Geometric aspects of ODEs and PDEs

Ari Stern
astern@math.ucsd.edu

CSME Seminar, UCSD
April 9, 2009

Overview

- Symplectic geometry is fundamental to the study of Hamiltonian mechanics.
- How can we generalize symplectic geometry from ODEs to PDEs?
- *Problem:* Time is “special” in Hamiltonian mechanics, so this presents a conceptual obstacle to studying relativistic/covariant field theories.
- Instead, begin by developing symplectic geometry on the *Lagrangian* side, from the perspective of variational principles. Then generalize to the Lagrangian densities for field theoretic PDEs.
- The result is *multisymplectic geometry*, whose key property is the vanishing of a certain n-dimensional boundary integral around regions of $(n+1)$-dimensional spacetime. When $n = 0$, this is just ordinary symplectic geometry, and the vanishing of the boundary integral says that the symplectic form is equal at the endpoints of any time interval.
- Other key results, such as Noether’s theorem, can be generalized naturally.
- *Not on today’s agenda:* the entire framework can be discretized, yielding structure-preserving numerical integrators for both ODEs and PDEs.

Lagrangian Mechanics

Let Q be a smooth manifold, called the configuration space, and let TQ be its tangent bundle, called the phase space. The *Lagrangian* is a function $L\colon TQ \to \mathbb{R}$. For a path $q\colon [0,T] \to Q$, with initial time 0 and final time T, define the *action functional*

$$S[q] = \int_0^T L\left(q(t), \dot{q}(t)\right)\, \mathrm{d}t.$$

Suppose we wish to find the trajectory $q(t)$ whose values at the initial and final time are given to be $q(0) = q_0$ and $q(T) = q_T$. *Hamilton's principle of stationary action* says that this trajectory must satisfy $\delta S[q] = \mathbf{d}S[q] \cdot \delta q = 0$, where δq is any variation of the path that preserves the initial- and final-time conditions, i.e., $\delta q(0) = \delta q(T) = 0$:

$$\begin{aligned} 0 = \mathbf{d}S[q] \cdot \delta q &= \int_0^T \left[\frac{\partial L}{\partial q}(q, \dot{q}) \cdot \delta q + \frac{\partial L}{\partial \dot{q}}(q, \dot{q}) \cdot \delta \dot{q} \right] \mathrm{d}t \\ &= \int_0^T \left[\frac{\partial L}{\partial q}(q, \dot{q}) - \frac{\mathrm{d}}{\mathrm{d}t}\frac{\partial L}{\partial \dot{q}}(q, \dot{q}) \right] \cdot \delta q \, \mathrm{d}t. \end{aligned}$$

This holds when $q(t)$ is a solution to the *Euler–Lagrange equations*

$$\boxed{\frac{\partial L}{\partial q}(q, \dot{q}) - \frac{\mathrm{d}}{\mathrm{d}t}\frac{\partial L}{\partial \dot{q}}(q, \dot{q}) = 0.}$$

Euler–Lagrange Flows are Symplectic

Let us define the one-parameter flow map $F_t \colon TQ \to TQ$, taking $(q_0, \dot{q}_0) \mapsto (q(t), \dot{q}(t))$, where q is the solution to the Euler–Lagrange equations with initial conditions $(q(0), \dot{q}(0)) = (q_0, \dot{q}_0)$. Then define the *restricted action function* $S_{TQ} \colon TQ \to \mathbb{R}$, using the flow map to pull the Lagrangian back to the space of initial conditions

$$S_{TQ}(q_0, \dot{q}_0) = \int_0^T F_t^* L(q_0, \dot{q}_0)\, \mathrm{d}t = \int_0^T L(q, \dot{q})\, \mathrm{d}t.$$

Now, let us vary with respect to the initial conditions, along variations δq_0 and $\delta \dot{q}_0$

$$\begin{aligned}
\mathrm{d}S_{TQ}(q_0, \dot{q}_0) \cdot (\delta q_0, \delta \dot{q}_0) &= \int_0^T \left[\frac{\partial L}{\partial q}(q, \dot{q}) \cdot \delta q + \frac{\partial L}{\partial \dot{q}}(q, \dot{q}) \cdot \delta \dot{q} \right] \mathrm{d}t \\
&= \int_0^T \left[\frac{\partial L}{\partial q}(q, \dot{q}) - \frac{\mathrm{d}}{\mathrm{d}t} \frac{\partial L}{\partial \dot{q}}(q, \dot{q}) \right] \cdot \delta q \, dt + \left[\frac{\partial L}{\partial \dot{q}}(q, \dot{q}) \cdot \delta q \right]_0^T \\
&= \left[\frac{\partial L}{\partial \dot{q}}(q, \dot{q}) \cdot \delta q \right]_0^T.
\end{aligned}$$

(Since these variations do not have the fixed-endpoint property as before, we pick up a boundary term when integrating by parts in step 2. Moreover, the nonboundary term vanishes in step 3, since q is a solution of the Euler–Lagrange equations.)

Euler–Lagrange Flows are Symplectic (continued)

Now, let us define the *canonical* 1*-form* $\theta_L \in \Omega^1(TQ)$ such that $\theta_L \cdot \delta q = \frac{\partial L}{\partial \dot{q}} \cdot \delta q$, or in coordinates, $\theta_L = \frac{\partial L}{\partial \dot{q}^i}\, \mathrm{d}q^i$. Then the previous equation can be rewritten as

$$\mathrm{d}S_{TQ}\,(q_0, \dot{q}_0) \cdot (\delta q_0, \delta \dot{q}_0) = \left[\frac{\partial L}{\partial \dot{q}}\,(q, \dot{q}) \cdot \delta q \right]_0^T = \left[\theta_L \cdot \delta q \right]_0^T,$$

or simply

$$\mathrm{d}S_{TQ} = F_T^* \theta_L - \theta_L.$$

Taking the exterior derivative of both sides and noting that $\mathrm{d}\mathrm{d}S_{TQ} = 0$, we get

$$F_T^*\,(\mathrm{d}\theta_L) = \mathrm{d}\theta_L.$$

Therefore, defining the *symplectic* 2*-form* $\omega_L = -\mathrm{d}\theta_L = \mathrm{d}q^i \wedge \mathrm{d}\frac{\partial L}{\partial \dot{q}^i}$, it follows that the Euler–Lagrange flow map preserves ω_L and hence is a symplectic flow.

Remark. Symplecticity is typically discussed in the context of Hamiltonian flows on T^*Q, where one defines the canonical 1-form $\theta_H = p_i\, \mathrm{d}q^i$ and symplectic 2-form $\omega_H = -\mathrm{d}\theta_H = \mathrm{d}q^i \wedge \mathrm{d}p_i$. This is equivalent to the Lagrangian formulation with respect to the Legendre transformation $TQ \to T^*Q$, which takes $(q, \dot{q}) \mapsto (q, p) = \left(q, \frac{\partial L}{\partial \dot{q}}\,(q, \dot{q})\right)$.

Noether's Theorem and Momentum Maps

Suppose that a Lie group G with Lie algebra $\mathfrak{g}$ acts on Q, and that the Lagrangian is invariant with respect to this G-action. Then, if we take some element $\xi \in \mathfrak{g}$ with infinitesimal generator $\xi_{TQ} \in TQ$, this invariance says that

$$\mathrm{d}S_{TQ}\left(q_0, \dot{q}_0\right) \cdot \xi_{TQ} = 0.$$

However, we have previously shown that $\mathrm{d}S_{TQ} = F_T^* \theta_L - \theta_L$, so it follows that

$$F_T^*\left(\theta_L \cdot \xi_{TQ}\right) = \theta_L \cdot \xi_{TQ},$$

and so the quantity $\theta_L \cdot \xi_{TQ}$ is conserved by the Euler–Lagrange flow.

Therefore, we can define a conserved *momentum map* $J \colon TQ \to \mathfrak{g}^*$, given by

$$\left\langle J\left(q, \dot{q}\right), \xi \right\rangle = \theta_L\left(q, \dot{q}\right) \cdot \xi_{TQ}.$$

This fundamental result is known as *Noether's theorem*.

The Laplace–Beltrami Operator

Let X be a smooth, orientable $(n+m)$-dimensional pseudo-Riemannian manifold, with metric signature (n,m). Let $\Omega^k(X)$ denote the space of differential k-forms on X, where $\mathrm{d}\colon \Omega^k(X) \to \Omega^{k+1}(X)$ is the exterior derivative and $*\colon \Omega^k(X) \to \Omega^{n+m-k}(X)$ is the Hodge star associated to the metric on X.

Next, we define the codifferential operator $\delta\colon \Omega^{k+1}(X) \to \Omega^k(X)$ to be

$$\delta = (-1)^{(n+m)k+1+m} * \mathrm{d} * .$$

Finally, we define the *Laplace–Beltrami operator*, applied to scalar functions on X, to be

$$\delta\mathrm{d}\colon \Omega^0(X) \to \Omega^0(X).$$

One can also apply the Laplace–Beltrami operator to k-forms for general k; another generalization is the *Laplace–de Rham operator*

$$(\delta\mathrm{d} + \mathrm{d}\delta) : \Omega^k(X) \to \Omega^k(X),$$

which reduces to Laplace–Beltrami when $k = 0$.

Laplace–Beltrami as an Elliptic Operator

Suppose that X is an n-dimensional Riemannian manifold, and let $u \in \Omega^0(X)$ be a scalar function on X. If $X = \mathbb{R}^n$ with the Euclidean metric, then the Laplace–Beltrami operator $\delta \mathrm{d}$ is precisely the usual Laplacian Δ, up to a sign. To see this, observe that in standard coordinates

$$*\mathrm{d}*\mathrm{d}u = *\mathrm{d}*\left(\partial_i u \, \mathrm{d}x^i\right) = *\mathrm{d}\left(\partial^i u \, \mathrm{d}^{n-1}x_i\right) = *\left(\partial_i \partial^i u \, \mathrm{d}^n x\right) = \partial_i \partial^i u = \Delta u.$$

More generally, suppose that X has the Riemannian metric g. Then in local coordinates, we have $g = g_{ij} \, \mathrm{d}x^i \otimes \mathrm{d}x^j$, and

$$\begin{aligned} *\mathrm{d}*\mathrm{d}u &= *\mathrm{d}*\left(\partial_i u \, \mathrm{d}x^i\right) \\ &= *\mathrm{d}\left(\sqrt{|g|}\partial^i u \, \mathrm{d}^{n-1}x_i\right) \\ &= *\left[\partial_i\left(\sqrt{|g|}\partial^i u\right) \mathrm{d}^n x\right] \\ &= \frac{1}{\sqrt{|g|}}\partial_i\left(\sqrt{|g|}\partial^i u\right). \end{aligned}$$

Therefore, a general elliptic operator—traditionally written as $\nabla \cdot a \nabla$ for some function a—can be expressed as the Laplace–Beltrami operator for a suitable metric.

Poisson's Equation with Differential Forms

One of the most fundamental elliptic PDEs is *Poisson's equation*

$$\Delta u = f,$$

which is called *Laplace's equation* in the special case $f = 0$. We can use the Laplace–Beltrami operator to translate this problem in terms of differential forms on a Riemannian manifold X.

Given some scalar function $f \in \Omega^0(X)$, we wish to find another scalar function $u \in \Omega^0(X)$ such that

$$\delta \mathrm{d} u = f.$$

This is the second-order formulation of Poisson's equation. One can also introduce an auxiliary field $v \in \Omega^1(X)$, and then rewrite this as a first-order system of equations

$$\mathrm{d} u = v, \qquad \delta v = f.$$

Alternatively, this can be written as

$$\begin{pmatrix} 0 & \delta \\ \mathrm{d} & 0 \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} f \\ v \end{pmatrix}.$$

“Reduced” Form of Poisson’s Equation

Now, suppose that u is simply a scalar potential for the 1-form $v = \mathrm{d}u$, and we only care about solving for the values of v. In this case, we can eliminate the explicit dependence on u by observing that $\mathrm{d}v = \mathrm{d}\mathrm{d}u = 0$, and therefore we can simply solve the system

$$\mathrm{d}v = 0, \qquad \delta v = f.$$

This is equivalent to finding a vector field $v^\sharp \in \mathfrak{X}(X)$, which is curl-free and has divergence f. If, later, one needs to reconstruct u from a solution for v, then this can be done (at least locally) by using the Poincaré lemma.

Laplace–Beltrami as a Hyperbolic Operator

Suppose now that X is an $(n+1)$-dimensional Lorentzian manifold, i.e., a pseudo-Riemannian manifold with metric signature $(n,1)$. In this case, since the metric is positive-definite in n dimensions and negative-definite in 1 dimension, the Laplace–Beltrami operator $\delta\mathrm{d}$ becomes a *hyperbolic operator*. In particular, the negative term arises when taking the first Hodge star, in order to raise the index $\partial^i = g^{ij}\partial_j$.

Consider the simple $(1+1)$-dimensional example of $X = \mathbb{R}^{1,1}$, where in coordinates the metric is given by

$$g = \mathrm{d}x \otimes \mathrm{d}x - \mathrm{d}t \otimes \mathrm{d}t.$$

Then, if $u \in \Omega^0(X)$, we have

$$\begin{aligned} *\mathrm{d}*\mathrm{d}u &= *\mathrm{d}*\left(u_x\,\mathrm{d}x + u_t\,\mathrm{d}t\right) \\ &= *\mathrm{d}\left(u_x\,\mathrm{d}t + u_t\,\mathrm{d}x\right) \\ &= *\left(u_{xx}\,\mathrm{d}x \wedge \mathrm{d}t + u_{tt}\,\mathrm{d}t \wedge \mathrm{d}x\right) \\ &= u_{xx} - u_{tt} \\ &= \Box u, \end{aligned}$$

where $\Box$ denotes the *d'Alembertian wave operator*.

The Wave and Transport Equations

The *wave equation* is the hyperbolic PDE

$$\Box u = 0.$$

In $\mathbb{R}^{1,1}$, it is common in the PDE literature to “factor” the wave operator as

$$\Box u = \left(\frac{\partial^2}{\partial x^2} - \frac{\partial^2}{\partial t^2} \right) u = \left(\frac{\partial}{\partial x} - \frac{\partial}{\partial t} \right) \left(\frac{\partial}{\partial x} + \frac{\partial}{\partial t} \right) u,$$

so by substituting $v = \left(\frac{\partial}{\partial x} + \frac{\partial}{\partial t} \right) u$, one can simply study the first-order equation

$$v_t = v_x,$$

which is also called the *transport equation*.

The Wave and Transport Equations (continued)

On a general Lorentzian manifold, the wave equation is simply written as

$$\delta \mathrm{d} u = 0,$$

which is formally identical to Laplace's equation. Again, the wave speed is determined by the metric, which manifests in the Hodge star operator. Just as before, we can rewrite this as the first-order system

$$\mathrm{d} u = v, \qquad \delta v = 0,$$

or simply in terms of $v \in \Omega^1(X)$ as

$$\mathrm{d} v = 0, \qquad \delta v = 0.$$

This latter formulation can be seen as a more *geometric* version of the transport equation.

Furthermore, because the "factorization" $\square = \delta \circ \mathrm{d}$ works for any n, not just $n = 1$, this can be seen as a coordinate-free generalization of transport and conservation laws for any number of spatial dimensions.

Lagrangian Density for Laplace–Beltrami

As before, let X be an orientable pseudo-Riemannian manifold, with $u \in \Omega^0(X)$ a scalar field, and consider the Lagrangian density

$$\mathcal{L}(u, \mathrm{d}u) = -\frac{1}{2}\mathrm{d}u \wedge *\mathrm{d}u.$$

This expression is closely related to the L^2 inner product $(\cdot, \cdot) : \Omega^k(X) \times \Omega^k(X) \to \mathbb{R}$ on differential k-forms, which is defined by

$$(\alpha, \beta) = \int_X \alpha \wedge *\beta.$$

An important property of this inner product is that the codifferential δ is dual to the exterior derivative d (hence its name), corresponding to integration by parts, assuming that the boundary terms vanish. Specifically, let $\alpha \in \Omega^k(X), \beta \in \Omega^{k+1}(X)$, and then by the Leibniz rule,

$$\mathrm{d}(\alpha \wedge *\beta) = \mathrm{d}\alpha \wedge *\beta - \alpha \wedge *\delta\beta.$$

Integrating over X and applying Stokes' theorem,

$$\int_{\partial X} \alpha \wedge *\beta = (\mathrm{d}\alpha, \beta) - (\alpha, \delta\beta),$$

so when the left-hand side vanishes, we have $(\mathrm{d}\alpha, \beta) = (\alpha, \delta\beta)$.

Hamilton's Principle and Euler–Lagrange Equations

Therefore, the action associated to the Lagrangian density $\mathcal{L}$ is

$$S[u] = \int_X \mathcal{L} = -\frac{1}{2}(\mathrm{d}u, \mathrm{d}u).$$

Let $\tilde{u} \in \Omega^0(X)$ be a variation of u preserving boundary conditions, so that $\tilde{u}|_{\partial X} = 0$. Then the variation of the action along $\tilde{u}$ is

$$\mathbf{d}S[u] \cdot \tilde{u} = -(\mathrm{d}\tilde{u}, \mathrm{d}u) = -(\tilde{u}, \delta \mathrm{d}u).$$

Setting this variation equal to zero, we get the Euler–Lagrange equation

$$\boxed{\boldsymbol{\delta \mathrm{d} u = 0,}}$$

which coincides with Laplace's equation or the wave equation, respectively, when X is Riemannian or Lorentzian.

Remark: Two Approaches for Nonhomogeneous Systems

More generally, a nonhomogeneous linear system, such as Poisson's equation, can be obtained by adding a term to the Lagrangian density

$$\mathcal{L} = -\frac{1}{2}\mathrm{d}u \wedge *\mathrm{d}u + u \wedge *f,$$

for some $f \in \Omega^0(X)$, *where f does not depend on u.*

If f *does* depend on u (i.e., the system is nonlinear) then one must take a slightly different approach: keeping the Lagrangian density as before, $\mathcal{L} = -\frac{1}{2}\mathrm{d}u \wedge *\mathrm{d}u$, one can use the modified variational principle

$$\mathbf{d}S[u] \cdot \tilde{u} + (\tilde{u}, f) = 0,$$

which generalizes the Lagrange–d'Alembert principle from mechanics.

Multisymplectic Geometry

Returning to Hamilton's action principle, suppose now that we restrict u to the space of Euler–Lagrange solutions, while taking a variation η that no longer vanishes the boundary ∂X. In this case, if we vary the restricted action along η, we get only the boundary integral term

$$\mathbf{d}S[u] \cdot \eta = \int_{\partial X} \eta \wedge *\mathrm{d}u.$$

If we vary this yet again, along another variation v, this becomes

$$\mathbf{dd}S[u] \cdot \eta \cdot v = \int_{\partial X} \frac{1}{2} \left(\eta \wedge *\mathrm{d}v - v \wedge *\mathrm{d}\eta \right).$$

The left-hand side vanishes, since $\mathbf{dd} \equiv 0$, resulting in the identity

$$\int_{\partial X} \left(\eta \wedge *\mathrm{d}v - v \wedge *\mathrm{d}\eta \right) = 0$$

for all η and v.

Multisymplectic Geometry (continued)

Let us now rephrase these statements in the language of multisymplectic geometry. Take the canonical Cartan 1-form, associated to the Lagrangian density $\mathcal{L}$, to be $\theta_{\mathcal{L}} = \mathbf{d}u \wedge *\mathrm{d}u$, so that varying the restricted action gives

$$\mathbf{d}S[u] \cdot \eta = \int_{\partial X} \theta_{\mathcal{L}} \cdot \eta.$$

Then the multisymplectic 2-form is $\omega_{\mathcal{L}} = -\mathbf{d}\theta_{\mathcal{L}} = \mathbf{d}u \wedge \mathbf{d}*\mathrm{d}u$, and thus

$$\mathbf{dd}S[u] \cdot \eta \cdot \nu = \int_{\partial X} \omega_{\mathcal{L}} \cdot \eta \cdot \nu = 0.$$

This last equation is called the *multisymplectic form formula*.

Returning to the special case of mechanics, where $X = [0, T] \Rightarrow \partial X = \{0, T\}$, this reduces to the statement that the (multi)symplectic 2-form is equal at the two endpoints.

The Lagrangian structure of Maxwell's equations

Take a spacetime manifold X, with a Hodge star operator $*$ corresponding to a Lorentzian metric, and define a source 3-form $\mathcal{J}$ satisfying the continuity of charge condition $\mathrm{d}\mathcal{J} = 0$.

Then, given a potential 1-form A, the Lagrangian density is the 4-form

$$\boldsymbol{\mathcal{L} = -\frac{1}{2}\mathrm{d}A \wedge *\mathrm{d}A + A \wedge \mathcal{J},}$$

which has the action integral $S[A] = \int_X \mathcal{L}$.

Applying Hamilton's principle of stationary action, for all variations α of A vanishing on ∂X, we have the weak formulation

$$0 = \mathbf{d}S[A] \cdot \alpha = \int_X (-\mathrm{d}\alpha \wedge *\mathrm{d}A + \alpha \wedge \mathcal{J}) = \int_X \alpha \wedge (-\mathrm{d}{*}\mathrm{d}A + \mathcal{J})\,,$$

which gives the Euler–Lagrange equations $\mathrm{d}{*}\mathrm{d}A = \mathcal{J}$. Finally, substituting $F = dA$ and noting $\mathrm{d}F = \mathrm{dd}A = 0$, we arrive at Maxwell's equations:

$$\boxed{\mathbf{d}{*}\boldsymbol{F} = \boldsymbol{\mathcal{J}}, \qquad \mathbf{d}\boldsymbol{F} = \mathbf{0}.}$$

Gauge symmetry and constraints in Maxwell's equations

Maxwell's equations have 6 dynamical components (involving $\partial/\partial t$), and 2 constraints,

$$\nabla \cdot \vec{B} = 0, \qquad \nabla \cdot \vec{D} = \rho,$$

involving only spatial derivatives. Why are these automatically preserved by the time flow?

Answer: $\nabla \cdot \vec{B} = 0$ is the easy one, since it is taken care of by $\mathrm{d}F = \mathrm{dd}A = 0$, which corresponds to $\nabla \cdot \vec{B} = \nabla \cdot \nabla \times \vec{A} = 0$. The second is more subtle: it turns out that $\nabla \cdot \vec{D} - \rho$ is the *conserved momentum* associated to a gauge symmetry.

Maxwell's equations have a gauge symmetry with respect to transformations $A \mapsto A + \mathrm{d}f$, since $\mathrm{dd}f = 0$ for any scalar function f. If we choose a time coordinate, we can choose the gauge to be the *Weyl gauge*, so that the scalar potential (the t-component of A) is $\phi = 0$. This eliminates $\nabla \cdot \vec{D} = \rho$ from the Euler–Lagrange equations, but $\nabla \cdot \vec{D} - \rho$ is a conserved momentum for the remaining gauge symmetry $\vec{A} \mapsto \vec{A} + \nabla f$.

Conclusion

- In Lagrangian mechanics, symplectic geometry arises through the *boundary terms* for variations of the action integral.
- Symplecticity of Euler–Lagrange flows is an immediate consequence of $\mathrm{dd} = 0$.
- For Lagrangian PDEs, this can be extended to *multisymplectic geometry* on $(n + 1)$-dimensional spacetime, where the boundary terms become an n-dimensional boundary integral.
- For Euler–Lagrange solutions, $\mathrm{dd} = 0$ implies the *multisymplectic form formula*, which requires a boundary integral to vanish. This is the higher-dimensional, covariant analog of a "symplectic flow."