

The 19th Problem: Part III

De Giorgi's Theorem

A Leisurely Excursion

Introduction

Let's recap where we are. In Part I, we solved Hilbert's 19th problem in two dimensions. We had two methods: the "chopping argument," which used the energy cost of oscillation on boundary circles, and the quasiconformal argument, which forced the gradient map to have bounded distortion. Both gave us the same results of oscillation decay, and therefore Hölder continuity of ∇u . We also discussed that the chopping argument fails in three and higher dimensions because the energy integral converges, so the gradient can "afford" to keep oscillating across planes forever without violating the Caccioppoli bound. The quasiconformal argument fails because one scalar equation is too weak, as there's too much algebraic slack.

In Part II, due to the failure in dimensions greater than two, we changed our approach from a geometrical argument to a more analytical argument. To that end, we built an analytical toolkit (weak derivatives, Sobolev spaces, the Sobolev inequality, the Caccioppoli inequality) and by the end of Part II, we had bootstrapped our way to $\nabla u \in L^\infty$. The gradient is bounded.

But bounded is not the finish line. Recall the vicious circle from Part I: to prove the solution is smooth, the classical theory (Schauder estimates) demands that the coefficients $a_{ij}(x) = F_{ij}(\nabla u(x))$ be smooth. And those coefficients are the gradient we're trying to prove smooth. The way to break this circle is to get just a *little* regularity for free, without assuming anything about the coefficients. If we can show $\nabla u \in C^\alpha$ for some $\alpha > 0$, then Schauder takes over and bootstraps us to C^∞ , and when F is analytic, we eventually hope to get C^ω .

So the entire proof in n dimensions reduces to one question: can we prove that ∇u is Hölder continuous, without assuming anything about the smoothness of the coefficients?

The answer is yes. The tool that does it is De Giorgi's Theorem, published in 1957. Let's us dive right into this.

De Giorgi's Theorem

Before we get to De Giorgi's theorem, let's set up the problem. We start with our Euler–Lagrange equation from Part I:

$$\operatorname{div}(\nabla F(\nabla u)) = F_{ij}(\nabla u) u_{ij} = 0, \quad (1)$$

where F is analytic, convex, and uniformly convex (meaning D^2F has eigenvalues between λ and λ^{-1}).

Here's a natural idea we already saw in Part I: differentiate the whole equation with respect to x_k . If u is smooth¹, we can move ∂_k inside and get

$$\partial_i(F_{ij}(\nabla u) \partial_j(\partial_k u)) = 0, \quad k = 1, \dots, n. \quad (2)$$

Look at what just happened. Each component $v = \partial_k u$ of the gradient satisfies a *linear* equation in divergence form:

$$\partial_i(a_{ij}(x) \partial_j v) = 0, \quad (3)$$

where $a_{ij}(x) = F_{ij}(\nabla u(x))$. Because F is uniformly convex, a_{ij} has eigenvalues in $[\lambda, \lambda^{-1}]$. The equation for v is *uniformly elliptic*.

But we don't yet know whether ∇u is continuous. So, the coefficients $a_{ij}(x) = F_{ij}(\nabla u(x))$ might be discontinuous. We only know they're *bounded* and *measurable*. So, the entire problem boils down to:

If v solves $\partial_i(a_{ij}(x) \partial_j v) = 0$ where the eigenvalues of a_{ij} lie in $[\lambda, \lambda^{-1}]$, then

$$\|v\|_{C^\alpha(B_{1/2})} \leq C(n, \lambda) \|v\|_{L^\infty(B_1)}$$

for some $\alpha(n, \lambda) > 0$.

This is De Giorgi's theorem. Once we have it, $v = \partial_k u \in C^\alpha$, so $\nabla u \in C^\alpha$, and the vicious circle is broken. Let's prove it.

The $L^2 \rightarrow L^\infty$ Estimate

Before stating the first key result, we need a definition.

A *sub-solution* is a function v satisfying $\partial_i(a_{ij} \partial_j v) \geq 0$ in the weak sense.

What does this mean intuitively?

Recall that a solution satisfies $\partial_i(a_{ij} \partial_j v) = 0$, which is a generalization of $\Delta v = 0$. Replacing the equality with ≥ 0 is like saying the function is allowed to curve upward more than a solution would, but never more downward. In the special case of Laplace's equation, $\Delta v \geq 0$

¹To emphasize ideas, we'll assume u is smooth and derive estimates that don't depend on this assumption. These are called *a priori* ("before the fact") estimates. The rigorous extension to weak solutions uses standard approximation arguments, but the ideas are the same.

means v at any point is bounded above by its average on surrounding balls. Sub-solutions are the same idea, but for our rougher equation.

The first ingredient of De Giorgi's proof says something intuitive: sub-solutions can't have interior upward spikes. If a sub-solution is small in an L^2 sense, it must be small pointwise in the interior.

Theorem 1 ($L^2 \rightarrow L^\infty$ Estimate). *Assume that $\partial_i(a_{ij}(x)\partial_j v) \geq 0$ in $B_2 \subset \mathbb{R}^n$. Then*

$$\sup_{B_1} v \leq C(n, \lambda) \|v_+\|_{L^2(B_2)},$$

where $v_+ = \max\{v, 0\}$.

So the pointwise max on the interior ball is controlled by the L^2 size on the bigger ball. *Why should this be true?* We already know two things from Part II:

1. **Caccioppoli** says the gradient energy on a smaller ball is controlled by the function's L^2 norm on a larger ball:

$$\int_{B_{r/2}} |\nabla v|^2 dx \leq \frac{C}{r^2} \int_{B_r} v^2 dx.$$

It says the equation doesn't allow the solution to oscillate too wildly.

2. **Sobolev** says having $\nabla v \in L^2$ forces v into a strictly better L^p space:

$$\|v\|_{L^{2^*}} \leq C \|v\|_{W^{1,2}}, \quad \text{where } 2^* = \frac{2n}{n-2} > 2.$$

Now here's the key insight. Caccioppoli converts L^2 information about v into L^2 information about ∇v . Sobolev then converts that into L^{2^*} information about v , where $2^* > 2$. So composing the two, we go from $v \in L^2$ to $v \in L^{2^*}$. We've *gained integrability*. And since $v \in L^{2^*}$ is just another L^p space, nothing stops us from repeating the argument, but now starting from a better exponent.

Each trip around this loop, the exponent goes up. After enough iterations, the exponent is so large that L^p control effectively gives L^∞ control². This is the essence of the $L^2 \rightarrow L^\infty$ estimate: a *feedback loop* between the Caccioppoli and Sobolev that keeps tightening until the function is trapped.

The proof

Let's make this precise. For any $\kappa \in \mathbb{R}$, define the truncation $v_\kappa := (v - \kappa)_+ = \max(v - \kappa, 0)$. This just chops off everything below the height κ and shifts the rest down. Since v is a sub-solution and the map $t \mapsto \max(t - \kappa, 0)$ is convex and increasing, v_κ is also a sub-solution. So, Caccioppoli applies to v_κ .

²Recall that $\|v\|_{L^p} \rightarrow \|v\|_{L^\infty}$ as $p \rightarrow \infty$. So once the exponent is large enough, we basically have pointwise control.

Using Caccioppoli on v_κ (multiplied by a *cutoff function* ψ , a smooth function that equals 1 on a smaller ball and smoothly drops to 0 outside a larger ball³) and then applying Sobolev, we get after some work with Hölder's inequality:

$$\int v_\kappa^2 \psi^2 \leq C(n, \lambda) \left(\int v_\kappa^2 |\nabla \psi|^2 \right) |\{v_\kappa \psi^2 > 0\}|^{2/n}.$$

Now let $0 \leq \tau < \kappa \leq 1$. Choose ψ to be a standard cutoff: 1 on $B_{2-\kappa}$, zero outside $B_{2-\tau}$. Define

$$V(s) := \int_{B_{2-s}} v_s^2 dx.$$

Substituting, we arrive at the key recursive inequality:

$$V(\kappa) \leq \frac{C}{(\kappa - \tau)^{2+4/n}} V(\tau)^{1+2/n}. \quad (4)$$

Notice that exponent on the right is strictly greater than 1. This is the whole game.

If the power were exactly 1, we'd just get $V(\kappa) \leq CV(\tau)$ that tells us nothing. But because the power is *superlinear*, small inputs get crushed. If V is already small, then $V^{1+2/n}$ is *much* smaller than V . Each iteration shrinks V faster than the last. If $V(0)$ starts sufficiently small (depending on n and λ), iterating forces $V(1) = 0$.

$V(1) = 0$ means $v_1 = (v - 1)_+ = 0$ in B_1 , i.e., $v \leq 1$ in B_1 . And since we can always normalize $\|v_+\|_{L^2(B_2)}$ to be small (just divide v by a large constant and the equation is linear), we get

$$\sup_{B_1} v \leq C(n, \lambda) \|v_+\|_{L^2(B_2)}.$$

Oscillation Decay

The $L^2 \rightarrow L^\infty$ estimate tells us sub-solutions can't spike. That's useful, but we need the *oscillation* of a solution to shrink as we zoom in. This shrinking is what gives Hölder continuity.

In Part I, we got oscillation decay from the chopping argument. Here we'll get it by a completely different method, one that works in every dimension.

Proposition 2 (Oscillation Decay). *Assume $\partial_i(a_{ij}(x)\partial_j v) \geq 0$ in $B_1 \subset \mathbb{R}^n$. For any $\delta > 0$, there exists $\epsilon(n, \lambda, \delta) > 0$ such that if*

$$|\{v_+ = 0\} \cap B_1| \geq \delta |B_1|,$$

then

$$\sup_{B_{1/2}} v \leq (1 - \epsilon) \sup_{B_1} v_+.$$

³Why do we need cutoffs? We want to localize our estimates to a specific region. The cutoff ψ acts like a smooth "window" that lets us focus on one ball without worrying about boundary effects.

What does this say? If a sub-solution is zero (or negative) on a decent chunk of the ball (at least a fraction δ of the volume), then its maximum *must drop* when we step inside to $B_{1/2}$. It can't stay close to its boundary maximum throughout the interior. The interior has to separate.

This is a *quantitative* maximum principle. The classical version, which we used in Part I, says sub-solutions can't have interior maxima. Proposition 2 is stronger: they can't even get *close* to the maximum deep inside, and the gap depends only on n , λ , and δ , not on any smoothness of the coefficients.

Imagine a sub-solution v in B_1 with $\sup_{B_1} v_+ = 1$, and suppose the set where $v \leq 0$ takes up at least a δ -fraction of the ball. The function needs to climb from 0 to nearly 1 somewhere. But H^1 functions can't do this for free; they have to "pay" in gradient energy for every level set they cross. And Caccioppoli caps the total energy budget. So, if the function starts at zero on a δ -fraction of the ball, it simply can't afford to reach the full height deep in the interior. Something has to give, and what gives is the maximum.

The proof

We'll show a slightly weaker version to keep things clean. Assume v is a sub-solution in B_2 with $\sup_{B_2} v_+ = 1$ and $|\{v_+ = 0\} \cap B_1| \geq \delta|B_1|$. We want $\sup_{B_{1/2}} v \leq 1 - \epsilon$.

Define the distribution function

$$W(s) := \frac{|\{v \leq s\} \cap B_1|}{|B_1|}, \quad s \in [0, 1].$$

This W is nondecreasing. Our hypothesis says $W(0) \geq \delta > 0$, and we want to show it reaches 1 (or close to it) well before s reaches 1.

For $0 \leq s < t \leq 1$, define

$$w = \frac{(v - s)_+}{1 - s},$$

so $0 \leq w \leq 1$. Let $\bar{w} = \min\{w, (t - s)/(1 - s)\}$, which is w capped at a certain height.

Now we hit \bar{w} with two inequalities pulling in opposite directions:

Since $w \leq 1$ is a sub-solution, Caccioppoli bounds $\int |\nabla w|^2$, and Cauchy–Schwarz gives

$$\int_{B_1} |\nabla \bar{w}| \leq C(n, \lambda) (W(t) - W(s))^{1/2}.$$

The function \bar{w} equals 0 on $\{v \leq s\}$ and $\frac{t-s}{1-s}$ on $\{v \geq t\}$. It's stuck at two different heights on two sets of definite measure. To get between them, \bar{w} needs gradient:

$$\int_{B_1} |\nabla \bar{w}| \geq c(n) W(s) (1 - W(t)) \frac{t - s}{1 - s}.$$

Combining:

$$W(s) \left(1 + c(n, \lambda) \frac{(t - s)^2}{(1 - s)^2} W(s) (1 - W(t))^2 \right) \leq W(t).$$

Now set $b_k = W(1 - 2^{-k})$ and $a_k = 1 - b_k$. The inequality becomes

$$a_{k+1} \leq a_k - c a_k^2,$$

where $c > 0$ depends on n, λ, δ . This is a nice recurrence. A direct induction⁴ gives $a_k \leq \frac{1}{1+ck} \rightarrow 0$, so $W(1 - 2^{-k}) \rightarrow 1$.

In particular, there's some $\epsilon(n, \lambda, \delta) > 0$ such that the sub-solution $\bar{v} = \frac{[v - (1-2\epsilon)]_+}{2\epsilon}$ has L^2 norm small enough that Theorem 1 gives $\bar{v} \leq 1/2$ in $B_{1/2}$. Unwinding, $v \leq 1 - \epsilon$ in $B_{1/2}$. The maximum has dropped.

Hölder Continuity

We're almost there. Here's how we go from oscillation decay to full Hölder continuity.

If v solves $\partial_i(a_{ij}\partial_j v) = 0$, then both v and $-v$ are sub-solutions (since flipping the sign flips the inequality). So we can apply Proposition 2 to both. After shifting and rescaling so that $\inf_{B_1} v = -1$ and $\sup_{B_1} v = 1$, at least one of $\{v \leq 0\}$ or $\{v \geq 0\}$ has measure $\geq |B_1|/2^5$. If $\{v \leq 0\}$ is big, apply Proposition 2 to v : the sup drops to $1 - \epsilon$ on $B_{1/2}$. If $\{v \geq 0\}$ is big, apply it to $-v$: the inf rises to $-1 + \epsilon$ on $B_{1/2}$. Either way:

$$\text{osc}_{B_{1/2}} v \leq (1 - \epsilon/2) \text{osc}_{B_1} v.$$

Now, the beautiful thing is that this decay works at *every scale*. If v solves the equation in a ball of radius r , we can rescale by defining $\tilde{v}(x) = v(rx)$, which solves the same type of equation in B_1 (with rescaled coefficients that still have the same ellipticity bounds). So the oscillation decay applies to any ball, not just B_1 :

$$\text{osc}_{B_{r/2}} v \leq (1 - \epsilon) \text{osc}_{B_r} v.$$

Iterating k times:

$$\text{osc}_{B_{2^{-k}}} v \leq (1 - \epsilon)^k \text{osc}_{B_1} v.$$

Setting $r = 2^{-k}$:

$$\text{osc}_{B_r} v \leq r^\alpha \text{osc}_{B_1} v, \quad \text{where } \alpha = \frac{\log(1/(1 - \epsilon))}{\log 2} > 0. \quad (5)$$

Compare this to Part I: there we had the same kind of estimate, $\omega(r) \leq Cr^\alpha$. We get the same conclusion (oscillation decays as a power of the radius) but by a completely different route. In Part I it came from the geometry of 2D. Here it comes from the interplay of Caccioppoli and Sobolev, which works in *every* dimension.

This power-law decay is exactly Hölder continuity:

$$\|v\|_{C^\alpha(B_{1/2})} \leq C(n, \lambda) \|v\|_{L^\infty(B_1)}.$$

That's De Giorgi's theorem. The key estimate is proven.

⁴If $0 \leq a_{k+1} \leq a_k - c a_k^2$ with $a_0 \leq 1$, then $a_k \leq \frac{1}{1+ck}$. The idea: $a_k \geq \frac{1}{1+ck}$ implies $a_k^2 \geq a_k/(1+ck)$, so the drop ca_k^2 is at least $ca_k/(1+ck)$, which is exactly the decrease we need.

⁵Pigeonhole: the two sets cover all of B_1 , so at least one has half the measure.

Putting It Together

We've now proved all three steps of De Giorgi's theorem: the $L^2 \rightarrow L^\infty$ estimate, the oscillation decay, and Hölder continuity. So how does the full proof of Hilbert's problem come together? Here's the complete chain:

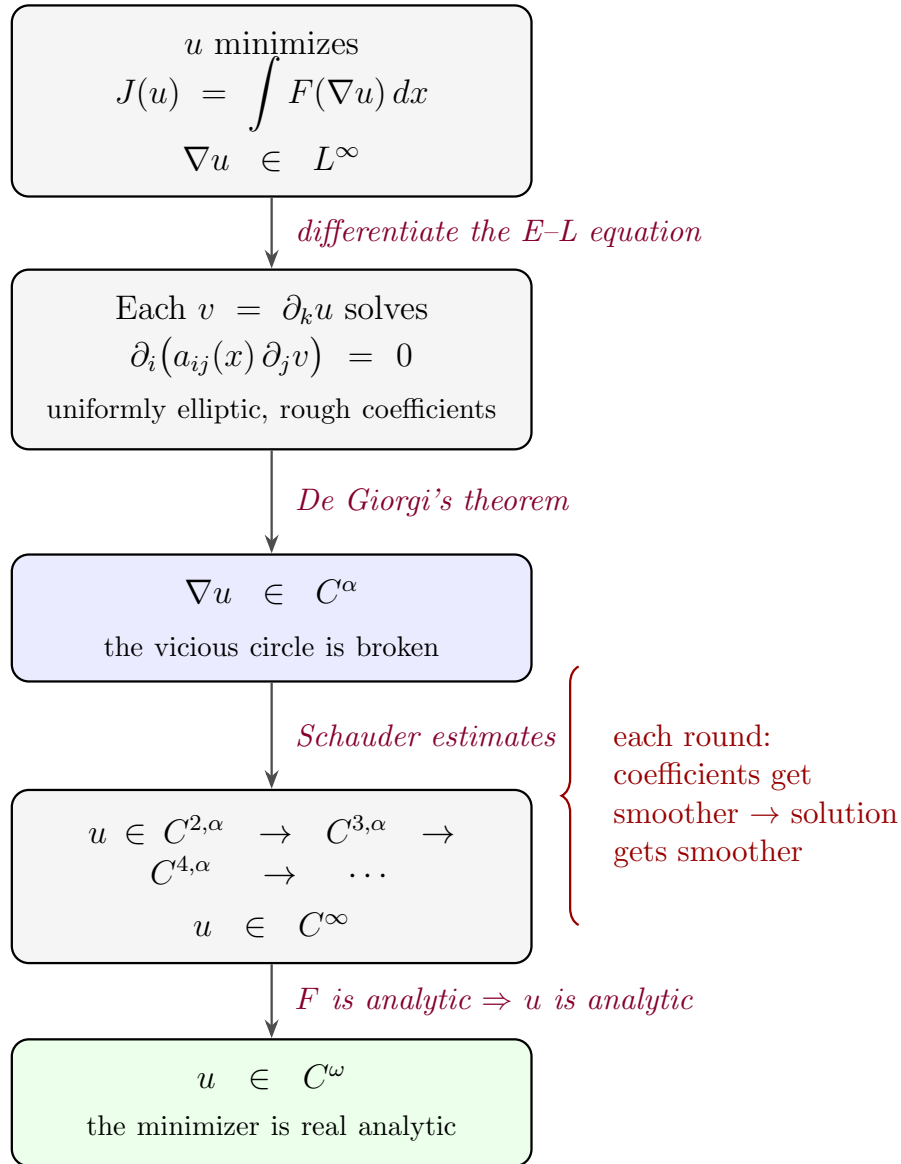


Figure 1: The full proof of Hilbert's 19th problem in n dimensions.

The only steps we haven't discussed in detail are the last two. The Schauder estimates⁶ work like this: once De Giorgi gives us $\nabla u \in C^\alpha$, the coefficients $F_{ij}(\nabla u)$ are C^α too, so Schauder gives $u \in C^{2,\alpha}$. But now $\nabla u \in C^{1,\alpha}$, so the coefficients are $C^{1,\alpha}$. Schauder again: $u \in C^{3,\alpha}$. Each round we gain a derivative, and this cycle never stops. Since $u \in C^{k,\alpha}$

⁶Roughly, solutions to elliptic equations are "two derivatives more regular" than their coefficients. C^α coefficients give $C^{2,\alpha}$ solutions.

for every k , we get $u \in C^\infty$. The final step from C^∞ to C^ω uses the fact that F itself is analytic: one can carefully estimate the growth of higher derivatives $\|D^k u\|_{L^\infty} \leq C^k k!$, which is exactly the condition that guarantees the Taylor series of u converges. And that's it. Hilbert's 19th problem is solved, in all dimensions.