

Harm Minimization Under Autonomous Driving

by [Nick Clark](#) | Published April 25, 2026

Why Harm Ordering Cannot Be Left to Implementation

When autonomous driving must choose between consequences (collision with vehicle vs collision with pedestrian, lane departure vs hard braking, evasive steering vs maintain lane), the choice operates against an implicit harm-ordering. Implementation-level harm-ordering is socially-contested and regulatorily-fragile.

Germany's 2017 Ethics Commission report on autonomous driving explicitly required that harm-minimization rules be transparent and externalized rather than embedded in implementation. Most current AV stacks have not externalized them.

Policy-Configurable Ordering

Harm-minimization deviation as architectural primitive supports policy-declared harm-ordering. The ordering is signed by a policy authority (regulatory body, jurisdictional authority, deployment operator); execution proceeds against the declared ordering; deviation from the declared ordering produces credentialed deviation events.

Different jurisdictions with different ethical frameworks (Germany, France, U.S. federal vs state, China, Japan) declare different harm orderings. The architecture

supports the jurisdictional variation structurally.

Post-Incident Audit

Audit reconstruction traverses: which harm-ordering was active, what observations the actuation operated against, what the resulting harm-minimization deviation was, what alternative would the prior orderings have produced. The reconstruction operates against architectural records.

Liability investigation, regulatory review, and ethical analysis all proceed against architecturally-supported records rather than reconstructed implementation traces.