

Lecture 4: The Riemann Curvature Tensor


To physically motivate the commutator of covariant derivatives, we must translate the macroscopic geometric concept of parallel transport into infinitesimal, localized calculus.

From Macroscopic Loops to Infinitesimal Parallelograms

Recall the previous example of parallel transporting a vector around the macroscopic closed loop on a 2-sphere. The vector returns rotated by an angle $\Delta\theta$. This rotation is the undeniable signature of intrinsic curvature.

To capture this curvature at a specific, localized point p using differential calculus, we shrink this macroscopic loop down to an infinitesimally small parallelogram. Imagine starting at point p with a vector V and parallel transporting it to a diagonally opposite corner q using two different paths:

1. **Path A:** Move infinitesimally in the x^ν direction, and then in the x^μ direction.
2. **Path B:** Move infinitesimally in the x^μ direction, and then in the x^ν direction.



In a perfectly flat space, these two paths form a closed, flat rectangle. The vector you end up with will be identical regardless of whether you took Path A or Path B. However, in a curved manifold, the two resulting vectors (V_A and V_B) will point in slightly different directions.

The Commutator as the Measure of Curvature

The covariant derivative ∇_μ is exactly the operator that calculates the change in a tensor as it is parallel transported. Applying two covariant derivatives in succession tells us what happens when we move along these sequential paths. Because covariant derivatives account for the twisting and bending of the coordinate basis, they **do not commute** on a curved manifold.

The difference between taking Path A and Path B is mathematically captured by taking the commutator of the covariant derivatives:

$$[\nabla_\mu, \nabla_\nu]V^\rho = \nabla_\mu \nabla_\nu V^\rho - \nabla_\nu \nabla_\mu V^\rho \quad (1)$$

This commutator quantifies exactly how much the vector V^ρ has rotated or stretched after completing the infinitesimal closed loop. Because the operation is strictly linear with respect to the initial vector V^ρ , the resulting discrepancy must be proportional to V^σ . The mathematical object that acts as this proportionality machine—eating the two directional vectors of our infinitesimal parallelogram and the original vector itself to spit out the deviation—is the **Riemann Curvature Tensor**:

$$[\nabla_\mu, \nabla_\nu]V^\rho = R^\rho_{\sigma\mu\nu}V^\sigma \quad (2)$$

By explicitly expanding the covariant derivatives using the Christoffel symbols, the Riemann tensor can be written in terms of the connection:

$$R_{\sigma\mu\nu}^{\rho} = \partial_{\mu}\Gamma_{\nu\sigma}^{\rho} - \partial_{\nu}\Gamma_{\mu\sigma}^{\rho} + \Gamma_{\mu\lambda}^{\rho}\Gamma_{\nu\sigma}^{\lambda} - \Gamma_{\nu\lambda}^{\rho}\Gamma_{\mu\sigma}^{\lambda} \quad (3)$$

The Riemann tensor is a massive object (a rank-4 tensor). In 4-dimensional spacetime, it has $4^4 = 256$ components, though geometric symmetries reduce this to 20 independent components. If and only if $R_{\sigma\mu\nu}^{\rho} = 0$ everywhere, the manifold is perfectly flat.

0.1 The Ricci Tensor and Ricci Scalar

While the Riemann tensor contains the complete curvature information of the manifold, it is often useful to construct simpler, contracted tensors for physical applications.

Contracting the first and third indices of the Riemann tensor yields the symmetric **Ricci Tensor**, $R_{\mu\nu}$:

$$R_{\mu\nu} = R_{\mu\lambda\nu}^{\lambda} \quad (4)$$

Contracting the Ricci tensor further with the inverse metric yields a single invariant number at every point, known as the **Ricci Scalar** (or scalar curvature), R :

$$R = g^{\mu\nu} R_{\mu\nu} \quad (5)$$

Concept Check 1: In standard flat Cartesian coordinates (x, y, z) , the covariant derivative ∇ reduces to the standard partial derivative ∂ . Given that partial derivatives commute for smooth functions ($\partial_x \partial_y f = \partial_y \partial_x f$), what must the value of the Riemann Tensor be in flat space?

1 Worked Example: Curvature of a 2-Sphere

Let us apply this curvature machinery to the 2-sphere (S^2) of constant radius R to calculate its Ricci scalar. The metric is $ds^2 = R^2 d\theta^2 + R^2 \sin^2 \theta d\phi^2$. The only non-zero Christoffel symbols are:

$$\Gamma_{\phi\phi}^{\theta} = -\sin \theta \cos \theta, \quad \Gamma_{\theta\phi}^{\phi} = \Gamma_{\phi\theta}^{\phi} = \cot \theta \quad (6)$$

Let's calculate the specific Riemann component $R_{\phi\theta\phi}^{\theta}$:

$$R_{\phi\theta\phi}^{\theta} = \partial_{\theta}\Gamma_{\phi\phi}^{\theta} - \partial_{\phi}\Gamma_{\theta\phi}^{\theta} + \Gamma_{\theta\lambda}^{\theta}\Gamma_{\phi\phi}^{\lambda} - \Gamma_{\phi\lambda}^{\theta}\Gamma_{\theta\phi}^{\lambda} \quad (7)$$

Substituting our known Christoffel symbols:

$$\begin{aligned} \partial_{\theta}\Gamma_{\phi\phi}^{\theta} &= \partial_{\theta}(-\sin \theta \cos \theta) = \sin^2 \theta - \cos^2 \theta \\ \partial_{\phi}\Gamma_{\theta\phi}^{\theta} &= 0 \\ \Gamma_{\theta\lambda}^{\theta}\Gamma_{\phi\phi}^{\lambda} &= 0 \\ \Gamma_{\phi\lambda}^{\theta}\Gamma_{\theta\phi}^{\lambda} &= \Gamma_{\phi\phi}^{\theta}\Gamma_{\theta\phi}^{\phi} = (-\sin \theta \cos \theta)(\cot \theta) = -\cos^2 \theta \end{aligned}$$

Combining these terms yields:

$$R_{\phi\theta\phi}^{\theta} = (\sin^2 \theta - \cos^2 \theta) - 0 + 0 - (-\cos^2 \theta) = \sin^2 \theta \quad (8)$$

To find the Ricci scalar, we must calculate the Ricci tensor components ($R_{\theta\theta}$ and $R_{\phi\phi}$) and contract them with the inverse metric components $g^{\theta\theta} = 1/R^2$ and $g^{\phi\phi} = 1/(R^2 \sin^2 \theta)$. Through symmetry and raising/lowering indices, we find $R_{\phi\phi} = \sin^2 \theta$ and $R_{\theta\theta} = 1$. Evaluating the full trace yields:

$$R = g^{\theta\theta} R_{\theta\theta} + g^{\phi\phi} R_{\phi\phi} = \frac{1}{R^2}(1) + \frac{1}{R^2 \sin^2 \theta}(\sin^2 \theta) = \frac{2}{R^2} \quad (9)$$

This confirms rigorously that a sphere has a strictly positive, constant scalar curvature that is inversely proportional to the square of its radius.

2 Gaussian Curvature and the *Theorema Egregium*

Before tackling the full Riemann tensor in higher dimensions, it is instructive to look at 2-dimensional surfaces. Extrinsicly, if we embed a 2D surface in 3D space, we can measure how much it bends in different directions at any given point. There will always be a direction of maximum bending and a direction of minimum bending. These are called the principal curvatures, denoted κ_1 and κ_2 .

Carl Friedrich Gauss defined the **Gaussian Curvature** K as the product of these two principal curvatures:

$$K = \kappa_1 \kappa_2 \quad (10)$$


Gauss's most famous geometric discovery, which he named the *Theorema Egregium* (the “Remarkable Theorem”), proved that K depends *only* on the intrinsic metric $g_{\mu\nu}$ of the surface, not on how the surface is embedded in 3D space. An observer restricted entirely to the 2D surface can determine K purely by measuring local distances and angles. Mathematically, the Riemann curvature tensor for a 2D manifold has only one independent component, and it is directly proportional to this Gaussian curvature.

Concept Check 2: Consider a cylinder embedded in 3D space. One of its principal curvatures corresponds to a circle around the cylinder, and the other to a straight line along its length. What is the value of the straight-line principal curvature? Consequently, what is the Gaussian curvature K of the cylinder?

3 Mathematical Examples of Intrinsic Curvature

To understand how intrinsic curvature manifests geometrically, we can examine three fundamental mathematical spaces.

- **Positive Curvature (The Sphere S^2):** On a sphere, the surface bends in the same direction along both principal axes, meaning κ_1 and κ_2 share the same sign. Therefore, $K > 0$. In positively curved spaces, geodesics that start parallel will eventually converge. Furthermore, the interior angles of a triangle formed by geodesics sum to strictly *more* than 180° . For instance, a triangle drawn from the North pole to the equator, along the equator for a quarter circumference, and back up to the pole has three 90° angles, summing to 270° .



Positive Curvature ($K > 0$)

- **Zero Curvature (The Plane and the Cylinder):** A flat Euclidean plane clearly has $K = 0$. Surprisingly, a cylinder also has strictly zero intrinsic curvature. While it looks curved from the outside, one of its principal curvatures (the straight line running parallel to the vertical axis) is zero. Thus, the product $K = \kappa_1 \kappa_2 = 0$. You can physically unroll a cylinder into a flat rectangle without stretching, shrinking, or tearing the material. Because the metric is unchanged by unrolling, the Riemann tensor of a cylinder evaluates to zero everywhere.
- **Negative Curvature (The Saddle and the Hyperbolic Plane):** On a saddle-shaped surface (like a hyperbolic paraboloid), the principal curvatures bend in opposite directions (one curves up, one curves down). This makes the product $K < 0$. In negatively curved geometries, initially parallel lines diverge away from one another, and the interior angles of a geodesic triangle sum to strictly *less* than 180° .

4 Sectional Curvature in Higher Dimensions

For manifolds of dimension three or higher, a single scalar value like K is insufficient to capture the full geometry. The Riemann tensor $R^\rho_{\sigma\mu\nu}$ contains all the necessary information, but it can be conceptually broken down into localized 2D slices.

If we pick a point p and choose a 2D plane Π within the tangent space $T_p\mathcal{M}$ (spanned by two linearly independent tangent vectors U and V), we can look at all the geodesics shooting out from p in directions that lie within Π . These geodesics sweep out a 2D subsurface within the larger manifold.

The intrinsic Gaussian curvature of this specific 2D subsurface at p is called the **Sectional Curvature**, denoted $K(U, V)$. It is computed directly from the Riemann tensor:

$$K(U, V) = \frac{R_{\rho\sigma\mu\nu}U^\rho V^\sigma U^\mu V^\nu}{g(U, U)g(V, V) - [g(U, V)]^2} \quad (11)$$

By calculating the sectional curvature for all possible 2D planes passing through the tangent space at p , we completely determine the full Riemann tensor.

Concept Check 3: The denominator in the formula for Sectional Curvature, $g(U, U)g(V, V) - [g(U, V)]^2$, looks very similar to a familiar concept from linear algebra. What geometric property of the two vectors U and V does this denominator represent the square of? (Hint: Think about areas spanned by vectors.)

5 The Gauss-Bonnet Theorem: Linking Geometry to Topology

One of the most profound mathematical results in differential geometry connects local curvature to the global, overarching topology of the manifold. For any compact, two-dimensional Riemannian

manifold $\mathcal{M}$ without a boundary, the total integral of the Gaussian curvature over the entire surface is a topological constant:

$$\int_{\mathcal{M}} K dA = 2\pi\chi(\mathcal{M}) \quad (12)$$

Here, χ is the Euler characteristic, a topological integer defined as $2 - 2g$, where g (the genus) is the number of “holes” in the surface.

This theorem reveals a striking rigidity: no matter how you stretch, deform, or dent a sphere (which has zero holes, so $\chi = 2$), the total integrated curvature will always exactly equal 4π . If you create a region of intense positive curvature by pinching a section into a sharp peak, the surrounding area must flatten out or develop negative curvature to perfectly balance the integral. Geometry can be altered locally, but the total curvature is locked by the manifold’s topology.

Concept Check 4: A torus (donut) has exactly one hole ($g = 1$). According to the Gauss-Bonnet theorem, what is the value of the total integrated curvature $\int K dA$ for a torus? If a torus is shaped so that the “outside” part has positive Gaussian curvature, what must be true about the curvature on the “inside” part (near the hole)?

The Lie derivative is a tool to compare tensors across different tangent spaces, leading directly to the study of continuous symmetries (isometries) and Killing vectors.

6 Comparing Tensors: A Gentle Introduction to the Lie Derivative

On a curved manifold, a fundamental problem arises when we try to evaluate how a tensor field (like a vector or a metric) changes from one point to another. The standard definition of a derivative requires us to subtract the value of the field at a point p from its value at a nearby point q .

However, a vector $U(p)$ lives in the tangent space $T_p\mathcal{M}$, while $U(q)$ lives in $T_q\mathcal{M}$. Because these are entirely different vector spaces, direct mathematical subtraction is meaningless—it is like trying to add a matrix to a polynomial.

To compare tensors at different points without relying on any external structure, we need a natural way to “drag” a tensor from q back to p . The **Lie derivative** achieves this by using the *flow* generated by a second vector field V .

Imagine the vector field V as a fluid flowing over our manifold. If you drop a test particle at point p , the fluid carries it along a path known as an integral curve. The collection of all these curves generates a smooth mapping of the manifold onto itself, often denoted ϕ_t , representing the flow parameterized by a “time” t .

If we have another tensor field T , we can evaluate T at the downstream point $q = \phi_t(p)$. Using the mathematical mapping of the flow, we can “pull back” the tensor $T(q)$ along the flow lines of V against the current, bringing it into the original tangent space $T_p\mathcal{M}$. Now that both the pulled-back tensor and the original tensor $T(p)$ reside in the *same* tangent space, we can finally subtract them.

The Lie derivative of T along V , denoted $\mathcal{L}_V T$, is formally the limit of this difference as $t \rightarrow 0$.

6.1 Coordinate Expressions

For a scalar function f , the Lie derivative is simply the directional derivative:

$$\mathcal{L}_V f = V^\mu \partial_\mu f \quad (13)$$

For a vector field U , dragging it along the flow of V yields the *Lie bracket* or commutator of the two vector fields. In coordinates:

$$(\mathcal{L}_V U)^\mu = [V, U]^\mu = V^\nu \partial_\nu U^\mu - U^\nu \partial_\nu V^\mu \quad (14)$$

When we apply the Lie derivative to a $(0, 2)$ -covariant tensor, such as the metric tensor $g_{\mu\nu}$, the product rule generates additional terms depending on the derivatives of V . The formula in local coordinates is:

$$\mathcal{L}_V g_{\mu\nu} = V^\sigma \partial_\sigma g_{\mu\nu} + g_{\sigma\nu} \partial_\mu V^\sigma + g_{\mu\sigma} \partial_\nu V^\sigma \quad (15)$$

This specific expression tells us precisely how the metric (and thus the geometry itself) changes as we slide the manifold along the vector field V . This brings us directly to the concept of symmetries.

7 Isometries and Killing Vectors

When analyzing manifolds, we are intensely interested in symmetries. Symmetries vastly simplify solving the geodesic equations because they generate conserved quantities. A continuous symmetry of the metric geometry is termed an **isometry**.

Using the language we just established, an isometry is a transformation along a vector field K that leaves the metric completely unchanged. In other words, if we drag the metric along the flow of K , we see no difference. Mathematically, this requires the Lie derivative of the metric along the vector field K to strictly vanish:

$$\mathcal{L}_K g_{\mu\nu} = 0 \quad (16)$$

Using the coordinate expression for the Lie derivative derived in the previous section, this condition is:

$$K^\sigma \partial_\sigma g_{\mu\nu} + g_{\sigma\nu} \partial_\mu K^\sigma + g_{\mu\sigma} \partial_\nu K^\sigma = 0 \quad (17)$$

While this equation is correct, it is often more elegantly written using covariant derivatives ∇_μ instead of partial derivatives ∂_μ . By rewriting the partials and grouping terms, this requirement simplifies beautifully into **Killing's Equation**:

$$\nabla_\mu K_\nu + \nabla_\nu K_\mu = 0 \quad (18)$$

Any vector field K^μ satisfying this equation is a **Killing vector**. It points in a direction along which the intrinsic geometry of the manifold does not change.

7.1 Conservation Laws

The power of Killing vectors lies in Noether's theorem for geometry: continuous symmetries generate conservation laws. If $u^\mu = \frac{dx^\mu}{d\lambda}$ is the tangent vector to an affinely parameterized geodesic, and K^μ is a Killing vector, their inner product is a conserved quantity along that geodesic:

$$Q = K_\mu u^\mu = \text{constant} \quad (19)$$

To prove this rigorously, we differentiate Q with respect to the affine parameter λ :

$$\frac{dQ}{d\lambda} = u^\nu \nabla_\nu (K_\mu u^\mu) = u^\nu u^\mu \nabla_\nu K_\mu + K_\mu (u^\nu \nabla_\nu u^\mu) \quad (20)$$

We analyze the two terms on the right-hand side separately:

1. The second term vanishes because $u^\nu \nabla_\nu u^\mu = 0$ is simply the geodesic equation written in covariant derivative notation (zero proper acceleration).
2. The first term, $u^\nu u^\mu \nabla_\nu K_\mu$, vanishes due to symmetry. The product $u^\nu u^\mu$ is perfectly symmetric under the exchange of μ and ν . However, Killing's equation ($\nabla_\mu K_\nu = -\nabla_\nu K_\mu$) dictates that $\nabla_\nu K_\mu$ is perfectly antisymmetric. The complete contraction of a symmetric tensor with an antisymmetric tensor is identically zero.

Thus, we conclude:

$$\frac{dQ}{d\lambda} = 0 \quad (21)$$

7.2 Example: Killing Vectors on a 2-Sphere

To see Killing's equation in action, let us explicitly find the Killing vectors for a 2-sphere (S^2) of radius R . The line element in standard spherical coordinates $(x^1, x^2) = (\theta, \phi)$ is:

$$ds^2 = R^2 d\theta^2 + R^2 \sin^2 \theta d\phi^2 \quad (22)$$

The non-zero metric components are $g_{\theta\theta} = R^2$ and $g_{\phi\phi} = R^2 \sin^2 \theta$.

First, we require the non-zero Christoffel symbols. Using the definition, one readily finds:

$$\Gamma_{\phi\phi}^\theta = -\sin \theta \cos \theta, \quad \Gamma_{\theta\phi}^\phi = \Gamma_{\phi\theta}^\phi = \cot \theta \quad (23)$$

We expand Killing's equation, $\nabla_\mu K_\nu + \nabla_\nu K_\mu = 0$, which evaluates to:

$$\partial_\mu K_\nu + \partial_\nu K_\mu - 2\Gamma_{\mu\nu}^\lambda K_\lambda = 0 \quad (24)$$

This gives us a system of three coupled partial differential equations for the covariant components K_θ and K_ϕ :

1. The $\theta\theta$ -component ($\mu = \theta, \nu = \theta$):

$$2\partial_\theta K_\theta - 2(0) = 0 \implies \partial_\theta K_\theta = 0 \quad (25)$$

This implies K_θ is a function of ϕ only: $K_\theta = f(\phi)$.

2. The $\phi\phi$ -component ($\mu = \phi, \nu = \phi$):

$$2\partial_\phi K_\phi - 2\Gamma_{\phi\phi}^\theta K_\theta = 0 \implies \partial_\phi K_\phi + \sin \theta \cos \theta f(\phi) = 0 \quad (26)$$

Integrating with respect to ϕ gives:

$$K_\phi = -\sin \theta \cos \theta \int f(\phi) d\phi + g(\theta) \equiv -\sin \theta \cos \theta F(\phi) + g(\theta) \quad (27)$$

where $F'(\phi) = f(\phi)$ and $g(\theta)$ is an arbitrary function of θ .

3. The $\theta\phi$ -component ($\mu = \theta, \nu = \phi$):

$$\partial_\theta K_\phi + \partial_\phi K_\theta - 2\Gamma_{\theta\phi}^\phi K_\phi = 0 \quad (28)$$

Substituting our expressions for K_θ and K_ϕ :

$$\partial_\theta [-\sin \theta \cos \theta F(\phi) + g(\theta)] + f'(\phi) - 2 \cot \theta [-\sin \theta \cos \theta F(\phi) + g(\theta)] = 0 \quad (29)$$

Applying the θ derivative, $\partial_\theta(-\sin \theta \cos \theta) = \sin^2 \theta - \cos^2 \theta$, and simplifying:

$$(\sin^2 \theta - \cos^2 \theta)F(\phi) + g'(\theta) + f'(\phi) + 2 \cos^2 \theta F(\phi) - 2 \cot \theta g(\theta) = 0 \quad (30)$$

$$F(\phi) + f'(\phi) + g'(\theta) - 2 \cot \theta g(\theta) = 0 \quad (31)$$

Because the first two terms depend only on ϕ and the final two depend only on θ , they must separately equal constants C and $-C$:

$$F(\phi) + F''(\phi) = C \quad \text{and} \quad g'(\theta) - 2 \cot \theta g(\theta) = -C \quad (32)$$

The ϕ equation is a simple harmonic oscillator with a constant driving force, yielding:

$$F(\phi) = A \cos \phi + B \sin \phi + C \implies f(\phi) = -A \sin \phi + B \cos \phi \quad (33)$$

Solving the θ equation using an integrating factor of $\csc^2 \theta$ yields:

$$g(\theta) = C \sin \theta \cos \theta + D \sin^2 \theta \quad (34)$$

Substituting $f(\phi)$, $F(\phi)$, and $g(\theta)$ back into our covariant components K_μ , and then raising the indices using $K^\mu = g^{\mu\nu} K_\nu$, we find the contravariant Killing vectors. Absorbing the overall $1/R^2$ scaling factor into our arbitrary constants A, B , and D , we isolate the three independent vector fields:

$$K_{(1)} = \cos \phi \partial_\theta - \cot \theta \sin \phi \partial_\phi \quad (35)$$

$$K_{(2)} = \sin \phi \partial_\theta + \cot \theta \cos \phi \partial_\phi \quad (36)$$

$$K_{(3)} = \partial_\phi \quad (37)$$

These three vectors correspond to the three generators of the $SO(3)$ rotational symmetry group, representing rotations about the Cartesian x , y , and z axes respectively.