HYPERSPHERE HEALTH, INC.
Working Paper · May 2026 · For Review and Comment

The Architecture Problem

Thirty Years of Clinical Evidence That Healthcare Data Was Never Built for the Person — and the Architecture That Finally Is

David D. Nethaway, MBA, CSM
Principal, Hypersphere Health, Inc. · Naya Advisory Services, Inc.
Patent-Pending HPDM™ Architecture · Provisional Filed May 5, 2026

> This paper synthesizes existing peer-reviewed literature, federal agency documentation, and published population health data. It does not introduce new clinical claims. It organizes what has already been documented into a coherent argument about why a specific architectural problem persists — and what solving it requires.

Abstract

Healthcare's most persistent failures are not mysteries. Preventable admissions, avoidable readmissions, undertreated complexity, fragmented care — these outcomes have been documented, quantified, and cited in peer-reviewed literature for three decades. The clinicians seeing them have been right all along. The data systems built to support them have not been.

Not because the technology failed. Because the architecture underneath every tool, every platform, and every analytics solution deployed in the last thirty years was built around a billing code, not a person.

This paper traces the documented history of that gap — from the Institute of Medicine's foundational quality reports through current Medicare Advantage payment accuracy data — and argues that closing it requires not a better algorithm running on the same architecture, but a different architecture entirely. One built around the person as the organizing unit of healthcare data.

The Hyper-Spherical Patient Data Model (HPDM™) is a patent-pending architecture designed around that premise. It appears in the final section of this paper — not as a product pitch, but as the logical architectural conclusion the preceding evidence demands.

Section 1 — What Clinicians Have Been Saying

1.1 The IOM's Documented Indictment

In 2001, the Institute of Medicine published Crossing the Quality Chasm, one of the most cited documents in the history of American healthcare policy. Its central finding was not that the healthcare system lacked good intentions, skilled clinicians, or capable technology. Its finding was structural: the system was fundamentally not designed to deliver the care that clinical science had made possible.

The IOM identified six dimensions of quality failure — safety, effectiveness, patient-centeredness, timeliness, efficiency, and equity — and documented that failures across all six were not random or isolated. They were systematic. They arose from the same root cause: a delivery system designed around institutional and financial units rather than around the patient as a continuous, whole person moving through time and across care settings.

> "The American health care delivery system is in need of fundamental change... The current system cannot do the job... trying harder will not work. Changing systems of care will." — Institute of Medicine, Crossing the Quality Chasm, 2001

A decade later, the IOM's 2012 report Best Care at Lower Cost documented that the conditions described in 2001 had not fundamentally changed. The report estimated that the United States wasted approximately $750 billion annually on care that failed to improve outcomes — including $130 billion in unnecessary services, $190 billion in inefficient care delivery, and $55 billion in missed prevention opportunities. The underlying diagnosis was the same: fragmentation, lack of care coordination, and data systems that could not support the clinical knowledge that existed.

These are not the findings of discontented critics. They are the documented conclusions of the nation's most authoritative healthcare quality body, affirmed across two decades of evidence.

1.2 The Population Cost Reality

The Agency for Healthcare Research and Quality's Medical Expenditure Panel Survey has consistently documented one of the most important and underappreciated facts in American healthcare: a very small fraction of patients generates a disproportionate share of total cost.

Top 1% of patients	**~21% of total health expenditure**	AHRQ MEPS — consistent finding across multiple survey years

Top 5% of patients	**~50% of total health expenditure**	AHRQ MEPS — persistent concentration of cost in high-complexity population

Bottom 50% of patients	**<3% of total health expenditure**	AHRQ MEPS — inverse concentration confirms structural driver, not random variation

This concentration is not random. It reflects a population of patients with multiple chronic conditions, complex social circumstances, and clinical trajectories that interact in ways that standard care models — and standard data systems — are not designed to see. The AHRQ has documented that patients in the top cost decile are disproportionately older, have higher rates of multiple chronic conditions, and have greater social determinant burden than patients in lower cost deciles.

Critically: this population is not unpredictable. Its members are identifiable in advance — if the data architecture can see them whole.

1.3 Preventable Admissions and Avoidable Readmissions

The AHRQ's Prevention Quality Indicators (PQI) measure the rate of hospitalizations that clinical evidence suggests could have been prevented with appropriate outpatient care. Across conditions including congestive heart failure, chronic obstructive pulmonary disease, diabetes, and hypertension, the PQI data consistently documents that a substantial proportion of hospitalizations are potentially preventable.

The Centers for Medicare and Medicaid Services' Hospital Readmissions Reduction Program (HRRP) was established in 2012 precisely because readmission rates — documented at 15-20% within 30 days for high-complexity patients — represented a measurable and persistent quality failure. CMS has published readmission rate data across conditions and hospitals for over a decade. The data shows improvement in some areas and persistent challenge in others, particularly for patients with multiple comorbidities.

What the HRRP data does not show is why readmissions persist despite a decade of focused attention and financial penalties. The answer, consistently documented in the implementation literature, is care coordination failure — specifically, the inability to maintain a coherent clinical picture of a high-complexity patient across care settings, across time, and across the multiple providers who touch them.

1.4 Social Determinants — Documented but Data-Invisible

The World Health Organization's Commission on Social Determinants of Health (2008) documented that 30 to 55 percent of health outcomes are attributable to social and environmental factors — housing stability, food security, transportation access, educational attainment, community safety, and social connection. The Centers for Disease Control and Prevention, Healthy People 2030, and the National Academies of Sciences, Engineering, and Medicine have all affirmed and refined this evidence base.

The clinical implication is direct: a patient's social circumstances compound, modify, and sometimes override the clinical interventions applied to their conditions. A diabetic patient who cannot afford insulin behaves differently in a risk model than a diabetic patient who can. A heart failure patient who lives alone, on the third floor of a building without an elevator, has a different readmission trajectory than one with family support and ground-floor access.

These facts are not contested. They are documented in peer-reviewed literature across every major clinical specialty. They are absent from the data architecture that most healthcare organizations use to make care management decisions.

> The documentation gap is not clinical ignorance. Clinicians have known this for decades. The documentation gap is architectural: the systems built to support care management were not designed to ingest, organize, or act on social determinant data alongside clinical data — for the same patient, at the same time, with the same analytical rigor.

Section 2 — Why the Data Never Caught Up

2.1 The Billing Code as the Organizing Unit

The International Classification of Diseases (ICD) system was developed as a mortality classification system. Its evolution into the primary organizing unit of clinical data in the United States was an administrative decision, not a clinical one. ICD codes are the currency of the claims system — the mechanism by which services are documented, submitted, adjudicated, and paid.

They were not designed to represent a patient's clinical reality. A patient with diabetes, chronic kidney disease, heart failure, depression, and housing instability does not appear in the claims system as a whole person with compounding, interacting conditions. They appear as a set of codes — each independently valued, each independently processed, each assigned a risk weight as if the others did not exist.

The Health Information Technology for Economic and Clinical Health (HITECH) Act of 2009 drove rapid adoption of electronic health records across the United States, with meaningful use incentive payments accelerating EHR penetration to near-universal levels by 2015. This was widely anticipated to solve the data fragmentation problem. It did not.

EHR systems captured more data than paper records. But they organized that data around the same unit — the encounter, the code, the transaction — rather than around the patient as a continuous, longitudinal entity. Data produced in one EHR system does not communicate naturally with data in another. A patient's complete clinical picture remains fragmented across the systems of every institution that has touched them, each holding a partial view.

2.2 The HCC Risk Adjustment Model — Designed for Payment, Not Prediction

The Centers for Medicare and Medicaid Services' Hierarchical Condition Category (HCC) model was developed as a risk adjustment mechanism for Medicare Advantage payment. Its explicit purpose, documented in CMS's technical specifications, is to predict relative healthcare expenditure across a population — not to identify individual patients who need clinical intervention.

The HCC model assigns risk scores additively. Each condition is identified from claims data, mapped to a hierarchical condition category, and assigned a risk weight. A patient's total risk score is the sum of these weights. This methodology is appropriate for its designed purpose: actuarial population-level risk adjustment for plan payment.

It is structurally inappropriate as a clinical intelligence tool. The additive model treats each condition as if it exists in isolation. A patient with diabetes, chronic kidney disease, and heart failure does not carry three independent risks. Clinical literature consistently documents that these conditions interact multiplicatively — each compounding the severity, progression, and treatment complexity of the others. An additive model, by design, cannot capture this interaction. It was never intended to.

The Medicare Payment Advisory Commission has documented in multiple annual reports that Medicare Advantage plans systematically underestimate the cost of their highest-complexity members — precisely the population where additive scoring fails most severely. The MedPAC analyses show persistent gaps between HCC-predicted costs and actual costs for members with multiple chronic conditions, particularly those with comorbidity profiles that interact at the clinical level.

> MedPAC has consistently documented that the HCC model underpredicts costs for the sickest beneficiaries — particularly those with complex comorbidity combinations that the additive architecture cannot resolve into a defensible risk estimate. — Medicare Payment Advisory Commission, March Reports (multiple years)

2.3 The Comorbidity Interaction Problem — Documented in Clinical Literature

The clinical literature on comorbidity interactions is extensive and consistent. Conditions do not co-exist independently. They interact at the physiological, pharmacological, and care management levels in ways that make the combined burden substantially greater than the sum of individual burdens.

Cardiorenal syndrome — the bidirectional interaction between cardiac and renal dysfunction — is documented as producing worse outcomes for patients with both conditions than would be predicted by either condition alone. Diabetes with chronic kidney disease produces accelerated renal progression beyond what either condition drives independently. COPD with cardiovascular disease creates respiratory and circulatory interactions that complicate both medication management and care escalation decisions. Depression complicating any chronic condition consistently worsens adherence, self-management, and clinical outcomes across the literature.

These interaction effects are not obscure research findings. They are documented in standard clinical textbooks, major journal reviews, and specialty society guidelines. Every clinician treating patients with multiple chronic conditions is aware of them. No standard risk scoring architecture incorporates them.

> The gap between what clinicians know about comorbidity interactions and what risk scoring systems can model is not a knowledge problem. It is a math problem. Additive models cannot be patched to capture multiplicative interactions. The architecture must change.

Section 3 — The Cost of the Gap

3.1 What Preventable Cost Looks Like at Scale

The Office of Inspector General of the Department of Health and Human Services has published multiple reviews of Medicare Advantage payment accuracy. These reviews consistently document patterns of both underpayment and overpayment attributable to the limitations of claims-based risk adjustment — with high-complexity patients frequently underscored and lower-complexity patients occasionally overscored.

The IOM's 2012 estimate of $750 billion in annual healthcare waste — approximately 30 percent of total US healthcare spending at the time — included components directly attributable to the data architecture gap: unnecessary services ordered because clinicians lacked visibility into prior workup, inefficient care delivery caused by fragmentation across settings, and missed prevention opportunities because risk stratification tools could not identify the right patients early enough.

The AHRQ Prevention Quality Indicators quantify preventable admissions for specific conditions at the county and state level. Across conditions including congestive heart failure, bacterial pneumonia, urinary tract infection, and chronic obstructive pulmonary disease, the PQI documentation establishes that a substantial and measurable proportion of hospitalizations follow from failures of outpatient management that better clinical intelligence could have intercepted.

Annual US healthcare waste	**~$750B (2012 est.)**	IOM Best Care at Lower Cost (2012) — significant proportion attributable to fragmentation and missed prevention
30-day readmission rate (high-complexity)	**15–20%**	CMS Hospital Readmissions Reduction Program data — persistent challenge for multi-morbid patients
SDOH contribution to health outcomes	**30–55%**	WHO Commission on Social Determinants of Health (2008); affirmed by CDC, NAM

3.2 The Risk Score Accuracy Gap in Medicare Advantage

Medicare Advantage plans are paid based on the predicted risk of their enrolled population. The HCC model is the mechanism by which that prediction is made. When the model systematically underpredicts risk for high-complexity members, plans are underpaid for the care those members require — creating financial pressure that ultimately affects care delivery for the most vulnerable members.

The mechanism of underprediction is structural. Conditions that interact clinically are not coded to reflect that interaction. Social determinants that modify clinical trajectories are not captured in claims at all. Behavioral and functional factors that predict utilization are invisible to a system that sees only what is billed. The gap between what a high-complexity patient actually costs and what the HCC model predicts they will cost is not random variation — it is a systematic consequence of architectural limitations.

Across synthetic cohort analyses conducted using HPDM methodology on documented public population data, the average RAF delta — the difference between current HCC scoring and HPDM-identified scoring — has ranged from +0.225 to +0.238 per member for populations with high chronic condition burden. For a plan serving 100,000 members with this profile, this represents a substantial and documentable underpayment that compounds annually as the population ages and conditions progress.

Section 4 — What the Architecture Needs to Be

4.1 The Logical Requirements

The preceding sections describe a problem with a clear structure. The solution can be derived from that structure logically, without invoking any specific technology or vendor. What the evidence demands is an architecture that:

1. **Organizes around the person —** Not the encounter, the code, or the transaction. The patient as a continuous, longitudinal entity — the same person across all care settings, all providers, all time.
2. **Captures all four data domains simultaneously —** Clinical conditions. Social determinants. Behavioral and functional factors. Immutable biological characteristics. Not as separate feeds to be reconciled later — as integrated dimensions of a single patient record.
3. **Models condition interactions multiplicatively —** Not additively. The clinical literature documents that conditions compound. The architecture must reflect that mathematical reality.
4. **Scores at the individual level, not the population level —** Risk adjustment models designed for actuarial population averages cannot identify which specific patients need intervention. The architecture must produce individual-level scores that reflect the whole person's actual clinical profile.
5. **Updates continuously —** Not annually at coding time. Clinical conditions evolve, social circumstances change, and behavioral patterns shift. The architecture must reflect the patient as they are now, not as they were documented to be in last year's claims.
6. **Is transparent and defensible —** Clinicians, administrators, payers, and patients must be able to understand why a score is what it is. A black-box algorithm that produces an output without an auditable rationale cannot be trusted in clinical decision-making.

These are not aspirational design principles. They are the direct architectural implications of the evidence documented in the preceding sections. Any architecture that meets these requirements will perform better than any architecture that does not, on every outcome metric that the literature has identified as systematically underachieved.

4.2 Why Incremental Improvement Is Not the Answer

The natural response to documented inadequacy in an existing system is to improve it. Apply better algorithms. Add more data sources. Refine the weighting. The HCC model has been revised multiple times — v22, v24, v28 — with each revision improving predictive accuracy at the population level.

These improvements do not address the structural problem. A more accurate additive model is still an additive model. It still treats conditions as independent. It still cannot capture comorbidity interactions. It still cannot incorporate social determinants as a first-class clinical variable. It still scores at the population level rather than the individual level.

The same logic applies to EHR-embedded analytics, population health management platforms, and care management software. These systems can be refined, their algorithms improved, their data inputs expanded. But they are built on the same organizing principle — the billing code, the encounter, the transaction — that the IOM identified as the core structural failure in 2001.

> Trying harder with the existing architecture will not produce the outcomes the evidence shows are achievable. Only a different architecture — one that starts with the person as the organizing unit — can close the gap the literature has documented for thirty years.

Section 5 — The HPDM Architecture

This section introduces the Hyper-Spherical Patient Data Model (HPDM™) as a direct architectural response to the requirements derived in Section 4. It appears here — in the final section — because the preceding four sections constitute the argument. HPDM is the conclusion, not the premise.

5.1 The Organizing Principle

HPDM places the individual patient — not the billing code, not the encounter, not the institution — at the center of the data architecture. Every data element collected about a patient is organized as an attribute of that patient, classified according to how mutable it is and what domain of their life it reflects.

This organizational choice has a specific architectural consequence: when you compute a risk score, you are computing a property of a whole person, not aggregating independently weighted codes. The person is the unit. Everything else is a dimension of that person.

5.2 The Four Shell Classification Framework

HPDM classifies patient data into four shells, representing decreasing mutability — how much the data can change in response to intervention, and therefore how it should be weighted in clinical decision-making.

Shell	Classification	What it captures and why it matters
Shell 0	**Immutable Biological**	Genetic factors, birth history, sex, race/ethnicity, congenital conditions. These factors do not change. They form the permanent baseline risk floor that all other shells modify.
Shell 1	**Near-Immutable Environmental**	Geography, housing stability, community environment, and social determinants of health. These change slowly and constrain what clinical interventions in Shells 2 and 3 can achieve. A clinical protocol that ignores Shell 1 will consistently underperform.
Shell 2	**Semi-Controllable Personal**	Behavioral, cultural, lifestyle, and personal preference factors — health literacy, medication adherence patterns, dietary behavior. Can be influenced by care interventions but not controlled.
Shell 3	**Interactable Clinical**	Active diagnoses, medications, procedures, utilization, care gaps, lab results, vitals, and risk-adjustment-relevant coding. The layer where clinical interventions directly apply — and the only layer traditional risk models see.

5.3 The Three Scoring Axes

HPDM scores each patient across three axes, producing a three-dimensional clinical profile rather than a single linear score. The three axes correspond to the three dimensions of clinical complexity that the literature identifies as independently important.

The X-axis captures chronic condition burden — the breadth and severity of a patient's active conditions across a 221-condition library, scored using 59 validated clinical instruments. The Y-axis captures functional status and care utilization patterns — how the patient is actually using the healthcare system and how their functional capacity affects care delivery. The Z-axis is the architectural differentiator.

The Z-Axis: Multiplicative Comorbidity Interaction

The Z-axis models the interaction effects between conditions — the same interactions documented throughout the clinical literature that additive models structurally cannot capture. Where HCC v28 assigns diabetes a risk weight and chronic kidney disease a risk weight and adds them, HPDM's Z-axis computes the interaction between the two: the amplification of renal progression risk that diabetes creates, the modification of diabetes management that reduced renal clearance requires, and the combined effect on the patient's overall clinical trajectory.

The HPDM Z-axis interaction engine models 65 validated comorbidity pairs, 29,403 total Z-axis interaction paths, and up to 23,205 three-way combination paths for patients with three or more interacting conditions. These are not theoretical constructs. Each documented interaction path corresponds to a pattern in the clinical literature demonstrating that the conditions in question compound rather than simply co-exist.

The mathematical consequence is significant. A patient with 12 active conditions generates 66 pairwise interaction evaluations under HPDM's Z-axis — against 12 independent condition assessments in an additive model. The clinical picture that emerges is fundamentally different.

5.4 What HPDM Produces

For each patient, HPDM produces a three-dimensional clinical profile that reflects their whole-person complexity — clinical, social, behavioral, and biological — with condition interactions computed multiplicatively. This profile supports three categories of output that the preceding evidence identifies as necessary but currently unachievable with standard architectures.

RAF Accuracy: HPDM identifies conditions that are clinically present and documentable but absent from current HCC coding — not through upcoding, but through complete, accurate documentation of what the clinical record supports. Validated synthetic cohort analyses show RAF delta improvements of +0.225 to +0.238 for high-complexity populations.

Care Gap Identification: By computing the whole-person profile across all four shells, HPDM identifies care gaps that single-source analytics cannot see: the diabetic patient whose CKD progression is accelerating faster than their HbA1c would suggest; the heart failure patient whose housing instability makes medication adherence structurally impossible; the COPD patient whose depression is the primary driver of exacerbation frequency.

Protocol Generation: HPDM's scoring output feeds directly into protocol-generation logic — individualized clinical action recommendations based on the patient's actual three-dimensional profile, not a population-average care pathway.

5.5 Patent-Pending Status and Disclosure

The HPDM architecture — including the four-shell classification framework, the three-axis scoring model, and the Z-axis comorbidity interaction engine — is the subject of a provisional patent application filed May 5, 2026 by Hypersphere Health, Inc. The application was filed with micro-entity status. The 12-month window for filing the non-provisional application is open.

This paper describes the HPDM architecture at a level of detail appropriate for an academic and policy audience. It does not reproduce the specific patent claims or the technical implementation specifications contained in the provisional application.

Section 6 — Conclusion

The gap between what clinicians know and what healthcare data systems can see has been documented for thirty years. The IOM documented it in 2001 and again in 2012. The AHRQ has quantified it in preventable admission and readmission data. MedPAC has measured its financial consequences in Medicare Advantage underpayment analyses. The WHO has established its dimensions in social determinant research that shows 30 to 55 percent of health outcomes are invisible to clinical data systems.

The gap persists not because the clinical knowledge to close it is absent. Clinicians have always known what a whole-person view of their patients would require. The gap persists because the data architecture built to support clinical decision-making was organized around billing codes, not people — and that organizing choice has compounded across every system, platform, and analytics tool built on top of it.

Closing it requires not an incremental improvement to what exists, but a different architecture — one that starts with the person as the organizing unit, classifies data by how it relates to that person's complete clinical reality, and computes risk by modeling how a person's conditions interact rather than by adding their weights.

HPDM was built to be that architecture. The preceding sections contain the argument. This sentence is the conclusion: the data architecture that healthcare's clinical evidence has demanded for three decades is now patent-pending.

> The results the healthcare system has been unable to achieve — meaningful reductions in preventable admissions, avoidable readmissions, and undertreated complexity — do not require new clinical knowledge. They require a data architecture that finally reflects the clinical knowledge that has existed for decades.

References

The following sources are cited or directly referenced in this paper. All are publicly available through the indicated channels.

1. Institute of Medicine. (2001). Crossing the Quality Chasm: A New Health System for the 21st Century. National Academies Press. Washington, DC.
2. Institute of Medicine. (1999). To Err is Human: Building a Safer Health System. National Academies Press. Washington, DC.
3. Institute of Medicine. (2012). Best Care at Lower Cost: The Path to Continuously Learning Health Care in America. National Academies Press. Washington, DC.
4. Agency for Healthcare Research and Quality. Prevention Quality Indicators (PQI) Technical Specifications. AHRQ Quality Indicators. Rockville, MD. Available at: qualityindicators.ahrq.gov
5. Agency for Healthcare Research and Quality. Medical Expenditure Panel Survey (MEPS). Household Component Data. Rockville, MD. Available at: meps.ahrq.gov
6. Centers for Medicare and Medicaid Services. (2023). HCC Version 28 Model Documentation: Risk Adjustment Model for Medicare Advantage. Baltimore, MD: CMS.
7. Centers for Medicare and Medicaid Services. Hospital Readmissions Reduction Program (HRRP). Available at: cms.gov/medicare/payment/prospective-payment-systems/acute-inpatient-pps/hrrp
8. Medicare Payment Advisory Commission. Report to the Congress: Medicare and the Health Care Delivery System. March Reports (multiple years). Washington, DC: MedPAC. Available at: medpac.gov
9. Office of Inspector General, Department of Health and Human Services. Medicare Advantage: Questionable Use of Health Risk Assessments (multiple reports). Washington, DC: OIG.
10. World Health Organization. (2008). Closing the Gap in a Generation: Health Equity Through Action on the Social Determinants of Health. Commission on Social Determinants of Health. Geneva: WHO.
11. US Department of Health and Human Services. Healthy People 2030: Social Determinants of Health. Office of Disease Prevention and Health Promotion. Washington, DC.
12. National Academies of Sciences, Engineering, and Medicine. (2019). Integrating Social Care into the Delivery of Health Care. National Academies Press. Washington, DC.
13. Hypersphere Health, Inc. (2026). HPDM™ Clinical Intelligence Platform — Canonical Reference Document v2.0. Internal document. Lansing, MI.

Correspondence: David D. Nethaway, MBA, CSM · Hypersphere Health, Inc. · Lansing, Michigan