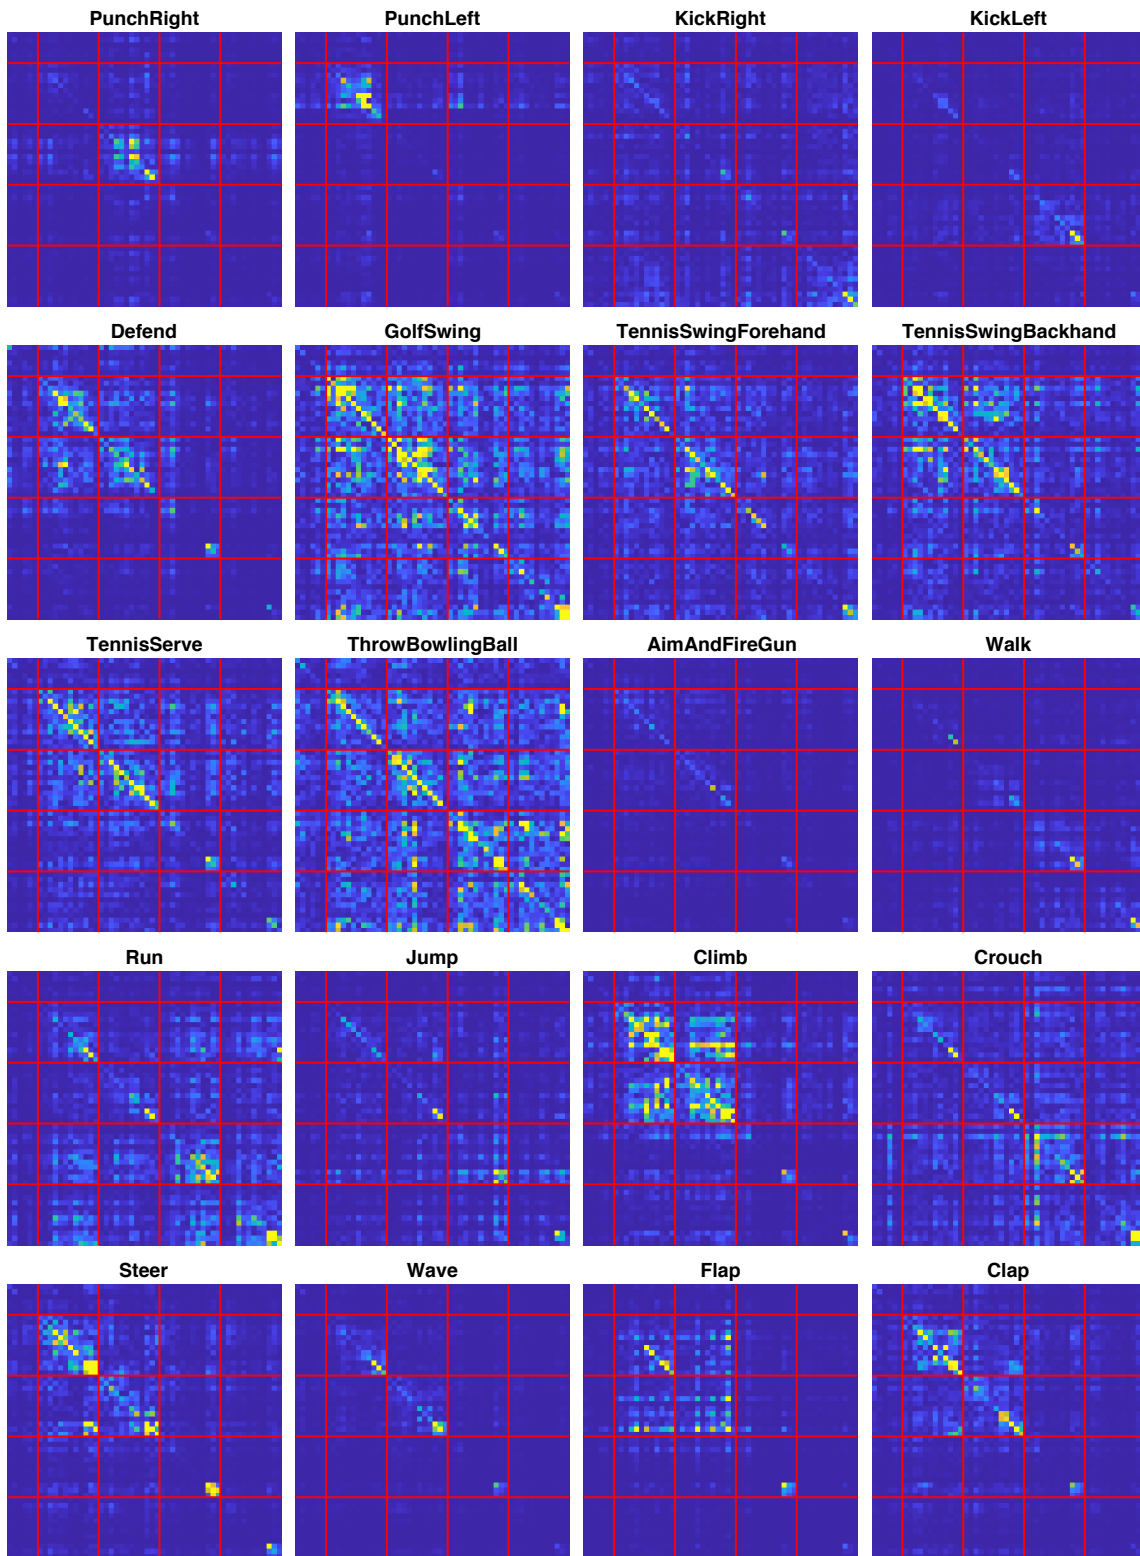


Figure 8: Averaged absolute  $S^2$  matrices for all actions in G3D dataset.

## 5.2 Random walks on Lie groups

In this experiment, we aim to take advantage of the characteristic property of the signature kernel and perform a kernel two-sample hypothesis testing problem, as described in Gretton et al. (2012). We aim to distinguish between two distinct discrete random walks on  $SO(3)$ , or in other words, between two distinct probability distribution on  $P(SO(3))$ . In addition, by using the matrix representation of  $SO(3)$ , we can also view these as probability distributions on  $P\mathbb{R}^9$ , and we also perform the hypothesis testing in this context. We find that for this problem, the path signature defined using left-invariant forms on  $SO(3)$  significantly outperforms the Euclidean signature.

### 5.2.1 HYPOTHESIS TESTING WITH MMD

Suppose  $\mu$  and  $\nu$  are two probability measures on  $PG$  and we have i.i.d. samples  $\{X^i\}_{i=1}^n$  of  $\mu$  and  $\{Y^i\}_{i=1}^n$  of  $\nu$ , each with  $n$  samples. Note that we use the superscript to index the separate samples, since the subscript is reserved for the time parameter. We wish to distinguish between the null hypothesis  $H_0 : \mu = \nu$  against the alternative hypothesis  $H_1 : \mu \neq \nu$ . This is done by computing the unbiased empirical MMD (Equation 23) between our two collections of samples, and comparing this against a chosen threshold. A Type I error occurs when the test falsely rejects the null hypothesis, and a Type II error occurs when the test falsely accepts the null hypothesis. The *level*  $\alpha$  of a test is a chosen upper bound on the probability of a Type I error.

The threshold for the MMD is chosen such that the test is *consistent*, meaning the test achieves a level  $\alpha$  and a Type II error of zero in the asymptotic large sample limit. One method for choosing the threshold is by using large deviation bounds for the MMD under the assumption that the null hypothesis is true, which are derived in Gretton et al. (2012). However, because such bounds must hold for all measures, the threshold chosen in this way is conservative and may not yield optimal results with a finite number of samples. An alternative method is to use a data-dependent threshold. Our method here will be to generate a null distribution for the MMD by performing a permutation test, and computing the  $1 - \alpha$  quantile for a level  $\alpha$  test. Alternative methods for deriving data-dependent thresholds are explored in Gretton et al. (2009, 2012).

### 5.2.2 RANDOM WALK GENERATION

For this experiment, we will consider discrete random walks on  $SO(3)$ . Each walk is initialized by sampling a point uniformly at random from  $SO(3)$ . A method to sample uniformly from  $SO(3)$  is given in Diaconis and Shahshahani (1987). Next, we sample  $T$  i.i.d. points  $\{v_i\}_{i=1}^T$  from the unit sphere in the Lie algebra  $\mathfrak{so}(3)$ . We use a unimodal probability distribution on  $S^2$ , called the von Mises-Fisher distribution. This is parametrized by the mean direction  $x \in S^2$ , and the concentration parameter  $\kappa \geq 0$ , which describes the concentration of the distribution about the mean.

We may think of this as a random walk with drift in the mean direction. A numerically stable method of sampling from this distribution is discussed in Jakob (2012). The random walk  $\hat{\gamma} : [T + 1] \rightarrow SO(3)$  is defined recursively as

$$\hat{\gamma}_{i+1} = \hat{\gamma}_i \exp(cv_i),$$




Figure 9: The von-Mises Fisher density on  $S^2$  with mean direction  $x = (0, 0, 1)$  and  $\kappa = 0.1$ .

where  $c > 0$  is the step size and  $\hat{\gamma}_1$  is a randomly sampled point on  $SO(3)$ .

### 5.2.3 RESULTS

We perform two classes of tests.

1.  $H_0$  is false. Here, we attempt to distinguish between distributions with different mean directions. We take  $x_\mu = (1, 0, 0)$  and  $x_\nu = (0, 1, 0)$ .
2.  $H_0$  is true. Here we take the mean directions of both distributions to be  $x_\mu = x_\nu = (1, 0, 0)$ .

All distributions will have a concentration parameter of  $\kappa = 0.1$ . For each test, we will sample  $n = 50$  random walks from each distribution, and each random walk will have  $L = 100$  steps. To generate the null distribution of the MMD, we perform a permutation test with 2,000 permutations for a given set of samples. A level of  $\alpha = 0.05$  is chosen, so we compute the 0.95 quantile of the null distribution as the threshold used for the MMD. The path signature is truncated at level 4 in the MMD computation. We perform 1,000 tests for each class, and each test is done using both Lie group and Euclidean path signatures.

The following table provides the error rates (false positive/negative) for the two classes of tests using the two methods.

	Lie Group	Euclidean
$H_0$ false	6.1%	83.7%
$H_0$ true	5.7%	4.9%

Table 3: Error rates for hypothesis testing. Each test was run 1,000 times.

Additionally, we provide histograms that summarize the test results. The test distributions show the distribution of  $MMD_u$  over the 1,000 independent trials. The null distribution shown is generated from a permutation test for a single trial. The red line shows the 0.95

quantile, and represents the threshold for that trial. The histograms for  $H_0$  being false are shown in Figure 10 and the histograms for  $H_0$  being true are shown in Figure 11.

We find that the Lie group path signatures significantly outperform Euclidean path signatures. This is due to the fact that the Euclidean representation of the data is ill-suited for this problem. We are aiming to detect a slight drift in the direction of the rotation, which is a translation invariant feature in  $SO(3)$ . However, this is not a translation invariant feature in the Euclidean representation of the problem, so the effect of the drift is confounded in the Euclidean path signature.

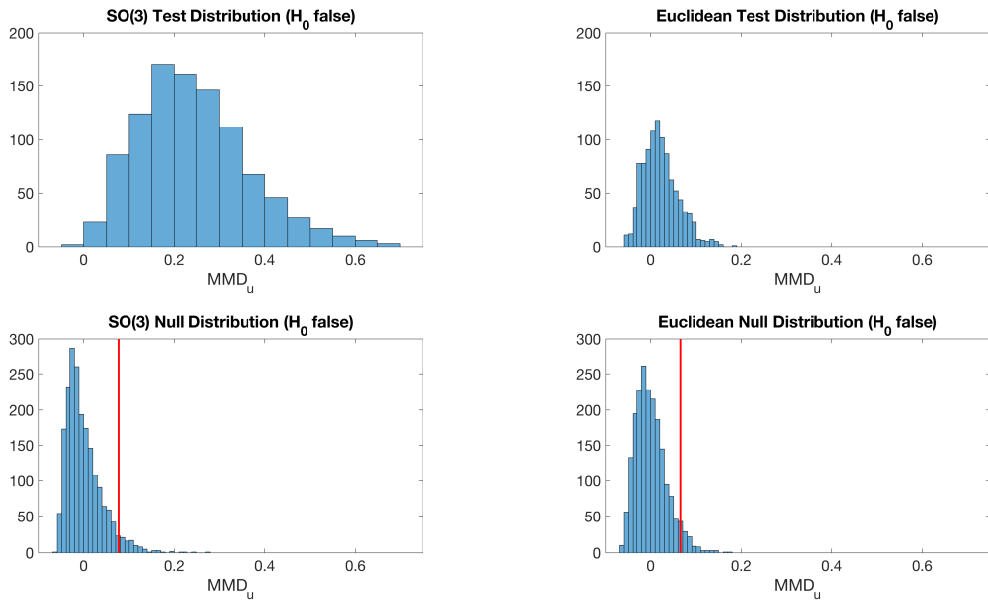


Figure 10: Test (top) and null (bottom) distributions of  $MMD_u$  when  $H_0$  is false.

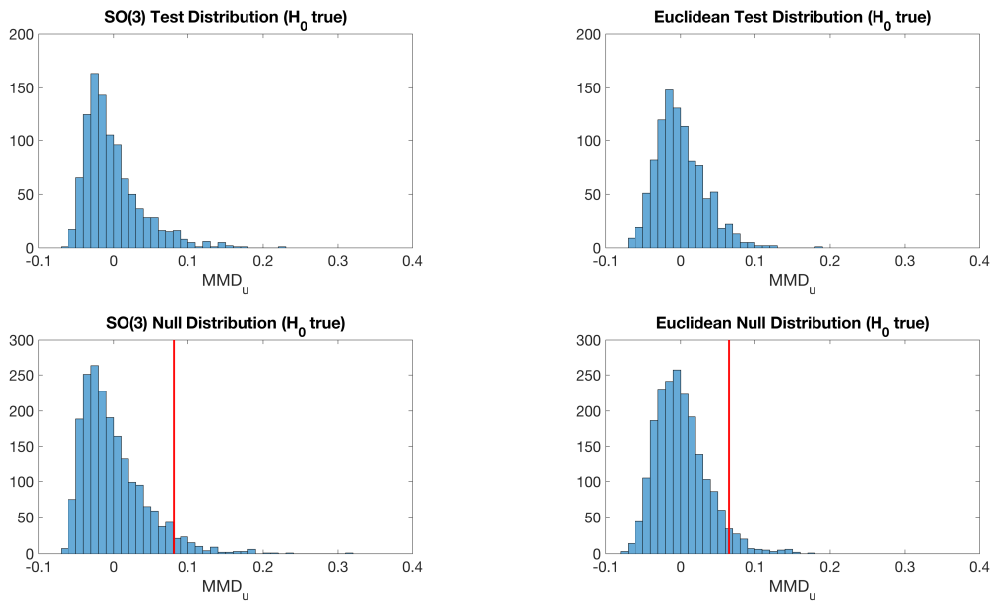


Figure 11: Test (top) and null (bottom) distributions of  $MMD_u$  when  $H_0$  is true.

## 6. Conclusion

We have defined path signature for Lie group valued time series, and studied several of its properties, the main result being the universal and characteristic properties of the signature kernel. By defining signature using only the derivative of the path, computational techniques from Euclidean valued paths to Lie group valued paths can be exported cleanly. Our theory is validated using two detailed experiments highlighting both the universal and characteristic properties, showing that the path signature has strong empirical performance, while providing an interpretable feature set which can be used to better understand the underlying phenomena.

Lie group valued data is ubiquitous; however, in previous studies analyzing such data, the analysis pipeline can be complicated due to the ostensible complexities when dealing with Lie groups (as described in Section 5.1). Our derivations show that Lie group valued data can be treated in a manner nearly identical to standard Euclidean valued data.

The work in this paper provides the foundations for further studying properties and applications of path signatures for Lie groups and more. We highlight two directions for possible future research.

1. Although the path signature is a complete characterization of paths in Euclidean space, the inversion problem of recovering a path given its path signature is difficult (Lyons and Xu, 2017; Chang et al., 2017). Pfeffer et al. (2019) studies a restricted and approximate inversion problem from an algebro-geometric perspective of the path signatures, which was initiated in Améndola et al. (2019).

One may consider a similar question for Lie groups. Suppose we fix a path  $\gamma \in PG$ , and we are given the path signature of a path  $\alpha \in PH$ , where  $\alpha = F\gamma$ , where  $F : G \rightarrow H$  is a Lie group homomorphism. Is it possible to recover the Lie group homomorphism  $F$ ?

2. In Chen (1958), the path signature is defined for a manifold  $M$  by choosing a collection of 1-forms  $\{\omega_i\}_{i=1}^N$  (as no natural Lie algebra basis is available). Choose  $x \in M$  to be a basepoint, and let  $S : PM_x \rightarrow T((\mathbb{R}^N))$  be the path signature defined with respect to these 1-forms and let  $\alpha, \beta \in PM_x$ . If the 1-forms span the cotangent bundle of  $M$  at every point  $x \in M$ , Chen's injectivity result states that  $S(\alpha) = S(\beta)$  if and only if  $\alpha$  and  $\beta$  are tree-like equivalent (in a slightly modified sense). The injectivity theorem given in Theorem 30 is a customization of Chen's result to the case  $M = G$  since a basis of  $\mathfrak{g}^*$  spans the cotangent bundle at every point.

Recently, there has been interest in studying time series evolving on manifolds, and the development of a path signature kernel for paths on manifolds would provide a powerful tool for geometric time series analysis. One of the difficulties of this definition would be the representation of data on manifolds and the choice of 1-forms used. One could begin with path signatures on *parallelizable* manifolds, which by definition admit a smooth basis of vector fields and 1-forms. This would encompass all orientable 3-dimensional manifolds.

## Acknowledgments

D.L. would like to acknowledge Chad Giusti and Jakob Hansen for several helpful discussions throughout this project. D.L. and R.G. are supported by the Office of the Assistant Secretary of Defense Research & Engineering through ONR N00014-16-1-2010. D.L. is also supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) PGS-D3.

## Appendix A. Implementation details of path signature computation

We provide implementation details of the path signature computation for Lie groups, which includes the case for Euclidean space. Several libraries containing path signature computations for Euclidean paths exist, listed below.

1. The `esig` package (Lyons) implemented in C++ and Python was one of the first libraries for path signature computations. It is a CPU-only and single-threaded implementation.
2. The `iisignature` package (Reizenstein and Graham, 2020) is also implemented in Python. It is also a CPU-only and single-threaded implementation.
3. The `Signatory` library (Kidger and Lyons, 2020) is the most recent and complete implementation of signature computations. It is written in C++ with a built-in Python wrapper. It includes functionality such as backpropagation which was not included in the previous packages. In addition, it provides GPU as well as both single and multi-threaded CPU support.

Our current package is the first Julia implementation of path signature computations, and is also the first package to support matrix Lie group valued time series. We provide functions for both the continuous and discrete path signatures. The package can be found in <https://github.com/ldarrick/PathSignatures>.

Suppose  $P : [T + 1] \rightarrow G$  is a discrete time series in an  $N$ -dimensional matrix Lie group  $G \subseteq GL_D(\mathbb{R})$  (elements are  $D \times D$  matrices), and let  $p : [T] \rightarrow \mathfrak{g}$  be the discrete derivative of the time series. Note our change in notation for the paths in order to match the code. The continuous signature computation is performed by using Chen’s identity (Equation 10) as described in Section 3.7. Specifically, once the discrete derivative is computed, the path signature is computed by tensor exponentiation and multiplication:

$$S(P) = \exp_{\otimes}(p_1) \otimes \exp_{\otimes}(p_2) \otimes \dots \otimes \exp_{\otimes}(p_T).$$

Our goal in this section is to compute the truncated discrete signature  $\hat{S}_M(P)$ . Recall that the discrete path signature with respect to the multi-index  $I = (i_1, \dots, i_m) \in [N]^m$  is computed as

$$\hat{S}^I(\hat{\gamma}) := \sum_{(t_1, \dots, t_m) \in \hat{\Delta}_T^m} \omega_{i_1}(p_{t_1}) \dots \omega_{i_m}(p_{t_m}).$$

We note that the path signature computation relies only on the derivative  $p$ , so we split the computation into two steps.

1. Compute the discrete derivative  $p$  from a discrete time series  $P$  in a Lie group  $G$  with respect to a chosen basis of  $\mathfrak{g}$ . Note that this step is dependent on both the Lie group  $G$ , and the choice of basis of  $\mathfrak{g}$ .
2. Compute the truncated discrete path signature  $\hat{S}_M(\hat{\gamma})$  given the discrete derivative  $p$ . Note that this step is independent of both the Lie group  $G$  as well as the choice of basis of  $\mathfrak{g}$ .

This abstraction allows us to write a single path signature function, though we must write a new discrete derivative function for each Lie group we wish to consider. Let us first fix some notation used in the pseudo-code.

1. **Matrix Multiplication.** For a  $T_1 \times T_2$  array  $A$  and a  $T_2 \times T_3$  array  $B$ , we define

$$AB[i, j] = \sum_{k=1}^{T_2} A[i, k]B[k, j].$$

2. **Element-wise Multiplication.** For two arrays  $A$  and  $B$  of the same size, we define

$$(A \cdot B)[i, j] = A[i, j]B[i, j].$$

3. **Transpose.** For an array  $A$ , we define

$$A^\top[i, j] = A[j, i].$$

4. **Sum.** For a  $T_1 \times T_2$  array  $A$ , we define

$$A[\Sigma, \Sigma] = \sum_{i=1}^{T_1} \sum_{j=1}^{T_2} A[i, j].$$

5. **Cumulative Sum.** For a  $T_1 \times T_2$  array  $A$ , we define the  $T_1 \times T_2$  array

$$A[\boxplus, \boxplus][i, j] = \sum_{k_1=1}^i \sum_{k_2=1}^j A[k_1, k_2].$$

6. **End.** For a length  $T$  array  $A$ , we define

$$A[\text{end}] = A[T].$$

7. **Splat.** For a length  $T$  array  $A$ , we define the splat operator

$$A\dots = A[1], A[2], \dots, A[T],$$

to use the entries of  $A$  as arguments for some function.



---

**Algorithm 2:** Discrete Derivative for Matrix Lie Groups

---

**Input** :  $P$  ( $((T + 1) \times D \times D)$  array): A path in the matrix Lie group  $G$   
**Output:**  $p$  ( $(T \times N)$  array): The discrete derivative

- 1 Initialize  $(T \times N)$  array  $p$  ;
- 2 Initialize  $(D \times D)$  array  $p$  ;
- 3 **for**  $t=1..T$  **do**
- 4 Compute the discrete derivative at time  $t$   
 $d \leftarrow \text{logm}(P[t+1, :, :]^{-1}P[t, :, :])$ ;
- 5 Convert the derivative to a representation in terms of a chosen basis of  $\mathfrak{g}$   
 $p[t, :] \leftarrow \text{convert\_to\_basis}(d)$ ;
- 6 **end**
- 7 Return  $p$ ;

---

The discrete derivative function for a path  $P : [T + 1] \rightarrow \mathbb{R}^N$ , which can be represented by a  $((T + 1) \times N)$  array, is simply a difference operation in the first dimension. Now, let's consider a path  $P : [T + 1] \rightarrow G$ , where  $G$  is an  $N$ -dimensional matrix Lie group  $G \subseteq GL_D(\mathbb{R})$ . This path can be represented by a  $((T + 1) \times D \times D)$  array. Suppose that we have implemented a function `convert_to_basis()` which takes in a  $(D \times D)$  array, which represents an element of  $\mathfrak{g}$ , and converts it into a length  $N$  vector. Let `logm()` be the matrix logarithm function.

Next, we will describe the discrete signature computation. Note that the following algorithm is not novel, and is similar to the implementation provided in Kidger and Lyons (2020). The idea is to compute the truncated signature using forward recursion, using the computations for lower levels to also compute higher levels. Thus, this computation is made up of two functions: the initialization and the recursion. The truncated signature will be an element of  $T^{(\leq M)}(\mathbb{R}^N)$ . This is represented as a 1D array of multidimensional arrays, which we label  $s$ . For example,  $s[m]$  is an  $m$ -dimensional array of length  $N$  in each dimension, and we denote access to the inner arrays using separate square brackets. Namely,  $s[2][3, 4]$  is the  $[3, 4]$  element of the 2D array  $s[2]$ .

Here, we may easily compute the complexity of the truncated signature algorithm in terms of the number of elementary operations. We note that there are  $N^M$  entries of the level  $M$  signature, and all lower levels are computed along the way. To compute an entry of the level  $M$  signature from the level  $M - 1$  signature, we require  $O(T)$  elementary operations. Thus, the complexity of this algorithm is  $O(TN^M)$ .

---

**Algorithm 3:** Discrete Signature Initialization
 

---

**Input** :  $p$  ( $(T \times N)$  array): The discrete derivative  
 $M$  (Int): The truncation level  
**Output:**  $s$  (Tensor polynomial): The truncated discrete signature

- 1 Initialize 1D array of multidimensional arrays  $s$ ;
- 2 Initialize length  $M$  array  $cur\_ind$ ;
- 3 Initialize length  $T$  array  $Q$ ;
- 4 **for**  $n=1..N$  **do**
- 5 Set the first index  
 $cur\_ind[1] \leftarrow n$ ;
- 6 Compute first level signature path and store in  $s$   
 $Q \leftarrow K[\boxplus, n]$ ;  
 $s[1][n] \leftarrow Q[end]$ ;
- 7 Perform forward recursion  
 $s \leftarrow sig\_forward(p, Q, cur\_ind, 2, M)$ ;
- 8 **end**
- 9 Return  $s$

---

---

**Algorithm 4:** Signature Forward Recursion (`sig_forward`)

---

**Input** :  $s$  (Tensor polynomial): The current signature  
 $p$  ( $(T \times N)$  array): The discrete derivative  
 $lastQ$  (Length  $T$  array): The previous signature path  
 $cur\_ind$  (Length  $M$  array): The current index  
 $cur\_level$  (Int): The current level  
 $M$  (Int): The truncation level

**Output:**  $s$  (Tensor polynomial): Partially computed truncated discrete signature

- 1 Initialize length  $T$  array  $Q$ ;
- 2 **if**  $cur\_level < M$  **then**
- 3     **for**  $n=1..N$  **do**
- 4         Update the current index  
            $cur\_ind[cur\_level] \leftarrow n$ ;
- 5         Compute the current signature path  
            $Q \leftarrow lastQ \cdot p[:,n]$  ;  
            $Q \leftarrow Q[\boxplus]$  ;
- 6         Store the current signature  
            $s[cur\_level][cur\_ind[1:cur\_level]...] = Q[end]$  ;
- 7         Perform forward recursion  
            $s \leftarrow sig\_forward(p, Q, cur\_ind, cur\_level+1, M)$  ;
- 8     **end**
- 9 **else**
- 10    **for**  $n=1..N$  **do**
- 11         \\ Last level: perform the same operations, but don't perform recursion  
            $cur\_ind[cur\_level] \leftarrow n$
- 12          $Q \leftarrow lastQ \cdot p[:,n]$
- 13          $s[cur\_level][cur\_ind...] \leftarrow Q(\Sigma)$
- 14    **end**
- 15 **end**
- 16 Return  $s$

---

## Appendix B. Tensor Normalization

In this appendix, we will discuss the construction of a tensor normalization, as shown in Appendix A of Chevyrev and Oberhauser (2018), as well as the computation of the normalization. The following proposition provides the construction that we will follow. Suppose  $H$  is a Hilbert space throughout this appendix.

**Proposition 65** *Let  $\psi : [1, \infty) \rightarrow [1, \infty)$  with  $\psi(1) = 1$ . For  $\mathbf{t} \in T_1((H))$  and define  $\lambda : T_1((H)) \rightarrow (0, \infty)$  to be the unique non-negative number such that  $\|\delta_{\lambda(\mathbf{t})}\mathbf{t}\|^2 = \psi(\|\mathbf{t}\|)$ . Define*

$$\begin{aligned} \Lambda : T_1((H)) &\rightarrow T_1((H)) \\ \mathbf{t} &\mapsto \delta_{\lambda(\mathbf{t})}\mathbf{t}. \end{aligned}$$

Denote further  $\|\psi\|_\infty = \sup_{x \geq 1} \psi(x)$ . Then the following holds.

1. The function  $\Lambda$  takes values in the set  $\{\mathbf{t} \in T_1((H)) : \|\mathbf{t}\| \leq \sqrt{\|\psi\|_\infty}\}$ .
2. If  $\psi$  is injective, then so is  $\Lambda$ .
3. Suppose that  $\sup_{x \geq 1} \psi(x)/x^2 \leq 1$ ,  $\|\psi\|_\infty < \infty$ , and that  $\psi$  is  $K$ -Lipschitz for some  $K > 0$ . Then

$$\|\Lambda(\mathbf{s}) - \Lambda(\mathbf{t})\| \leq (1 + \sqrt{K} + 2\sqrt{\|\psi\|_\infty})(\sqrt{\|\mathbf{s} - \mathbf{t}\|} \vee \|\mathbf{s} - \mathbf{t}\|).$$

**Corollary 66** *Let  $\psi : [1, \infty) \rightarrow [1, \infty)$  be injective satisfying  $\psi(1) = 1$  and the conditions of item (3.) in Proposition 65. Then, the function  $\Lambda$  constructed in Proposition 65 is a tensor normalization.*

This method of constructing a tensor normalization is done by defining the normalized norm  $\|\delta_{\lambda(\mathbf{t})}\mathbf{t}\|$  of an element  $\mathbf{t}$ , given its original norm  $\psi(\|\mathbf{t}\|)$ . In practice, we will be working with a truncated element of the tensor algebra. Thus, given an element  $\mathbf{t} \in T_1^{(\leq M)}((H))$ , we can obtain the value of  $\lambda(\mathbf{t})$  by solving the equation  $\|\delta_{\lambda(\mathbf{t})}\mathbf{t}\|^2 = \psi(\|\mathbf{t}\|)$ . This reduces to finding the zero of the polynomial equation

$$\sum_{k=1}^M \lambda^{2k} \|\mathbf{t}_k\|^2 - \psi(\|\mathbf{t}\|) = 0.$$

## References

- Marcos Alexandrino and Renato Bettiol. *Lie Groups and Geometric Aspects of Isometric Actions*. Springer International Publishing, 2015.
- Carlos Améndola, Peter Friz, and Bernd Sturmfels. Varieties of Signature Tensors. *Forum of Mathematics, Sigma*, 7:e10, 2019.
- Imanol Perez Arribas, Kate Saunders, Guy Goodwin, and Terry Lyons. A signature-based machine learning model for bipolar disorder and borderline personality disorder. *arXiv:1707.07124 [stat]*, July 2017.

- Yuliy Baryshnikov and Emily Schlafly. Cyclicity in multivariate time series and applications to functional MRI data. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pages 1625–1630, December 2016.
- Victoria Bloom, Dimitrios Makris, and Vasileios Argyriou. G3D: A gaming action dataset and real time action recognition evaluation framework. In *2012 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, pages 7–12, June 2012.
- Karsten M. Borgwardt, Arthur Gretton, Malte J. Rasch, Hans-Peter Kriegel, Bernhard Schölkopf, and Alex J. Smola. Integrating structured biological data by Kernel Maximum Mean Discrepancy. *Bioinformatics*, 22(14):e49–e57, July 2006.
- Annegret Y Burtscher. Length structures on manifolds with continuous Riemannian metrics. *New York Journal of Mathematics*, 21:273–296, 2015.
- Jiawei Chang, Nick Duffield, Hao Ni, and Weijun Xu. Signature inversion for monotone paths. *Electronic Communications in Probability*, 22, 2017.
- Kuo-Tsai Chen. Iterated Integrals and Exponential Homomorphisms†. *Proceedings of the London Mathematical Society*, s3-4(1):502–512, 1954.
- Kuo-Tsai Chen. Integration of Paths, Geometric Invariants and a Generalized Baker- Hausdorff Formula. *Annals of Mathematics*, 65(1):163–178, 1957.
- Kuo-Tsai Chen. Integration of Paths—A Faithful Representation of Paths by Noncommutative Formal Power Series. *Transactions of the American Mathematical Society*, 89(2): 395–407, 1958.
- Kuo-Tsai Chen. Iterated path integrals. *Bulletin of the American Mathematical Society*, 83 (5):831–879, September 1977.
- Ilya Chevyrev and Andrey Kormilitzin. A Primer on the Signature Method in Machine Learning. *arXiv:1603.03788 [cs, stat]*, March 2016.
- Ilya Chevyrev and Harald Oberhauser. Signature moments to characterize laws of stochastic processes. *arXiv:1810.10971 [math, stat]*, October 2018.
- Ilya Chevyrev, Vidit Nanda, and Harald Oberhauser. Persistence Paths and Signature Features in Topological Data Analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(1):192–202, January 2020.
- Persi Diaconis and Mehrdad Shahshahani. The Subgroup Algorithm for Generating Uniform Random Variables. *Probability in the Engineering and Informational Sciences*, 1(1):15–32, January 1987.
- Joscha Diehl and Jeremy Reizenstein. Invariants of Multidimensional Time Series Based on Their Iterated-Integral Signature. *Acta Applicandae Mathematicae*, 164(1):83–122, December 2019.

- Joscha Diehl, Kurusch Ebrahimi-Fard, and Nikolas Tapia. Time-Warping Invariants of Multidimensional Time Series. *Acta Applicandae Mathematicae*, May 2020.
- Adeline Fermanian. Embedding and learning with signatures. *arXiv:1911.13211 [cs, stat]*, November 2019.
- Peter K. Friz and Nicolas B. Victoir. *Multidimensional Stochastic Processes as Rough Paths: Theory and Applications*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2010.
- Robin Giles. A Generalization of the Strict Topology. *Transactions of the American Mathematical Society*, 161:467–474, 1971.
- Chad Giusti and Darrick Lee. Iterated Integrals and Population Time Series Analysis. In Nils A. Baas, Gunnar E. Carlsson, Gereon Quick, Markus Szymik, and Marius Thauale, editors, *Topological Data Analysis*, Abel Symposia, pages 219–246, Cham, 2020. Springer International Publishing.
- Arthur Gretton, Kenji Fukumizu, Zaïd Harchaoui, and Bharath K. Sriperumbudur. A Fast, Consistent Kernel Two-Sample Test. In Y. Bengio, D. Schuurmans, J. D. Lafferty, C. K. I. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems 22*, pages 673–681. Curran Associates, Inc., 2009.
- Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A Kernel Two-Sample Test. *Journal of Machine Learning Research*, 13: 723–773, March 2012.
- Lajos Gergely Gyurkó, Terry Lyons, Mark Kontkowski, and Jonathan Field. Extracting information from the signature of a financial data stream. *arXiv:1307.7244 [q-fin]*, July 2013.
- Ben Hambly and Terry Lyons. Uniqueness for the signature of a path of bounded variation and the reduced path group. *Annals of Mathematics*, 171(1):109–167, 2010.
- Zhiwu Huang, Chengde Wan, Thomas Probst, and Luc Van Gool. Deep learning on lie groups for skeleton-based action recognition. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- Wenzel Jakob. Numerically stable sampling of the von Mises Fisher distribution on  $S^2$  (and other tricks). 2012.
- Patrick Kidger and Terry Lyons. Signatory: Differentiable computations of the signature and logsignature transforms, on both CPU and GPU. *arXiv:2001.00706 [cs, stat]*, January 2020.
- Franz J. Kiraly and Harald Oberhauser. Kernels for Sequentially Ordered Data. *Journal of Machine Learning Research*, 20(31):1–45, 2019.

- Stefan Knerr, Léon Personnaz, and Gérard Dreyfus. Single-layer learning revisited: A stepwise procedure for building and training a neural network. In Françoise Fogelman Soulié and Jeanny Hérault, editors, *Neurocomputing*, NATO ASI Series, pages 41–50, Berlin, Heidelberg, 1990. Springer.
- Enrico Le Donne. A Primer on Carnot Groups: Homogenous Groups, Carnot-Carathéodory Spaces, and Regularity of Their Isometries. *Analysis and Geometry in Metric Spaces*, 5 (1):116, 2017.
- John M. Lee. *Introduction to Smooth Manifolds*. Graduate Texts in Mathematics. Springer-Verlag, New York, 2003.
- Yanshan Li, Tianyu Guo, Xing Liu, and Rongjie Xia. Skeleton-based Action Recognition with Lie Group and Deep Neural Networks. In *2019 IEEE 4th International Conference on Signal and Image Processing (ICSIP)*, pages 26–30, July 2019.
- Terry Lyons. esig. <https://esig.readthedocs.io/en/latest/>.
- Terry Lyons. Differential equations driven by rough signals. *Revista Matemática Iberoamericana*, 14(2):215–310, 1998.
- Terry Lyons. Rough paths, Signatures and the modelling of functions on streams. *arXiv:1405.4537 [math, q-fin, stat]*, May 2014.
- Terry Lyons and Zhongmin Qian. *System Control and Rough Paths*. Clarendon, Oxford, 2007.
- Terry Lyons and Weijun Xu. Hyperbolic development and inversion of signature. *Journal of Functional Analysis*, 272(7):2933–2955, April 2017.
- Terry Lyons, Michael Caruana, and Thierry Lévy. *Differential Equations Driven by Rough Paths: Ecole d’Été de Probabilités de Saint-Flour XXXIV-2004*. École d’Été de Probabilités de Saint-Flour. Springer-Verlag, Berlin Heidelberg, 2007.
- Terry Lyons, Hao Ni, and Harald Oberhauser. A feature set for streams and an application to high-frequency financial tick data. In *BigDataScience ’14*, 2014.
- P. J. Moore, T. J. Lyons, J. Gallacher, and Alzheimer’s Disease Neuroimaging Initiative. Using path signatures to predict a diagnosis of Alzheimer’s disease. *PloS One*, 14(9): e0222212, 2019.
- Max Pfeffer, Anna Seigal, and Bernd Sturmfels. Learning paths from signature tensors. *SIAM Journal on Matrix Analysis and Applications*, 40(2):394–416, 2019.
- Philipp Probst, Marvin N. Wright, and Anne-Laure Boulesteix. Hyperparameters and tuning strategies for random forest. *WIREs Data Mining and Knowledge Discovery*, 9 (3):e1301, 2019.
- Jeremy F. Reizenstein and Benjamin Graham. Algorithm 1004: The Iisignature Library: Efficient Calculation of Iterated-Integral Signatures and Log Signatures. *ACM Transactions on Mathematical Software*, 46(1):8:1–8:21, March 2020.

- Manel Rhif, Hazem Wannous, and Imed Riadh Farah. Action Recognition from 3D Skeleton Sequences using Deep Networks on Lie Group Features. In *2018 24th International Conference on Pattern Recognition (ICPR)*, pages 3427–3432, August 2018.
- J. M. Selig. Lie Groups and Lie Algebras in Robotics. In Jim Byrnes, editor, *Computational Noncommutative Algebra and Applications*, NATO Science Series II: Mathematics, Physics and Chemistry, pages 101–125, Dordrecht, 2004. Springer Netherlands.
- Carl-Johann Simon-Gabriel and Bernhard Schölkopf. Kernel Distribution Embeddings: Universal Kernels, Characteristic Kernels and Kernel Metrics on Distributions. *Journal of Machine Learning Research*, 19(44):1–29, 2018.
- Bharath K. Sriperumbudur, Arthur Gretton, Kenji Fukumizu, Bernhard Schölkopf, and Gert R. G. Lanckriet. Hilbert Space Embeddings and Metrics on Probability Measures. *Journal of Machine Learning Research*, 11(Apr):1517–1561, 2010.
- Steven H. Strogatz. From Kuramoto to Crawford: Exploring the onset of synchronization in populations of coupled oscillators. *Physica D: Nonlinear Phenomena*, 143(1):1–20, September 2000.
- Raviteja Vemulapalli and Rama Chellappa. Rolling Rotations for Recognizing Human Actions from 3D Skeletal Data. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4471–4479, June 2016.
- Raviteja Vemulapalli, Felipe Arrate, and Rama Chellappa. Human Action Recognition by Representing 3D Skeletons as Points in a Lie Group. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 588–595, June 2014.
- Weixin Yang, Lianwen Jin, and Manfei Liu. DeepWriterID: An End-to-End Online Text-Independent Writer Identification System. *IEEE Intelligent Systems*, 31(2):45–53, March 2016.
- Weixin Yang, Terry Lyons, Hao Ni, Cordelia Schmid, and Lianwen Jin. Developing the Path Signature Methodology and its Application to Landmark-based Human Action Recognition. *arXiv:1707.03993 [cs]*, December 2019.
- Benjamin J. Zimmerman, Ivan Abraham, Sara A. Schmidt, Yuliy Baryshnikov, and Fatima T. Husain. Dissociating tinnitus patients from healthy controls using resting-state cyclicity analysis and clustering. *Network Neuroscience*, pages 1–23, April 2018.