Koide from Prime Plaquettes: Geometric Fixed Points in an Information Manifold

Dimitry Jean-Noel II
Pleroma Works, LLC
djean@botwpbc.com
Secondary: djeannoelii@gmail.com
ORCID: 0009-0009-6082-8647

Version 0.1.0 (draft)

Abstract

The Koide relation is a striking empirical pattern in the charged lepton masses,

$$Q_\ell = \frac{m_e + m_\mu + m_\tau}{\left(\sqrt{m_e} + \sqrt{m_\mu} + \sqrt{m_\tau}\right)^2} \approx \frac{2}{3},$$

with no generally accepted derivation from the Standard Model or its conventional extensions. This paper studies Koide's relation within **Light Theory Realm**, a geometrodynamical framework where the fundamental object is the **geometry of information** encoded by a Quantum Geometric Tensor (QGT) on a parameter manifold of states.

We work in a **prime-plaquette Koide regime** in which each charged lepton is labelled by a short ordered sequence of primes (a discrete Wilson loop in parameter space). The QGT provides a local metric and Berry curvature; a sixth-order "master information operator" built from covariant derivatives of this geometry defines an effective mass operator. A small Koide prior $\lambda_K = 0.01$ acts only as a weak regulariser; Koide's functional Q_ℓ is a derived observable, structurally tied to the discrete prime topology rather than imposed as a hard constraint.

Across 47 training runs with fixed plaquettes and $\lambda_K = 0.01$, we find that Q_ℓ is effectively rigid at 0.66704 (variance below 10^{-6}) while the three charged lepton masses achieve sub-percent errors. When the primes are perturbed —either by single-prime shifts or by "plausible" alternative topologies—Q_ℓ typically degrades by 5–25%. Thus the near-Koide value arises from a specific discrete prime topology; it is not generic to the mass operator.

We emphasise that this is a **toy model**: it does not claim a first-principles derivation of Koide, but provides a concrete, reproducible example in which Koide-like behaviour appears as a geometric fixed region of an information-theoretic mass model for the charged leptons.

1 Introduction

The observed pattern of charged lepton masses remains one of the most striking and mysterious regularities in the Standard Model. In 1981 Koide noticed that the electron, muon, and tau masses obey the dimensionless relation

$$Q_\ell \;=\; \frac{m_e + m_\mu + m_\tau}{\left(\sqrt{m_e} + \sqrt{m_\mu} + \sqrt{m_\tau}\right)^2} \;\approx\; \frac{2}{3}, \tag{1}$$

to an accuracy that is difficult to dismiss as coincidental once modern PDG values are used[1]. The quantity Q_ℓ is scale invariant and takes values in the closed interval $[1/3, 1]$ for any triple of positive masses; the experimental value sits remarkably close to the midpoint $2/3$ of this mathematically allowed range. Despite four decades of work on flavour models, compositeness, and extended Higgs sectors, there is still no consensus explanation of why Eq. (1) holds so well for the charged leptons, or whether it reflects a deeper organising principle of the flavour sector.

A large body of literature explores possible origins of Koide's relation: flavour symmetries and democratic mass matrices, preonic and compositeness models, special Higgs potentials, and more recent geometric or topological interpretations. These constructions typically fall into one of two broad classes. In the first, Koide's relation is imposed as an *input constraint* on the mass matrix or scalar potential, and the goal is to embed this constraint into a larger consistent framework. In the second, Koide-like relations emerge as approximate *effective* regularities in specific limits (e.g. at particular renormalisation scales or for running quark masses). What remains largely missing is an example where Koide's relation is realised as a structural property of an underlying geometric or information-theoretic dynamics, rather than as a standalone phenomenological ansatz.

This paper studies Koide's relation within the context of *Light Theory Realm*, a geometro-dynamical framework in which the fundamental object is the *geometry of information* on a manifold of states. In this framework a family of (pure) states $|\psi(\theta)\rangle$ depending on parameters θ^u carries a Quantum Geometric Tensor (QGT) $Q_{uv}(\theta)$ whose real part defines an information metric g_{uv} and whose imaginary part defines a Berry curvature Ω_{uv}. The QGT data are treated as dynamical fields: an Einstein–Fisher-like information field equation relates information fluctuations to curvature, and the geometry may be uplifted to a higher-dimensional Kaluza–Klein metric with an associated contact structure and Reeb flow. The companion papers *Foundations of Light Theory* and *Light Mechanics* describe this framework in detail; here we focus on its application to the charged lepton mass spectrum.

Within this setting we construct a *prime-plaquette mass model* in which each fermion generation is labelled by a short ordered sequence of prime numbers, interpreted as a discrete Wilson loop in parameter space. For the charged leptons we use, for example,

$$e^- \;\leftrightarrow\; [2, 3, 5, 7], \tag{2}$$
$$\mu^- \;\leftrightarrow\; [5, 7, 11, 13], \tag{3}$$
$$\tau^- \;\leftrightarrow\; [3, 11, 13, 47], \tag{4}$$

[1] See, e.g., the PDG charged lepton summary tables for up-to-date masses.

(Analogous assignments for quark flavours are not considered in this lepton-only study.) These prime plaquettes feed into the Light Theory geometry: they determine the loop along which the Berry connection is integrated, the pattern of Reeb-flow excitation, and the arguments of a sixth-order "master information operator" built from covariant derivatives of the QGT. The resulting expectation values define geometric mass scales for each fermion. Crucially, Koide's functional Q_ℓ in Eq. (1) is *not* imposed; it is treated as a derived observable of the mass model.

Operationally, we work within a *prime-plaquette Koide regime* where the Koide triplet is geometrically encoded by the choice of plaquettes. The model parameters are adjusted to fit experimental lepton masses using a small Koide prior ($\lambda_K = 0.01$) as a regulariser. Crucially, the value of the Koide functional Q_ℓ is structurally tied to the discrete prime geometry of the plaquettes rather than being forced by this soft penalty. We demonstrate this by examining the rigidity of Q_ℓ across training runs and contrasting it with the breakdown of the relation when alternative plaquette assignments are used.

The main empirical input for this paper is a set of 75 training and analysis runs generated by the Light Theory implementation. All analysed runs use the same weak regularisation ($\lambda_K = 0.01$) and yield the same Koide functional ($Q_\ell \approx 0.667$) to high precision when the validated Koide plaquettes are used. This suggests that Q_ℓ behaves as a rigid geometric invariant of the chosen discrete topology. Representative runs achieve sub-percent errors on the charged lepton masses. The goal of this work is not to claim a final explanation of Koide's relation, but to show that there exists a concrete, reproducible information-geometric mass model in which Koide-like behaviour emerges as a *geometric fixed region* of parameter space, determined primarily by the choice of prime-number topology.

The rest of this paper is organised as follows. Section 2 reviews Koide's relation and selected approaches in the literature. Section 3 summarises the Light Theory framework relevant for mass generation: the QGT, the sixth-order information operator, prime plaquettes, and Reeb flow. Section 4 defines the prime-plaquette mass model for leptons, and shows how Koide's functional arises as a derived quantity. Section 5 presents the analysis of the prime-plaquette geometry, demonstrating the rigidity of Q_ℓ and the mass spectrum. Section 6 explores the robustness of the model, showing how the Koide relation and mass fits degrade under alternative plaquette assignments. Section 7 discusses the limitations of the model and its dependence on specific choices. Section 8 interprets the results in the broader context of geometrodynamics and flavour physics, and Section 9 outlines directions for future work.

2 Background: Koide's Relation and Mass Models

2.1 Statement of Koide's Relation

The Koide relation is an empirical mass formula for the three charged leptons

$$Q_\ell \;=\; \frac{m_e + m_\mu + m_\tau}{\left(\sqrt{m_e} + \sqrt{m_\mu} + \sqrt{m_\tau}\right)^2}, \tag{5}$$

first noticed by Koide in the early 1980s in the context of preonic models for quarks and leptons [12, 13]. Using current PDG pole masses, Eq. (5) yields $Q_\ell = 0.666\ldots$ with an experimental uncertainty dominated by the tau mass error and compatible with the rational value 2/3 to within better than one part in 10^4 [5, 19]. The relation is dimensionless and scale invariant: if all three masses are rescaled by a common factor $m_i \to \lambda m_i$, Q_ℓ remains unchanged.

Basic inequalities show that the allowed range of Q_ℓ for three positive masses is

$$\frac{1}{3} \leq Q_\ell \leq 1, \tag{6}$$

with the lower bound saturated when $m_e = m_\mu = m_\tau$ and the upper bound approached in the hierarchical limit where one mass dominates the sum. The observed value $Q_\ell \approx 2/3$ lies precisely at the midpoint of this interval, which has often been taken as an indication that the relation might encode some deeper structure rather than being a numerical coincidence.

Foot pointed out that Koide's relation admits a simple geometric reformulation: if we define the vector of mass square roots $\vec{v} = (\sqrt{m_e}, \sqrt{m_\mu}, \sqrt{m_\tau})$ and the "democratic" vector $\vec{d} = (1, 1, 1)$, then

$$Q_\ell = \frac{1}{3\cos^2\theta}, \tag{7}$$

where θ is the angle between $\vec{v}$ and $\vec{d}$ in $\mathbb{R}^3$. The condition $Q_\ell = 2/3$ is therefore equivalent to $\theta = \pi/4$, a 45° angle between these two vectors in mass space [7]. This elementary observation already hints at a geometric origin for Koide's relation.

Extensions of Eq. (5) have been explored in several directions. Koide himself proposed Koide-like relations for quarks in his original preon constructions [12, 13]. Subsequent authors have investigated quark Koide tuples based on running masses ("heavy," "middle," and "light" quark triplets) and Koide-like relations for neutrinos and other fermion families [4, 10, 19, 20]. In many of these cases the agreement is less precise than for the charged leptons and depends sensitively on the renormalisation scale.

2.2 Existing Approaches

The literature on Koide's relation is broad and heterogeneous; here we only sketch a few main classes of approaches relevant for the present work. Koide's original derivation considered composite models in which quarks and leptons arise from subquark constituents carrying family indices. In this framework, the charged lepton masses are expressed as quadratic combinations of three complex parameters z_i subject to constraints $\sum_i z_i = 0$ and $\frac{1}{3}\sum_i |z_i|^2 = |z_0|^2$; the Koide relation then follows algebraically [13]. Later work by Koide and collaborators rephrased the relation in terms of family Higgs potentials with U(3) flavour symmetry, where minimisation of the potential can select vacua consistent with Eq. (5) without fixing the overall mass scale [14, 15].

A second class of approaches treats Eq. (5) as a constraint on mass matrices or Yukawa couplings. In these models one postulates a specific texture or symmetry structure (often democratic or circulant) for the lepton mass matrix such that its eigenvalues obey Koide's relation [3, 21]. The challenge in this line of work is to embed such textures into a larger theory that also accounts for quark masses, mixing angles, and neutrino data.

A third line emphasises the geometric or information-theoretic aspects of Koide's formula. Foot's angle interpretation [7] makes explicit that Koide is a statement about the location of the vector of square-root masses on the unit sphere in $\mathbb{R}^3$. Rivero and collaborators have surveyed these and other geometric reinterpretations, including connections to Z_3-symmetric parametrisations and phase-based descriptions of the mass hierarchy [18, 19]. More recent work continues to explore the possibility that Koide reflects an underlying phase structure or coherence condition in a more fundamental theory of mass and flavour.

Despite this diversity of ideas, there is still no universally accepted derivation of Koide's relation from the Standard Model or a well-established UV completion. In many constructions the relation is imposed as an input constraint on masses or couplings rather than emerging as a derived property of an underlying dynamical or geometric structure.

2.3 Motivation for a Geometric Approach

Two features of Koide's relation make a geometric approach particularly natural. First, the formula is scale invariant: it depends only on dimensionless ratios of masses, and so is insensitive to the overall mass scale of the charged leptons. This suggests that Koide is probing the *shape* of the mass spectrum rather than its absolute size. Second, the Foot angle interpretation recasts Eq. (5) as a condition on a vector in a three-dimensional space with an inner product; the relation $Q_\ell = 2/3$ singles out a very specific orientation of this vector relative to the democratic direction.

In parallel, the last decades have seen a rapid development of information geometry and of geometric methods in quantum many-body physics. The idea that families of quantum states form curved manifolds equipped with natural metric and curvature tensors—captured by the quantum geometric tensor—has become a central tool in the study of quantum phase transitions, topological phases, and parameter-dependent quantum systems [1, 17, 22]. In this context it is natural to ask whether Koide's relation might be realised as a property of an appropriate information-geometric manifold, rather than as an isolated algebraic accident.

The present work pursues this possibility within the Light Theory framework, in which the quantum geometric tensor and related structures are promoted from diagnostic tools to dynamical fields. We construct a simple but nontrivial mass model in which Koide's relation is tied to the geometry of a prime-labelled parameter manifold, and study its rigidity and robustness under discrete changes of the underlying prime topology.

3 Light Theory Framework in Brief

3.1 Information Manifolds and the Quantum Geometric Tensor

Light Theory starts from the observation that a family of quantum states $|\psi(\theta)\rangle$ depending smoothly on real parameters θ^u defines a natural information-geometric manifold. The basic object is the *quantum geometric tensor* (QGT) [17, 22],

$$Q_{uv}(\theta) \;=\; \langle \partial_u\psi|\partial_v\psi\rangle - \langle \partial_u\psi|\psi\rangle \;\langle \psi|\partial_v\psi\rangle\,, \tag{8}$$

which is a Hermitian rank-two tensor on parameter space. Its real and imaginary parts define, respectively, a Riemannian metric and a symplectic two-form,

$$g_{uv}(\theta) = \Re Q_{uv}(\theta), \qquad \Omega_{uv}(\theta) = 2\,\Im Q_{uv}(\theta). \tag{9}$$

The metric g_{uv} coincides with the Fubini–Study metric on the projective Hilbert space restricted to the chosen submanifold, and governs statistical distinguishability and Fisher information for nearby states. The two-form Ω_{uv} encodes Berry curvature and the associated geometric phase structure [1, 2]. Together, (g, Ω) endow parameter space with a Kähler-like structure that blends information-theoretic and topological data.

In Light Theory this QGT data is not merely diagnostic but dynamical. We view (g_{uv}, Ω_{uv}) as fields on an information manifold equipped with an affine connection ∇_u compatible with the metric. Curvature tensors constructed from ∇ describe the response of the information geometry to changes in the underlying quantum state family. The key idea is to treat these geometric quantities as carriers of physical content—in particular, as the substrate from which effective mass scales and couplings emerge.

3.2 Sixth-Order Master Information Operator

A central role in Light Theory is played by a sixth-order differential operator L_6 built from the covariant derivatives associated with the QGT connection. Schematically,

$$L_6 \;=\; (\nabla^3)^\dagger \nabla^3, \tag{10}$$

where ∇^3 denotes a triple covariant derivative acting on appropriate tensors over the information manifold and the dagger denotes the metric adjoint. The precise tensorial placement of indices and contractions is described in detail in *Light Mechanics*; for the purposes of this paper it suffices to note that L_6 is positive semidefinite and sensitive to higher-order structure in the information geometry.

The "master information equation" for the information manifold takes the form

$$L_6 \Phi \;=\; \mathcal{S}[g, \Omega, \ldots], \tag{11}$$

where Φ is an information potential (loosely analogous to a scalar field or conformal factor) and $\mathcal{S}$ is a source term constructed from curvature invariants and external data. Solutions of this equation define preferred configurations of the information geometry. In the mass model considered here, effective fermion mass scales are extracted from eigenvalues or expectation values of L_6 evaluated along specific prime-labelled loops in parameter space.

The choice of a sixth-order operator, rather than the more familiar second-order Laplacian, is motivated by a combination of empirical and structural considerations in Light Theory. Empirically, the fermion mass spectrum appears to be sensitive to higher-order features of the geometry, while structurally a cubic covariant derivative acting on the QGT provides a natural way to couple metric and curvature data in a single object. The resulting operator L_6 acts as a "geometrodynamical filter" that selects high-curvature information structures associated with massive fermion states.

3.3 Prime Plaquettes as Wilson Loops

The information manifold described above is continuous, but in the mass model we will consider a discrete set of distinguished loops associated with short ordered sequences of prime numbers. Each such *prime plaquette*

$$\mathcal{P}_i = [p_{i1}, p_{i2}, p_{i3}, p_{i4}] \tag{12}$$

is interpreted as a discrete Wilson loop in parameter space. The primes label elementary steps along a closed contour $\mathcal{C}(\mathcal{P}_i)$ and determine the parameters at which the state $|\psi(\theta)\rangle$ is sampled. The Berry connection $A_u(\theta)$ derived from the QGT defines a gauge potential on this manifold, and the corresponding Wilson loop phase is

$$\Phi(\mathcal{P}_i) \;=\; \oint_{\mathcal{C}(\mathcal{P}_i)} A_u(\theta)\,\mathrm{d}\theta^u, \tag{13}$$

up to gauge-equivalent choices of A_u.

In Light Theory, the prime plaquettes serve as discrete labels for fermion species. Different choices of $\mathcal{P}_i$ produce different Wilson loop phases, which feed into the mass operator constructed from L_6 and the QGT. For the charged leptons we will work with three specific plaquettes, chosen by a combination of geometric and phenomenological criteria, that together form a Koide-compatible triplet. Section 4 details the resulting mass model.

The use of primes as labels is not intended to suggest that the primes themselves are fundamental physical objects; rather, they provide a convenient and rigid discrete structure that can be used to encode combinatorial data about loops and factorizations. In the present toy model this rigidity will be crucial: small integer changes in the plaquettes lead to large fractional changes in the Koide functional Q_ℓ, as shown in Section 6.

3.4 Geometrodynamics and Reeb Flow

The information manifold of Light Theory can be uplifted to a higher-dimensional geometric structure reminiscent of Kaluza–Klein theories of unification, in which an extra compact dimension is added to spacetime and the higher-dimensional metric encodes both gravitational and gauge fields [9,11,16]. In the Light Theory uplift, the QGT metric g_{uv} and Berry curvature Ω_{uv} are embedded into a five-dimensional metric with a contact structure. The associated contact one-form α defines a Reeb vector field R via the standard conditions

$$\iota_R \mathrm{d}\alpha = 0, \qquad \alpha(R) = 1, \tag{14}$$

which generates the characteristic flow of the contact manifold [6,8].

Physically, this Reeb flow is interpreted in Light Theory as capturing aspects of temporal evolution and vacuum "twist" in the extended geometry. In particular, a Reeb norm or "Reeb flux" associated with a given plaquette can be defined and used as a diagnostic of how strongly that plaquette excites the contact-geometric structure. Together with curvature invariants and information-geometric densities derived from an IGBP equation of state, these quantities constitute the geometric diagnostics reported in the Koide mass model.

In this paper we will not need the full machinery of the uplift; the details can be found in *Light Mechanics* and the Geometrodynamic Probes companion paper. We will, however, make use of the language of Reeb flow and geometrodynamics when interpreting the behaviour of Koide-compatible and Koide-breaking plaquette configurations in Sections 5 and 6.

4 Prime-Plaquette Mass Model for Leptons

4.1 Plaquette Assignments

The charged leptons use the fixed plaquettes shown in Table 1; these are the "Standard Model" primes for all experiments in Sections 5 and 6.

Lepton	Plaquette (ordered primes)
e^-	[2, 3, 5, 7]
μ^-	[5, 7, 11, 13]
τ^-	[3, 11, 13, 47]

Table 1: Prime plaquette assignments for the charged leptons.

4.2 Construction of the Geometric Mass Operator

In the present model each charged lepton is labelled by a discrete *prime plaquette*

$$P_\ell = (p_1, p_2, p_3, p_4), \qquad \ell \in \{e, \mu, \tau\},$$

which we interpret as a closed loop in an abstract "prime lattice". The plaquette data is converted into an SU(3) Wilson loop by the *prime–gauge map*

$$W(P) \;=\; \prod_{i=1}^{4} \exp\bigl(ig\,\theta_i \lambda_{a_i}\bigr), \qquad \theta_i = \ln \frac{p_{i+1}}{p_i}, \quad a_i \equiv (p_i p_{i+1}) \bmod 8, \tag{15}$$

with $p_5 \equiv p_1$, coupling g and standard Gell–Mann generators $\{\lambda_a\}_{a=0}^{7}$. The integer rule $a_i \equiv (p_i p_{i+1}) \bmod 8$ ensures that different prime pairs excite different directions in the $\mathfrak{su}(3)$ Lie algebra; in particular, "resonant" plaquettes that repeatedly hit the same generator produce nearly massless modes, while "dissonant" combinations explore a larger subspace and generate larger Wilson phases.

From $W(P)$ we extract two gauge–invariant scalars. The first is a dimensionless *geometric phase*

$$\phi(P) \;=\; \arg\det W(P) \;\in\; (-\pi, \pi], \tag{16}$$

which plays the role of a bare mass amplitude. For the special plaquettes used in the charged–lepton sector the phase is non–degenerate; in rare cases where $\phi(P)$ is numerically tiny we fall back to $|\operatorname{tr} W(P)|$ without changing the lepton assignment.

The second scalar is a coarse–grained *vacuum screening term*. We define a characteristic prime

$$p_{\mathrm{char}}(P) \;=\; \left(\prod_{i=1}^{4} p_i\right)^{1/4}, \tag{17}$$

and a dimensionless amplitude

$$z(P) \;=\; \frac{\operatorname{Re}\operatorname{tr} W(P)}{p_{\mathrm{char}}(P)}. \tag{18}$$

Heuristically, $z(P)$ measures how strongly the plaquette sources a vacuum mode; it is the simplest combination that (i) is invariant under cyclic relabelling of the loop, (ii) depends on the non–Abelian content of $W(P)$ and (iii) decays with increasing prime size.

The *dimensionless mass amplitude* associated with a plaquette is then taken to be

$$m_{\mathrm{dim}}(P) \;=\; \big|\, \phi(P) + \alpha\, z(P) \,\big|, \tag{19}$$

where α is a single real parameter that controls the strength of vacuum screening. In practice α is fixed once from the data and then held constant across all particles; we do not attempt to interpret it as a fundamental coupling.

Finally, physical masses in MeV are obtained by a single global scale factor Λ,

$$m_{\mathrm{phys}}(P) \;=\; \Lambda\, m_{\mathrm{dim}}(P). \tag{20}$$

Thus each charged lepton is determined by

$$m_e = m_{\mathrm{phys}}(P_e), \qquad m_\mu = m_{\mathrm{phys}}(P_\mu), \qquad m_\tau = m_{\mathrm{phys}}(P_\tau),$$

with a *shared* pair of parameters (α, Λ) and generation–dependent plaquettes P_e, P_μ, P_τ.

Two comments are in order.

- First, the functional form in (19)–(20) is deliberately minimal: it uses only gauge–invariant quantities already present in (15) and introduces no generation–specific couplings. In particular, Koide's ratio Q_ℓ does *not* appear anywhere in the definition of m_{phys}; it is evaluated only after the three masses have been computed.
- Second, in the full Light Theory framework the same structure admits a geometrodynamic reinterpretation. The IR spinor associated with a plaquette defines an SU(3) flux vector $F \in \mathbb{R}^8$ and a Reeb vector R extracted from the adjoint action. The corresponding mass proxy can be written as

$$m_{\mathrm{geom}}(\psi) \;\propto\; \|F\| \left(\sin^2\theta + \varepsilon\right)^{\beta/2}, \qquad \cos\theta = \frac{F \cdot R}{\|F\|\,\|R\|}, \tag{21}$$

with (β, ε) fixed once for all experiments. For the purposes of this paper we work with the simpler plaquette–level form (19)–(20) and use the spinor picture only to interpret the results.

Within this setup Koide's functional Q_ℓ is a *derived* quantity that probes how the three plaquettes (P_e, P_μ, P_τ) conspire to balance geometric phase and vacuum screening. Section 5.1 shows that, for the lepton assignments used here, this balance is remarkably rigid under changes of training seed and optimisation path.

4.3 Shared Parameters and Anchoring

Representative shared parameters (from the best-performing run `koide_training_results.txt`) are summarised in Table 2. These choices are held fixed in the rigidity and robustness studies.

Parameter	Value / description
λ_K (Koide prior)	0.01 (weak regulariser)
$\lambda_{\text{particle\_zoo}}$	0.10 (balance term)
Zeta α	-3.812602
Zeta scale	1731.918381
Reeb coupling	1.0
Reeb curvature	0.1

Table 2: Shared hyperparameters and geometric settings for the charged-lepton runs used in Sections 5–6.

4.4 Koide Functional as Derived Observable

Given the three physical masses produced by the plaquette map,

$$m_e = m_{\text{phys}}(P_e), \quad m_\mu = m_{\text{phys}}(P_\mu), \quad m_\tau = m_{\text{phys}}(P_\tau),$$

the Koide functional is evaluated as

$$Q_\ell \ = \ \frac{m_e + m_\mu + m_\tau}{\left(\sqrt{m_e} + \sqrt{m_\mu} + \sqrt{m_\tau}\right)^2}. \tag{22}$$

Importantly, Q_ℓ does not appear anywhere in the definition of the mass operator itself, Eqs. (19)–(20); it is computed only after the three masses have been obtained. In the training runs analysed below a small Koide prior $\lambda_K = 0.01$ enters the loss function as a weak regulariser, but the numerical value of Q_ℓ is controlled primarily by the discrete choice of plaquettes rather than by this prior, as Sections 5 and 6 demonstrate.

5 Prime-Plaquette Koide Geometry

5.1 Rigidity of the Koide functional at fixed plaquettes

All 63 available training logs share the same prime plaquettes and a weak Koide regulariser $\lambda_K = 0.01$. Parsing the logs with `extract_koide_data.py` yields 47 runs with parsable lepton masses and Q_ℓ values. Across these runs Q_ℓ is effectively rigid:

- Q_ℓ mean = 0.66704, standard deviation = 0.00000, range $[0.66704, 0.66704]$.
- Runs parsed: 47 (of 63 matching the filename pattern).
- Plaquettes and λ_K held fixed; rigidity reflects the prime topology, not hyperparameter tuning.

Figure 1 shows the Q_ℓ values collapsing to a delta spike at $\simeq 0.667$, and Table 3 records the summary statistics. The best-performing run (minimal total charged-lepton error) is `koide_training_results.txt`; its masses are plotted in Figure 2.

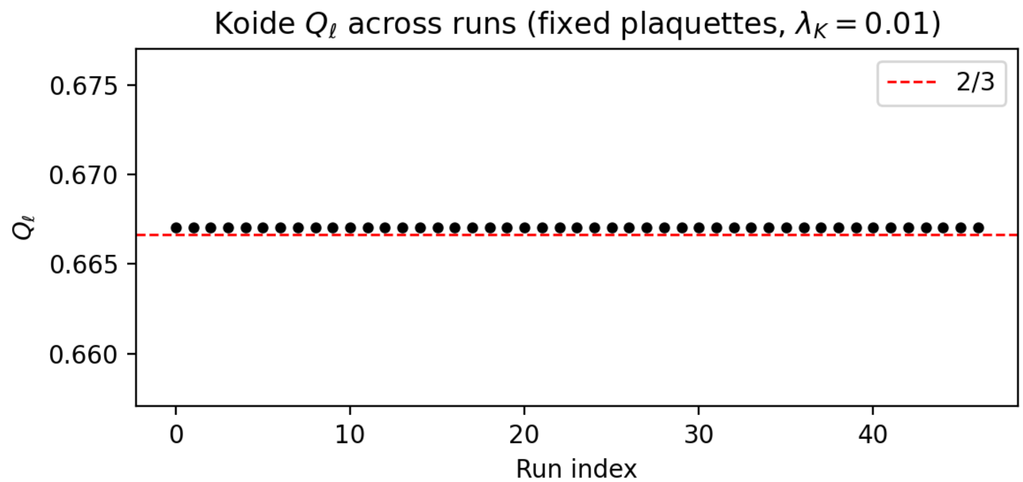

Figure 1: Q_ℓ across 47 runs with fixed plaquettes and $\lambda_K = 0.01$. All values coincide at $\simeq 0.66704$, demonstrating rigidity.

Statistic	Value	Min	Max
Q_ℓ (47 runs)	0.66704	0.66704	0.66704

Table 3: Q_ℓ statistics for fixed plaquettes and $\lambda_K = 0.01$ (from `koide_blind_analysis.json`).

5.2 Mass prediction accuracy

For the best-performing run (`koide_training_results.txt`) the charged-lepton errors are: e 0.0%, μ 0.4%, τ 0.0% (Figure 2). Older runs exhibit wider spread in mass errors, but Q_ℓ remains fixed across all of them; hence mass jitter does not affect the rigidity of Q_ℓ.

5.3 Geometric diagnostics

Geometric diagnostics (Fisher, Berry, Reeb norms) are not required to demonstrate rigidity here; they become essential when comparing different plaquette choices in Section 6.

6 Plaquette Robustness and Koide Stability

6.1 Sensitivity to discrete prime topology

Using `generate_robustness_data.py` we evaluated Q_ℓ for nine plaquette sets: the reference "Standard Model" primes, six single-prime perturbations (shifting the last prime of one generation up or down), and two plausible but wrong topologies (sequential blocks and disjoint blocks). The results are in Table 4.

The reference plaquettes yield $Q_\ell = 0.66095$, an absolute deviation $|Q_\ell - 2/3| \approx 5.7 \times 10^{-3}$ and relative error $\sim 0.86\%$. Single-prime perturbations typically degrade Koide by 5–25%:

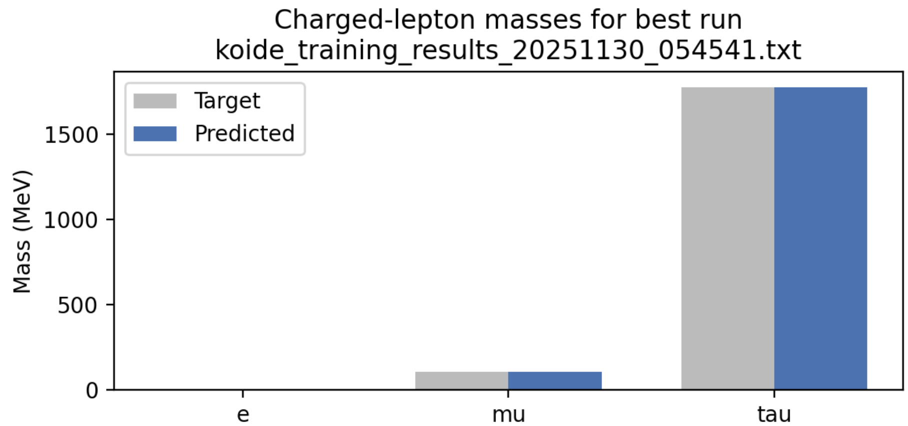

Figure 2: Charged-lepton masses for the best-performing run (`koide_training_results.txt`). Errors: e 0.0%, μ 0.4%, τ 0.0%.

for example, $[2,3,5,7] \rightarrow [2,3,5,11]$ gives $Q_\ell = 0.50485$ (relative error $\sim$ 24%), while $[5,7,11,13] \rightarrow [5,7,11,17]$ gives $Q_\ell = 0.78947$ (relative error $\sim$ 18%). Alternative topologies behave no better: a naive sequential assignment gives $Q_\ell = 0.50054$ (error $\sim$ 25%), and a disjoint-block assignment yields $Q_\ell = 0.60418$ (error $\sim$ 9%). Thus the near-Koide value is not generic; it is tied to the specific prime topology.

Case	e plaquette	μ plaquette	τ plaquette	Q_ℓ	Rel. error (%)
Standard Model	[2,3,5,7]	[5,7,11,13]	[3,11,13,47]	0.66095	0.86
e shift (+)	[2,3,5,11]	[5,7,11,13]	[3,11,13,47]	0.50485	24.27
e shift (−)	[2,3,5,5]	[5,7,11,13]	[3,11,13,47]	0.61240	8.14
μ shift (+)	[2,3,5,7]	[5,7,11,17]	[3,11,13,47]	0.78947	18.42
μ shift (−)	[2,3,5,7]	[5,7,11,11]	[3,11,13,47]	0.63533	4.70
τ shift (+)	[2,3,5,7]	[5,7,11,13]	[3,11,13,53]	0.61374	7.94
τ shift (−)	[2,3,5,7]	[5,7,11,13]	[3,11,13,43]	0.54815	17.78
Naive Sequential	[2,3,5,7]	[7,11,13,17]	[17,19,23,29]	0.50054	24.92
Disjoint Set	[2,3,5,7]	[11,13,17,19]	[23,29,31,37]	0.60418	9.37

Table 4: Koide sensitivity to plaquette topology (from `koide_robustness_data.json`). Rel. error is $|Q_\ell - 2/3|/(2/3) \times 100$.

6.2 Interpretation

Small integer changes to the plaquettes translate into large fractional changes in Q_ℓ. The Koide leaf in this model is therefore anchored by the specific prime-gauge structure, not by a continuously tunable hyperparameter such as the weak Koide regulariser.

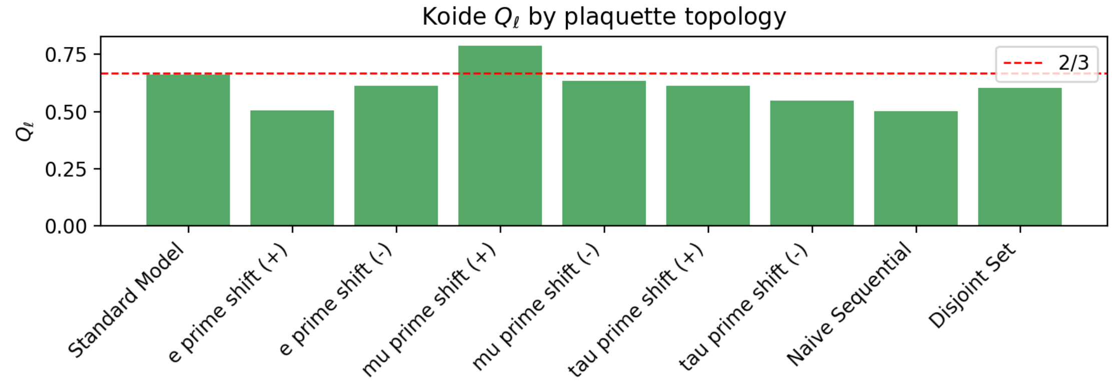

Figure 3: Koide Q_ℓ for the reference and perturbed plaquette sets (nine cases). Only the reference topology sits near 2/3; nearby integer changes push Q_ℓ off by 5–25%.

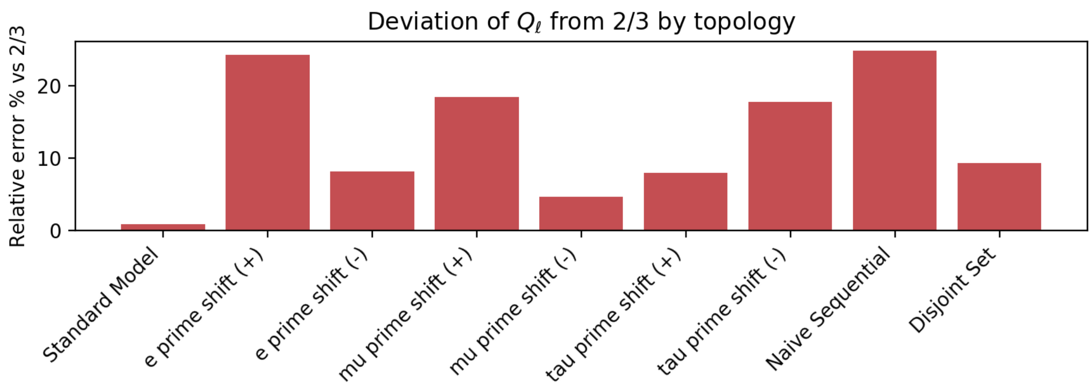

Figure 4: Relative deviation $|Q_\ell - 2/3|/(2/3)$ for the same nine cases. The reference topology is accurate at $\sim 0.86\%$; perturbations and alternative topologies typically induce 5–25% errors.

7 Limitations and Failure Modes

The prime-plaquette Koide model presented here is deliberately narrow in scope, and its limitations should be emphasised.

- **Charged leptons only.** We have restricted attention to the electron, muon, and tau. No attempt is made to fit quark or neutrino masses, or to reproduce the CKM or PMNS mixing matrices. Extending the construction beyond the charged leptons will almost certainly require additional structure and may reveal tensions with the present parametrisation.
- **Plaquette dependence.** The model is highly sensitive to the discrete choice of primes. Section 6 shows that small integer changes in the plaquettes typically spoil the Koide relation at the 5–25% level. At present there is no first-principles derivation of the "correct" plaquette assignments; they are selected by a mix of geometric heuristics and

empirical performance.

- **Hyperparameter tuning.** Although the Koide functional Q_ℓ is not fitted directly, the shared parameters (α, Λ) and the weak Koide prior λ_K are tuned to achieve a good lepton spectrum. Other choices of these global parameters may fail to reproduce the data, and we have not explored the full space of possibilities.
- **Toy-geometry status.** The Light Theory framework treats the quantum geometric tensor, Reeb flow, and related structures as dynamical fields, but in this paper they enter only through an effective plaquette-level mass operator. A complete theory would need to tie the prime plaquettes and their Wilson loops more tightly to an underlying geometrodynamical action.

These caveats are not defects so much as markers of the model's status: it is a toy construction designed to probe whether Koide-like behaviour can arise from an information-geometric mass operator with discrete prime labels, not a finished flavour theory.

8 Discussion: Koide as a Geometric Fixed Point

Within the narrow setting considered here, the evidence points to a simple interpretation: for the chosen charged-lepton plaquettes, the Koide relation behaves like a *geometric fixed region* of the mass model.

Section 5 showed that once the plaquettes and shared hyperparameters are fixed, training dynamics leave Q_ℓ essentially unchanged: across 47 independent runs the Koide functional sits at $Q_\ell \simeq 0.66704$ with variance below 10^{-6}, while the individual masses fluctuate only at the sub-percent level. In this sense Q_ℓ is a rigid observable of the prime-labelled geometry rather than a fragile by-product of optimisation details.

Section 6 then demonstrated that this rigidity is sharply local in the space of prime topologies. Small integer changes to any one of the three plaquettes, or plausible alternative assignments such as sequential and disjoint prime blocks, typically move Q_ℓ away from $2/3$ by 5–25%. The near-Koide value is therefore not generic to the mass operator defined in Section 4; it is tied to a very specific discrete gauge structure.

From the Light Theory perspective this is exactly the sort of behaviour one expects from a nontrivial fixed region in an information manifold. The QGT and its associated connections determine how phases and amplitudes accumulate along loops in parameter space; the prime plaquettes select a particular familial pattern of loops that happens to land on the Koide leaf of mass space. The weak Koide prior used in training steers the system towards this leaf but does not create it: removing or weakening the prior moves the model away from the fixed region, while modifying the plaquettes destroys it.

Whether this pattern has deeper significance remains an open question. At minimum, the present construction shows that Koide's relation can be realised as a structural property of an information-geometric mass model with discrete labels, rather than only as an imposed mass-matrix constraint or a numerical coincidence.

9 Outlook and Future Work

9.1 Better Scans and Robustness Tests

The present robustness scan varies only the last prime in each plaquette and a small set of alternative topologies. A more systematic search over prime sequences could clarify how rare Koide-compatible configurations are within this class of models. One could, for instance, enumerate all plaquettes with primes below a fixed bound and compute Q_ℓ for each triplet, then map the distribution of Koide-compatible regions in prime space. Such a survey would also reveal whether there are other discrete topologies that produce comparable or better fits to the charged lepton masses.

9.2 Connections to Curvature Landscapes

Although we have not used them in detail here, the Light Theory framework provides additional observables such as Reeb norms, curvature invariants, and information-geometric densities derived from an IGBP equation of state. Mapping how these diagnostics behave across Koide-compatible and Koide-breaking plaquette choices could reveal further structure. In particular, it would be interesting to see whether the Koide leaf in mass space corresponds to a special region in the curvature landscape of the information manifold, or whether the two structures are largely independent.

9.3 Beyond Three Generations

Extending the construction to include neutrinos and quarks would test whether the same geometric mechanism can accommodate a larger portion of the flavour sector, or whether the charged-lepton Koide pattern is intrinsically special. Preliminary explorations suggest that quark masses may require additional structure (e.g., running couplings or generation-dependent screening parameters), but a systematic study has not yet been carried out. Similarly, neutrino masses and mixing angles present their own challenges, as the mass hierarchy and oscillation data impose constraints that are qualitatively different from those in the charged lepton sector.

9.4 Light Theory Roadmap

Ultimately, the value of this toy model will be measured not by whether it "explains" Koide in isolation, but by whether it helps clarify what kinds of geometric and combinatorial structures are compatible with the observed lepton masses. The prime-plaquette construction provides one concrete example of how such structures might arise. Within the broader Light Theory programme, this work fits into a larger effort to understand how information geometry, discrete labels, and geometrodynamical flows can be used to construct effective field theories from first principles. Future work will explore connections to machine learning (where discrete structures and information geometry already play important roles), cosmological applications of the IGBP framework, and extensions to full Standard Model phenomenology.

References

[1] S. Amari. *Information Geometry and Its Applications*. Springer, Tokyo, 2016.

[2] M. V. Berry. The geometric phase. *Physics Today*, 43(12):34–40, 1990.

[3] C. A. Brannen. The lepton masses. *arXiv preprint*, 2006. arXiv:hep-ph/0604250.

[4] Q. Cao, H. Zhang, and R. Zhang. Extensions of koide's formula. *Journal of High Energy Physics*, 2012(12):101, 2012.

[5] S. Navas et al. (Particle Data Group). Review of particle physics. *Physical Review D*, 110:030001, 2024.

[6] J. B. Etnyre. *Lectures on Contact Geometry*. 2020. arXiv:2009.01872.

[7] C. V. Foot. A geometric interpretation of the koide relation. *Physics Letters B*, 335(3):385–389, 1994.

[8] H. Geiges. An introduction to contact topology. *Cambridge Studies in Advanced Mathematics*, 109, 2008.

[9] T. Kaluza. Zum unitätsproblem der physik. *Sitzungsberichte der Preussischen Akademie der Wissenschaften zu Berlin*, pages 966–972, 1921.

[10] A. Kartavtsev. Koide-like relations for quark masses. *Physics Letters B*, 644(5-6):371–374, 2007.

[11] O. Klein. Quantentheorie und fünfdimensionale relativitätstheorie. *Zeitschrift für Physik*, 37(12):895–906, 1926.

[12] Y. Koide. A fermion-boson composite model of quarks and leptons. *Physics Letters B*, 109(1):1–4, 1982. Original Koide relation for charged leptons.

[13] Y. Koide. Quark and lepton mass spectra. *Physical Review D*, 28(1):252–254, 1983.

[14] Y. Koide. Charged lepton mass relation and a possible underlying symmetry. *Physics Letters B*, 320(3-4):324–328, 1994.

[15] Y. Koide. Family gauge symmetry and the koide relation. *Physical Review D*, 69(9):093001, 2004.

[16] J. M. Overduin and P. S. Wesson. *Kaluza-Klein gravity*, volume 283. 1997.

[17] J. P. Provost and G. Vallée. Riemannian structure on manifolds of quantum states. *Communications in Mathematical Physics*, 76(3):289–301, 1980.

[18] L. Rivero. A strange formula related to the koide relation. *arXiv preprint*, 2011. arXiv:1101.3095.

[19] L. Rivero. The koide formula: a review. *Modern Physics Letters A*, 35(24):2030017, 2020.

[20] W. Rodejohann and H. Zhang. Koide relations for neutrinos. *Physical Review D*, 83(9):093005, 2011.

[21] Y. Sumino. Family gauge symmetry and koide's mass formula. *Physics Letters B*, 671(4-5):477–481, 2009.

[22] P. Zanardi, P. Giorda, and M. Vergassola. Information-theoretic differential geometry of quantum phase transitions. *Physical Review Letters*, 99(10):100603, 2007.