WHITE PAPER

Large Human Network Models (LHNM)

The Human Intelligence Bridge. A New Model Class Between Artificial Intelligence and Human Networks

A position paper on connecting expert knowledge, relationship networks and large language models

David A. Pieper
Investor · Entrepreneur · Family Office

First Edition · November 2026
© 2026 David A. Pieper. All rights reserved. This white paper is not investment advice.

Contents

Abstract
1 · The question you cannot google
2 · Three sources, three model classes
3 · Definition and distinctions
4 · The data
5 · Architecture
6 · How experts build their Large Human Network Model
7 · Operation
8 · Transferability
9 · Impact
10 · Counter-thesis
11 · Limits and testable hypotheses
Annex

Abstract

Large language models answer questions whose answer is in texts. A substantial share of the economically decisive questions, however, has no answer in texts. The answer lies with people: in what they can do, own, want and decide, and in what an experienced expert knows about them without ever having written it down. This white paper describes a model class built for this knowledge. A Large Human Network Model (LHNM) connects three things that have so far existed separately: the unpublished domain, access and behavioral knowledge of an expert, a team or a house, a leading expert in their field; the relationship network these people have built over decades; and the reasoning capability of the best large language models. The result is not a language model and is not meant to replace one. It is a layer above it that learns from a different source than any language model: not from what humanity has written, but from what a person or a team knows about their world and whom they can reach in it. The paper defines the class; distinguishes it from Vertical AI, internal enterprise AI, relationship intelligence software, expert networks, expert clones and AI matchmaking services; describes the data basis, architecture, construction and operation of a model that emerged from a live case in the world of family offices and is being developed further there; and shows why the pattern transfers to any domain in which decisive knowledge lies with a few people. What is described here could not be built until a few years ago: the knowledge and the relationship network of a person or a house in one place, connected, retrievable and able to learn. Such a model can be built in two directions: outward, as access to knowledge and people that can be monetized with third parties, and inward, as the memory of a group of experts, a family or a company that would otherwise be lost with the people who hold it. The counter-thesis and its various facets are examined at the end; the paper closes with testable hypotheses. To the author's knowledge, it is the first description of this model class under this name.

1 • The question you cannot google

Anyone who asks a large language model, Google or Perplexity today how to value, for example, a large residential complex with 300 units gets a usable answer. Anyone who asks how to sell it does not get a usable answer: which three families in the US, in Germany or in Japan would buy exactly this example complex, in this location, in this quality, at this price, in the next six weeks, why, and through whom they have to be approached. This is a product without a current, liquid, transparently available price tag, one that has no daily quote like a stock and that only a few people, some of them not even known, are able to buy. The same applies to a stake in a company, a share in a venture capital fund, a work of art, a car collection and much more. The answer you get is rarely simply empty. At first glance it often looks good, but it is half-baked in parts, wrong in parts and not infrequently points in the wrong direction; and it depends on which model you use and how skilled you are at asking it. Not because the model is too small, but because the answer is not in any text a model has ever been trained on. It is in the heads of a few people: an inner circle of experts who have known these buyers and investors, for example family offices, for twenty years, and a few firms that work with them, where again it is individual people with special access, such as the broker who has worked for the family for years, the auditor or the lawyer. Nothing these people know can be found online.

The pattern repeats in every industry in which decisions are made by a few. Which director works again with which producer even though the last film failed? Which promoter books an unknown artist on whose word? Which bank finances a family

business against covenants instead of against equity, and who there makes that decision? Which supplier to the defense industry actually delivers when it matters? These are not knowledge gaps that the next, larger model will close. They are questions about people, about their abilities, their capital and their decisions. The answer lies in access to them and in knowledge about them. Any high-net-worth individual can call a major bank. But who there handles lending against assets, such as a Lombard loan with a cross-border element, whether the bank does such a thing at all and whom to talk to about it: that is the valuable information, and it is written down nowhere.

You can picture the world's knowledge as an iceberg (Fig. 1). Above the waterline is everything that has been written, published, scanned and trained into models. Below it is what practitioners know and have never written down: who is reliable, who pays, who keeps commitments, why a family office never takes part in bidding processes, which surgeon stays calm in which complication, which minister trusts which ambassador. In 1966 Michael Polanyi coined the phrase for this: we can know more than we can tell.[1] No language model contains this part of the iceberg, no matter how many billions go into the next training run. It was never written, and it will not be, because its value lies precisely in the fact that it was not written. It is an entire world of access knowledge, context and relationship history that lies outside every training corpus. If you like: the red pill from the film "The Matrix".

The thesis of this paper is that this part of the iceberg can be opened up, not through more training on more text, but through a different kind of model that learns from a different source. Peter Thiel described the difference between the better and the new as the difference between 1 to n and 0 to 1:[2] 1 to n means improving or multiplying something that exists; 0 to 1 means creating something that did not exist before. A larger language model is 1 to n. A model that learns from people instead of texts starts at zero.

A word on why this paper is being written now and not ten years ago. Anyone who has worked with such networks for a long time has tried to capture them: in databases, in client systems, with their own developers, at great expense and always only in parts, because the connections between the entries had to stay in someone's head. A database knows that two names exist. It does not know that one has owed the other a favor for twelve years. Since language models have been able to read, connect and draw conclusions, this knowledge can for the first time be brought together in one place, with all its connections, and it keeps learning as it goes. Those who would benefit most from it do not yet know that it works, or how.

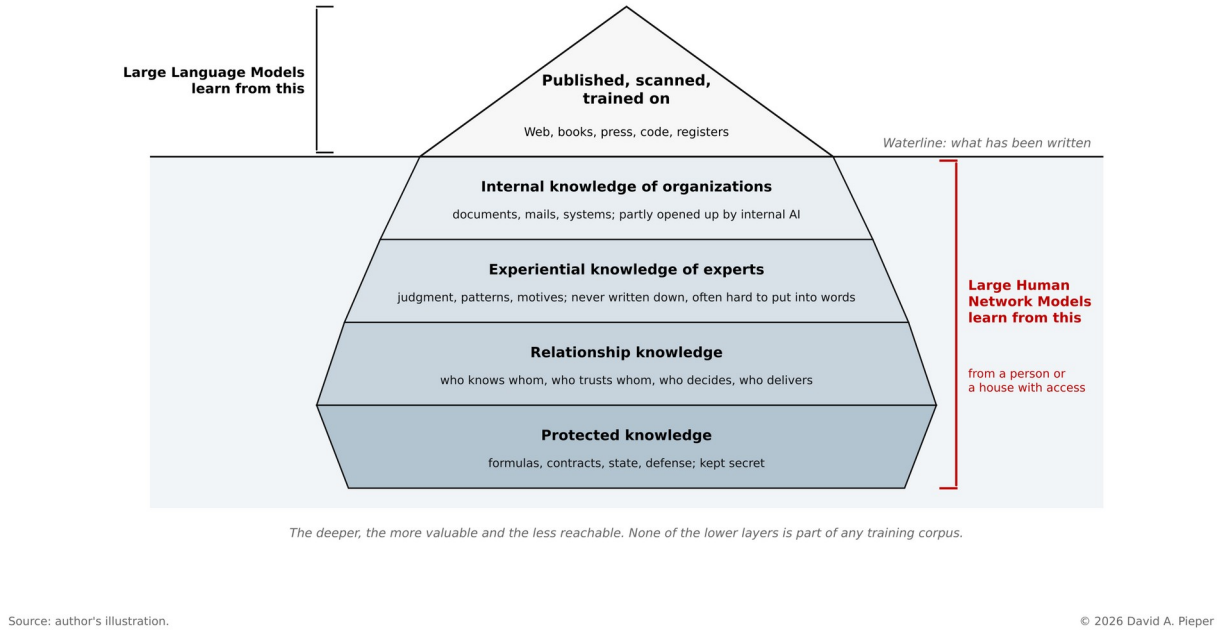

Fig. 1: The iceberg of knowledge. Above the waterline, what has been written and trained into language models; below it, what was never written. Author's illustration. © 2026 David A. Pieper.

2 · Three sources, three model classes

The major model classes of today can be ordered by what they learn from. This axis matters more for understanding them than the question of how a model is built, because it determines which questions it can answer (Fig. 2).

Large Language Models learn from text, and from speech that has been transcribed and is therefore text again. The transformer architecture[3] and scaling to hundreds of billions of parameters[4] have produced models that reconstruct a surprising amount of what humanity knows from what it has written. Their limit is the limit of what has been written.

World Models learn from the physical world. The term goes back to David Ha and Jürgen Schmidhuber, who in 2018 described generative models of an environment in which an agent acts.[5] In 2022 Yann LeCun made it the core of an architecture for autonomous machines.[6] Since 2025 World Models have been an industry: NVIDIA presented Cosmos, a platform for "Physical AI",[7] and World Labs, the company co-founded by Fei-Fei Li, is to be acquired by AMD for around 8.2 billion US dollars in stock under an agreement announced in September 2026.[8] World Models learn how things behave in space so that robots and vehicles can act in it. Their limit is the limit of what can be measured. The current state shows how early all of this still is: in 2025 around 13,000 humanoid robots were shipped worldwide,[9] and the bridge into the physical world is essentially being built by people. Since 2024 Tesla has paid employees to perform movements in motion capture suits and VR headsets for the robot to imitate;[10] data companies have thousands of people in more than fifty countries film everyday hand movements for robots;[11] in 2023 Google DeepMind pooled demonstration data from 22 robot types across 33 labs.[12] Critics such as the roboticist Rodney Brooks, co-founder of iRobot, consider the first profitable deployment of humanoid robots to be more than ten years away.[13]

Large Human Network Models learn from people. They learn who knows, can do, owns, wants and decides what, with whom and since when, and how to reach that

person. Their source is not humanity, but a person, a team or a house with decades of access to a world that cannot be written down. Their limit is the limit of this access.

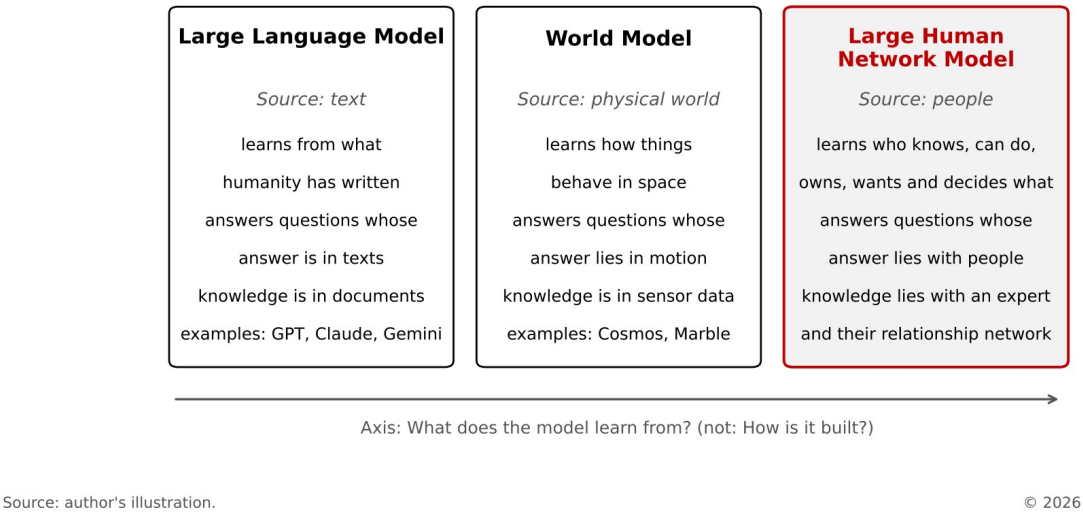

Fig. 2: Three model classes, ordered by their learning source. Author's illustration. © 2026 David A. Pieper.

Two other developments do not belong in this series, although they are often mentioned alongside it. The first is agents: language models that act, operate tools and complete tasks on their own over hours. In September 2026 OpenAI presented "Dots," continuously running agents,[14] and in July 2026 Meta gave its assistant agent capabilities with Muse Spark 1.1.[15] Agents will take over jobs and change industries. At their core, however, they change nothing about what the model learns from. An agent can query registers, write emails and set up meetings. It cannot call a confidant, cannot open a filing cabinet, cannot read a database that no one gives it access to, and it cannot have lunch and discuss confidential matters over it. The second development is self-improving AI, referred to in research as recursive self-improvement. In 2025 Google's AlphaEvolve sped up computing kernels that shorten the training of its own foundation,[16] and Dario Amodei, co-founder of Anthropic, wrote in September 2026 that this process is already beginning across the entire industry.[17] Both describe how models act and how they get better, not what they learn from. The same applies to quantum computers and the software that will run on them. These are other axes, and they can accelerate each of the three classes, including the one described here.

Two gaps

The situation can be drawn as terrain (Fig. 3). In the middle is what AI reaches today: everything that has ever been written, and increasingly also acting in software. To the left is the physical world, separated by a gap that World Models and Physical AI are now trying to bridge, built from demonstrations, simulation and sensors, among other things. To the right is the human world: confidants, former clients, filing cabinets, closed systems, experiential knowledge, the knowledge of who knows, can do, wants and decides what. So far no bridge leads across this second gap, and agents reach as far as its edge and no further.

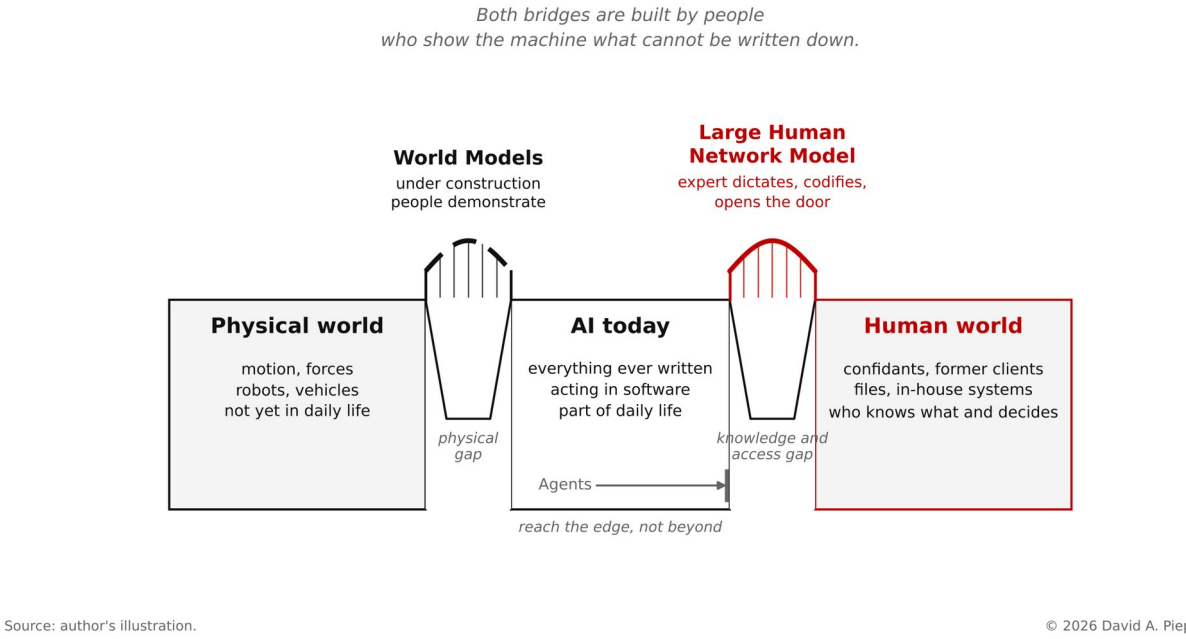

Fig. 3: Two gaps. What AI reaches today, what it does not reach, and which bridge leads where. Author's illustration. © 2026 David A. Pieper.

A Large Human Network Model is the bridge across the second gap, and it is built on the same principle as the first: a person shows the machine what cannot be written down. With the robot, the person demonstrates the movement. With the LHNM, the expert dictates and describes his judgment, "trains" it in a figurative sense, and opens the door. This parallel is more than an image. It is also the answer to the question of why an LHNM is a model, to which the next section returns.

3 · Definition and distinctions

Definition

A Large Human Network Model is an AI system with five properties. First, it is built on the unpublished domain, access and behavioral knowledge of a recognized expert, a team or a house. Second, it contains that expert's real relationship network as a structured data basis in which every entry carries its origin, time and context. Third, it uses a large language model as an interchangeable reasoning engine. Fourth, it answers questions whose answer lies with people. Fifth, it actually opens the way to these people, through the expert, the team or the house itself or through persons they have authorized to do so. Outward, it does not hand out the network, but answers and opened doors. The raw data never leaves the vault.

The core statement is: **Large Language Models answer questions whose answer is in texts. A Large Human Network Model answers questions whose answer lies with people: in their abilities, their capital, their decisions.** It is the bridge between artificial and human intelligence, and in both directions: the machine opens up the knowledge of the person, and the person gives the machine access to a world it would never reach on its own. Fig. 4 shows this image as a suspension bridge. Two pylons carry the load: the expert with his judgment and his set weights, and the network with its relationships over decades. The rented language model is the main cable, which distributes the load and remains interchangeable. Connectors, skills, memory and rulebooks are the hangers. Across the deck, requests go in one direction,

and answers and opened doors in the other, never raw data. The raw data lies in the foundations under water, in its own vault.

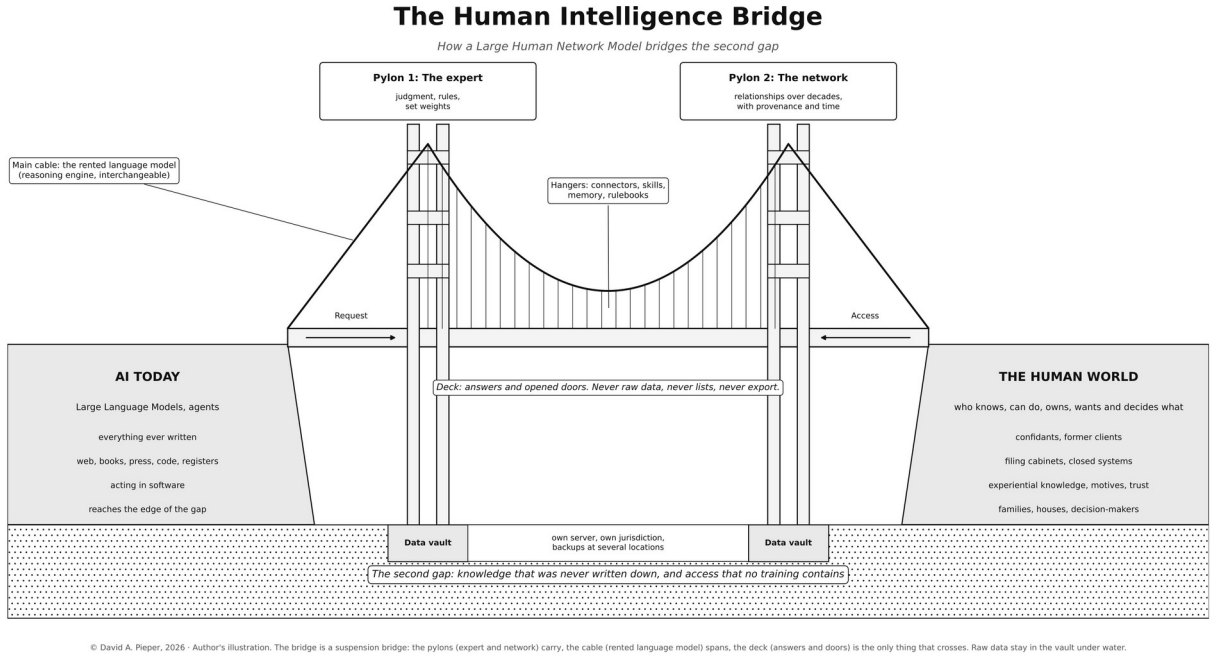

Fig. 4: The Human Intelligence Bridge. A Large Human Network Model as a suspension bridge across the second gap. Author's illustration. © 2026 David A. Pieper.

Set weights, living parameters

The obvious question about this definition is what, mathematically, makes such a system a model if it does not train weights of its own. The precise answer is: it does not train learned weights, but it has set ones. Formally, an LHNM is a weighted, directed graph with a scoring function over it.[18] The nodes are people and firms with attributes: access tier (three tiers in the case described here), role, decision-making power, capital, motives. The edges are relationships with strength, direction, trust, reciprocity and date of last contact. Above this sits a function that determines, for a specific request, the best path through this network, similar to the way PageRank derives the importance of a page from the structure of links,[19] except that here the weights are not computed from a structure but set by the expert: contact with an owner weighs more than contact with an intermediary, an intermediary is approached only through the owner, and below three matches nothing is displayed. What exists today in the case described here as a rulebook in text form is this function, just not yet written as a formula.

Three things follow from this. First, there are few parameters, all readable, all verifiable, all owned by one person, one team or one house. This is the opposite of the billions of weights of a language model, none of which means anything on its own and whose interplay not even their own developers and researchers can trace in detail. In an LHNM, a single parameter decides whether a family with a fortune in the billions is approached at all about a request, and through whom. It is not the number of parameters that makes the model, but what a single parameter moves, and behind it stands a data basis that exists nowhere else: text from the internet is available in abundance; several tens of thousands of verified high-net-worth individuals with relationship history, as in the case described here, are not. Second, the schema is general and the values are individual. Access tiers, value classes, relationship strength

and trust apply to every expert in every domain; who stands on his first tier and what counts as valuable to him, he alone decides. The schema is therefore what transfers; the values are the trade secret. Third, the weights are set, but not static. They change in three ways, and in this the system differs from a language model, whose weights change only through new training on new data and which no one can read afterward: the expert sets them and changes them when his judgment changes; the work updates them, because every request, every introduction and every reaction shifts relationship strength and trust; and the machine proposes changes when it sees in new data that a contact has changed roles, a family has sold or an intermediary has not delivered, and the expert confirms. This is learning without backpropagation, more like a trust score that keeps updating. In this sense an LHNM is a living parameter system. At the same time it grows with every advance in language models, because it can draw on their entire world and context knowledge at any time.

Added to this is a second level that only becomes visible above a certain amount of data: patterns. With thirty contacts there are no patterns, only individual cases. With thirty thousand, regularities emerge that no expert holds completely in his head: which family office sells in which phase, which intermediaries actually deliver, which route of access works with which type of counterparty, which requests fail and why. This is statistics over the corpus, not memory, and the language model detects it across the mass, while the expert can check it immediately because he knows the cases. The surgeon who has operated on a thousand celebrities is not necessarily better at the craft than the one who has operated on one. He knows the distribution: what this type of patient expects, pays, keeps quiet about, where he drops out. This knowledge does not lie in any one contact but between the contacts, and it is another reason why "Large" is in the name. Below a threshold, such a system is an address book with notes. Above it, it is a model.

A critic will object that this is a scoring system, not a model. Statistics sees it differently: a model is a parameterized function over data, and whether the parameters are estimated or set changes nothing about that.[20] The paper therefore keeps the word, but says exactly what is meant: a small, readable parameter system owned by a person, over a scarce data basis, connected to a large, rented language model that finds the patterns in it.

Why "Model" and why "Large"

The second objection is obvious: an LHNM does not train the large weights that make up the language model. It rents them. Is it then a model at all, or just an application? The answer lies in what the system represents. A World Model is called that because it contains a learned, continuously updated representation of a world, not because it runs on a particular chip. An LHNM contains a learned, continuously updated representation of a human network: who, what, with whom, since when, from what motive, with what result. This representation is the model. The language model is the machine that reads it, queries it and connects it with the knowledge of the world. Together they form a system that learns with every request, every judgment and every result, which is exactly what one expects of a model. And it comes about in the same way as a World Model: a person shows the machine what is in no text, there by demonstrating, here for example by dictating and deciding. To put it plainly: the weights belong to the provider. The model of the network and the judgment encoded in it belong to the expert. **Anyone can rent the best AI. No one can rent twenty years of relationships.** On the physical side of the gap, this division of labor is already a category of its own with its own capital market. The Munich company

microagi, founded in 2025, by its own account builds neither robots nor foundation models, but a hardware-agnostic and model-agnostic layer that adapts existing models to the customer's factory using the customer's operating data and a proprietary data corpus, without being tied to a provider;[21] in July 2026 it received 55 million US dollars for this, the largest seed round Germany has seen to date, led by Hummingbird Ventures with Northzone, LocalGlobe, Village Global and redalpine.[22] There this is called post-training, and microagi is regarded as a research lab. An LHNM is the same principle on the human side: not the foundation model, but the layer that connects it with what only the owner has.

"Large" here does not refer to the number of parameters, but to the size of the corpus and the reach of the network: correspondence over decades on the order of several hundred thousand messages and attachments, structured holdings on the order of a thousand tables and a hundred thousand records, thousands of decision-makers known personally. That is small for a language model and extraordinarily large for a human network.

What an LHNM is not

The building blocks of an LHNM exist individually. The paper does not claim to have invented them. It argues that their combination forms a class of its own. The following therefore lists the related variants, models and forms and sets out where, in the author's assessment, the differences lie (Table 1).

Vertical AI is the venture capital term for applying large models to an industry with its own data and workflows.[23] Harvey, for law firms, was valued at 15.5 billion US dollars in September 2026,[24] Abridge, for doctor-patient conversations, at 5.3 billion in 2025.[25] With these valuations, the dismissive label "wrapper" has had its day.[26] The investor logic behind it is notable and supports this paper: models are becoming interchangeable, probably a commodity; the moat is data that no one else has.[27, 28] In this sense an LHNM is a Vertical AI. The difference lies in what is encoded. Vertical AI encodes the documents and processes of an industry. An LHNM encodes the people and relationships of an expert, team or house and opens the door to them.

Internal knowledge AI in law firms, banks and corporations answers from its own document base and says who in-house knows about a question. The answer is a document or a colleague, the system serves its own people, and normally no firm sells access to its relationship graph, because the relationships are the business itself. A bank knows internally who knows whom, and normally does not turn that into a product for third parties.

Relationship intelligence software such as Affinity or Louisa AI maps from relationship data, in Louisa's case explicitly from emails and calendars, who in the firm has the warmest connection to a target, and suggests the route.[29, 30] Technically, this is the closest relative of the data layer of an LHNM. What is missing is the judgment of an expert, access from outside and the door.

Expert networks such as GLG, AlphaSights or Third Bridge sell conversations with outside specialists by the hour. They have no relationship data, no model and no expert who vouches for the connection with his name. Services such as Pearl now put an AI in front of this kind of brokering.[31]

Expert clones such as Delphi, which raised a round led by Sequoia in 2025, train a replica of a person on their texts, podcasts and posts and sell conversations with it.[32] This is the closest relative on the knowledge side. What is missing is the network, the

knowledge that was never published, and the door. A clone can say what an expert has said. It cannot call anyone or meet with anyone.

AI matchmaking services such as Boardy[33] or AYU Capital, a capital matchmaking service launched in 2026 by the family office community AYU and supported by the AI "Cupid",[34, 35] create connections between their own users. The network is built from self-described profiles and mandates, not from decades of shared business; the judgment lies with a team or with the machine, not with a person whose name means something to the other side. Marketplaces such as PartnerHQ go one step further and let network owners sell warm introductions for a fee,[36] without a model and without knowledge. This is the closest precursor for monetizing one's own network outward, and at the same time the opposite of an LHNM: it sells the connection and has no judgment.

Table 1: Neighbors of a Large Human Network Model compared. Author's classification based on public product descriptions, as of October 2026.

	Proprietary relationship data	Expert knowledge, never published	AI judgment	Door opened personally	Accessible from outside
Internal knowledge AI (Harvey and others)	partly	no	yes	no	no
Relationship intelligence (Affinity, Louisa)	yes, internal	no	partly	no	no
Expert network (GLG, AlphaSights)	no	no	no	no	yes
Expert clone (Delphi, Viven)	no	no, only published material	yes	no	yes
AI matchmaking service (Boardy, AYU)	no, user profiles	no	partly	no	yes
Intro marketplace (PartnerHQ)	yes, many individuals	no	no	yes, for a fee	yes
Large Human Network Model	yes	yes	yes	yes	yes

The second attempt

The question of whether the knowledge of an expert, team or house can be put into a machine is not new. Artificial intelligence has tried it once before. Edward Feigenbaum of Stanford, in his 1977 foundational text on the art of knowledge engineering, described acquiring the expert's knowledge as the bottleneck problem in building intelligent systems and wrote of the expert that much of this knowledge is private, and that he knows more than he is aware of knowing.[37] The expert systems of that era, MYCIN for medicine, DENDRAL for chemistry, XCON for configuring computers at Digital Equipment, hit their limit at exactly this point: knowledge engineers had to translate the knowledge into rules by hand, the experts could not say what they knew, and the rule bases decayed. At XCON the rule base grew to 6,200 rules, about half of which changed every year.[38] Management science later described the problem under the name externalization, the conversion of tacit into explicit knowledge, without having a machine that could do it.[39]

Today expert knowledge is harvested in three forms, and none of them is an LHNM. The first is the harvest for everyone: an entire industry pays specialists to turn their

knowledge into training data for the large models. In June 2026 Mercor reached a revenue run rate of over two billion US dollars with more than 30,000 weekly active experts,[40] and OpenAI reportedly had more than a hundred former investment bankers build financial models for its models to learn from.[41] This is the exact counterparty to an LHNM: the knowledge becomes generic, flows into a model that anyone can rent, and none of it stays with the expert, least of all his network. The second form is inward-facing tools: law firms use Harvey to turn their in-house knowledge into reusable workflows,[42] and corporations let systems such as Glean index their documents and find employees who know the subject. This stays in-house and encodes processes, not a person's judgment. The third form is clones and door openers, separate from each other: Delphi and MasterClass sell conversations with the AI replica of an expert, trained on what he has published or provided to them;[32, 43] Boardy and Affinity broker introductions without judgment. No one combines the two.

Research from the last two years shows that the second attempt can turn out differently from the first. An AI agent that interviewed employees of an organization in a simulation fully reconstructed undocumented knowledge in about 95 percent of runs,[44] and a 2026 study found that language models extract tacit knowledge from the behavior of experts better than the experts themselves.[45] Feigenbaum's bottleneck is opening because the machine understands the expert instead of waiting for his rules. That is the technical reason why an LHNM can be built today and could not be in 1985. The second reason is that no one back then thought of the network: an expert system was supposed to know what the expert knows. An LHNM also knows whom the expert knows.

On the term

The name "Large Human Network Model" is deliberately placed alongside Large Language Model and World Model. The word "Human" describes the function. Without this word, "Network" in the technology world first means a communication network or a neural network: in 2025 a Japanese research group used "Large Network Model" for a language model that simulates communication networks,[46] and the short form LNM is used for "Large Numerical Models"[47] and for Refik Anadol's "Large Nature Model".[48] The network at issue here consists of people, and the name emphasizes that. The closest relative by name is a 2025 essay that, under the heading "Large Relationship Models", sketches the vision of a platform model meant to learn from the aggregated relationship data of many users which connections create value.[49] The difference is fundamental: there a model learns from the behavior of many; here it learns from the knowledge and the network of one person, one team or one house. There is no expert knowledge there, no door opener and no outward monetization. None of these uses refers to an AI class that learns from the knowledge and the network of an expert. The claim of this paper is therefore narrowly drawn: to the author's knowledge, it is the first description of this class under this name. The parts are known. Their combination and its name are not. The name is also meant to help this knowledge and the possibilities it holds get around, be discussed critically and be developed further.

4 · The data

An LHNM stands or falls with its data, and it is worth saying exactly how that data differs from what is otherwise traded as data. Four sources come together.

First, correspondence. Over decades, an expert, a team or a house accumulates several hundred thousand messages and attachments. They contain not addresses but

transactions: who approached whom, when and with what request, what came of it, who replied and who did not, who turned down an offer and why. This is the chronicle of a network, and it can normally be neither bought nor used for training.

Second, the structured holdings: contact databases, spreadsheets maintained over years, in the underlying case on family offices, families, family relationships, foundations, companies, acquisition profiles, financing partners and private wealth, often half finished, often outdated, but with an age stamp on every row that is itself a piece of information. Much of this is proprietary knowledge that could not be rebuilt in a month with all the money in the world, because the knowledge holders for the most part expressly do not share their knowledge publicly, and so it cannot be extracted anywhere. This kind of knowledge, usually the bottom two layers of the iceberg, exists in similar form in many industries. The pattern is independent of industry; what it needs is the expert who can read the holdings. Without him, the value and quality of an LHNM are sharply reduced from the start.

Third, the documented expertise of the expert, the team and the house: their own research, their own papers, documented experience, rulebooks, models of why a wealthy family buys something, why it never buys, or why it invests in certain things, how decisions are made among owners, which patterns precede a failure. In its domain, this knowledge often exceeds what is taught at universities and published in studies, not because the research is poor, but because it observes from the outside. Someone who has negotiated, traveled and invested with the families for twenty years knows the part of the iceberg below the waterline. The same holds in other domains: the best producers, surgeons or negotiators in the world as a rule do not teach at universities and therefore do not pass their knowledge on to students who could develop it further. One learns about them secondhand. The decisive point is this: this knowledge is not published and therefore cannot flow into the training of language models.

Fourth, the experiential knowledge that is only put into words during the build. A large part of what makes an expert is written down nowhere, not even in his own files. It only comes out when he dictates cases, discusses them, gives reasons for decisions, formulates rules and records misjudgments. This part is the most valuable and the most laborious, much like the early days of building language models. That language models help here for the first time, because they can draw parts of this knowledge out of an expert's behavior, was shown in section 3; in an LHNM this works in both directions: the machine learns from the expert and at the same time helps him put into words what he knows.

Five properties distinguish this data from what data brokers sell. It is not public: no part of it was ever extracted automatically, because it was never online. It is longitudinal: it shows relationships over years, not states on a given date. It is relational: its value lies in the connections between entries, not in the attributes of individual entries. An example makes the difference clear: data brokers count roughly 510,000 people worldwide with assets of more than 30 million US dollars,[50] and such a list can be bought. It says who is rich. It does not say who does business with whom, who trusts whom, who will never call whom again, and who has access to whom or can create it. The first is a table, the second a network. It carries motives: it contains why someone acts, and with that, what behavioral finance describes for markets, for individual decision-makers. And it has provenance: every entry knows where it comes from and how old it is.

Sociology described the mechanics of such networks long before AI. Mark Granovetter showed in 1973 that the decisive information flows through weak ties, not close ones,[51] and Ronald Burt showed that the advantage goes to whoever bridges structural holes between otherwise separate groups.[52] An expert with decades of access sits in hundreds of such holes. Toby Stuart added in 2025 that status arises through anointment: someone introduced by a high-ranking connector is read differently by the other side.[53] An LHNM makes these three effects queryable for the first time without using them up.

5 · Architecture

The architecture of an LHNM follows from a single principle: the raw data never leaves its own vault, and only answers and doors go out. Everything else follows from that (Fig. 5).

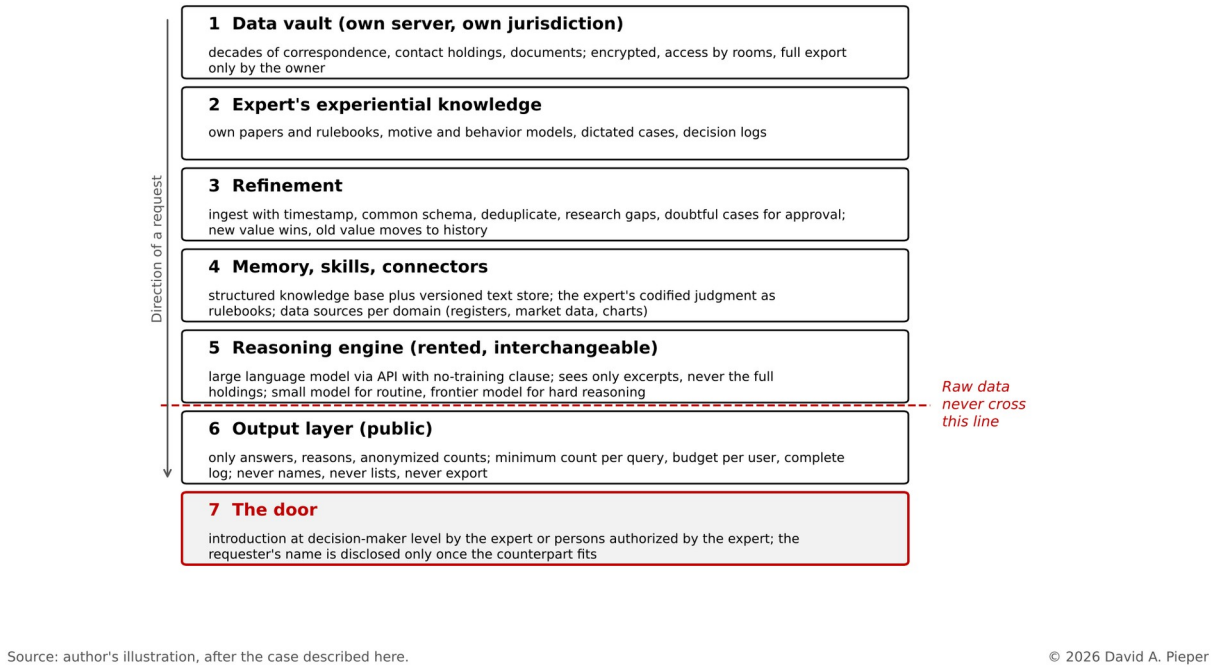

Fig. 5: Layers of a Large Human Network Model. Author's illustration based on the case described here. © 2026 David A. Pieper.

The bottom layer is the data vault: a dedicated server in one's own jurisdiction or a deliberately chosen one, encrypted, with access by rooms rather than by individual permissions. Those who work in the firm may contribute but not extract; a full export is possible only for the owner, and every access is logged, with an alarm on mass retrieval. Above it sits the expert's experiential knowledge in documented form. The refinement pipeline turns both into a knowledge base: ingest with time stamps, a common schema, deduplication, filling gaps through follow-up research via connectors, submitting doubtful cases for approval. A new value wins; the old one moves into the history.

The fourth layer is what the machine actually reads: a structured knowledge base for anything filterable, a versioned text store for theses, cases and rules, the expert's codified judgments as executable rulebooks (in the language of the tools: skills), and the connectors and interfaces to the data sources of the respective domain, including special data providers that are central to the individual LHNM and that the author deliberately does not list in full here. In the case described here, these include, for

example, the commercial register, shareholder data and market data. For a music producer it would be the streaming services' charts, for a film producer casting and rights databases, for a surgeon his own case collection, and in technical domains patent databases as well. These sources enrich the model and in part are what make it effective at all, and they too are not freely accessible: the data behind such interfaces is licensed, limited and expensive. The LinkedIn case shows how hard this boundary is defended: for years the network has taken legal action against the scraping of its profiles and opens its interfaces to only a few partners, for exactly the reason this paper is about.[54]

The fifth layer is the reasoning engine: one or more large language models, connected through an interface with a contractual exclusion from training. They see only the excerpts a query needs, never the holdings. This layer can also be run with open models if, for reasons of data protection and discretion, one wants to send nothing at all outside, at the cost of giving up the most powerful model of the moment and some convenience and ease of implementation. The models are interchangeable, and that is intentional: small, cheap models handle routine work; the demanding conclusions go to the strongest one, which incidentally gets better year by year. The provider can change without anything changing in the LHNM. Those who do not trust the providers can, as described, run this layer on their own server with an open model or one accepted in their own jurisdiction; given the formation of blocs worldwide, for some firms this will not be a matter of taste.

Above the red line lies the output layer. It knows no names. It delivers assessments, reasons and anonymized counts ("14 families in the holdings fit this request, for three reasons"), with a minimum number per count, a budget per user and a complete log. Never names, never lists, never exports. Only then comes the door: the introduction, the bridge into the human world, is made in person, at decision-maker level, by the expert or by people he has authorized, and the name of the person asking is only given once the other side is a fit.

This separation is not only security; it is data protection as architecture. An LHNM processes data of third parties who were not themselves asked. A system that releases no raw data and generates no lists is the sound starting point for this in every legal system; in the European Union, legal basis, information obligations and rights of access come on top and have to be clarified with lawyers before operation.[55] A remark on this is necessary, because this paper is not written for a single jurisdiction. Models of this kind can contribute to the knowledge and technology lead of a region or of entire industries, especially where the most valuable knowledge lies, for example, in family businesses and with individual minds. Regulation that makes even the build difficult would not protect this lead but move it elsewhere: the business would then arise where it is permitted, as is already happening today. The same applies to World Models: where every employee who demonstrates a movement to a robot must first sign a consent form and the recordings may not be stored, these models will be built in the countries that allow progress first and regulate afterwards.

6 · How experts build their Large Human Network Model

The word "train" is used cautiously here, and on purpose. An expert or a team of experts building an LHNM does not train weights in the sense of a language model. It sets the weights of the network (section 3) and does something more delicate: it puts its knowledge into a form that a machine can read, check and apply. Technically, this consists of retrieval (the machine finds the relevant excerpts), memory (it keeps what

it has learned), tool use (it queries connectors) and procedural instructions (it follows the codified rules). Fine-tuning the language model on one's own corpus is an option for later, not part of the basic build. The build proceeds in five stages.

First, ingest. Everything that exists goes into the vault with a time stamp: correspondence, holdings, documents, photos of business cards, notes. Nothing is cleaned up in the original; it is the copy that gets sorted.

Second, structure. The refinement pipeline extracts people, firms, relationships and transactions from the raw material, deduplicates them and stores the source and date with every entry. This is computing work, mostly done by small models, and it takes weeks.

Third, draw out experiential knowledge. This is the expert's real work. He dictates cases, gives reasons for decisions, names patterns, puts into words what he sees at first glance and why. In the case described here, several hundred working hours produced, for example, rulebooks for matching property and buyer, for vetting a counterparty, for the motives of families and family offices, and for handling requests. An experienced surgeon would do the same with thousands of cases: which nose suits which face, which expectation is unrealistic, which patient should not be operated on. What a practitioner knows after three data points and two phone calls has to be spelled out here sentence by sentence so that the machine can learn it.

Fourth, codify judgment as rules. The dictated cases become executable rulebooks: an acquisition profile is the main direction but not an exclusion criterion; an intermediary only comes in through the owner; a count below three matches is not displayed. These rules are what the machine learns from the expert, and they are proprietary. The method itself is again independent of industry; the rules are not.

Fifth, test and adjust. The model is measured against shadow cases: would it have found the buyers the expert found? Every decision is logged, with the outcome months later, and a "who was right" log shows where the expert's judgment and the machine's diverge. The target in the case described here is a hit rate of 90 to 95 percent for suggestions, because in this world a wrong suggestion costs more than a missed one.

The productivity this creates can only be estimated today, but the direction is clear. In the case described here, the machine takes over the groundwork that in a traditional consulting firm is done by a team of analysts, assistants and sales staff: searching, checking, matching, reaching out, following up, documenting. In the author's estimate, implemented properly, this equals the work of 20 to 30 full-time employees, with the difference that the machine checks every request against the entire holdings and not against what an employee happens to have in mind. A firm with few people can thus do the work of a much larger organization. This is an estimate from practice, not a measured value; it is formulated as a hypothesis in section 11, and the bottleneck remains the expert himself, which section 10 returns to.

Who else trains it

An LHNM does not learn only from its owner. Depending on the model, everyone who uses it contributes. In the case of a network of families and family offices, the model learns from every request what is currently being sought, from every introduction who fits with whom, and from every outcome whether the judgment was right; members of a closed circle also bring in their own contacts. In a ministry, all employees would enter what they know, and the institution would keep what today is lost with every change of staff. For government organizations in particular, this holds a potential that

has hardly been recognized so far. For the surgeon, it would be the patients' cases, with their consent. This is a data network effect in the strict sense: every user makes the model better for the next one, without anyone seeing anyone else's data. The data barrier remains absolute: what a user enters improves the model, but the owner never uses it for his own business against the user, for instance by pursuing a request himself before brokering it.

7 · Operation

A query to an LHNM runs in six steps (Fig. 6). The person asking describes the request in his own words. The machine classifies it, provides context and asks a follow-up question. It matches it against network and knowledge and reports, anonymized, whether there are suitable people, how many, and for what reasons. This takes minutes, and up to this point the expert is not involved. Then comes his judgment: is the request viable, what has to be clarified first, in what order are people approached? Then the door, and finally the outcome, which flows back into the central memory.

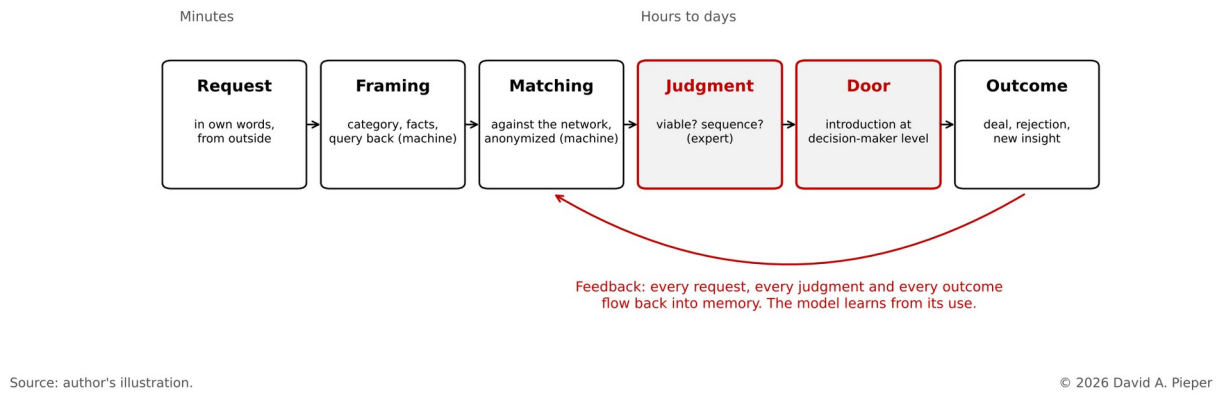

Fig. 6: Path of a query through a Large Human Network Model, with feedback. Author's illustration. © 2026 David A. Pieper.

The sequence is the point. In the conventional way of working, the expert stands at the beginning: he takes the meeting, researches, thinks, makes calls and then declines or accepts. His time goes into groundwork, and most of it is wasted. In an LHNM he stands at the end. The machine does the collecting, screening and pre-sorting, and the expert only judges what has passed the check. What is automated is the collecting, not the judging. That is the difference between an advisor who scales and one who would have to multiply himself, and it applies not only to advisors but to every function whose day consists of groundwork: sales, marketing, financing, purchasing, HR. The model stays current because the people behind it feed it daily, from any device, in passing: a sales channel in one country no longer works after a change of government, a supplier has delivered, a family has sold. Whoever joins the circle of contributors brings his knowledge along; whoever leaves loses access. What he contributed stays, and the question of ownership, to which section 9 returns, hinges on exactly that.

In the case described here, a closed funnel has the job of enforcing this difference externally: there is no other way to the expert than through the system, every request has a price, externally in money, internally as a quota, and a rejection costs the person asking nothing. Internally, in the family or the company, each employee would, for example, have a fixed quota of requests. The value lies in three dimensions. First, time: answers that otherwise take days come in minutes. Second, possibility: without

the model's knowledge, the person asking might not have been able to make the request at all, because he did not know who fits with whom. Third, stakes: in the lower layers of the iceberg, the value of a piece of information is very high, at times decisive, because it leads to contracts worth millions, carries weight for the economy as a whole or, in medicine, defense and national security, is a matter of life and death. The result can be a precision and speed that this world has not seen before: a family office and a private bank are identified as a fit within minutes, or an otherwise inaccessible piece of information is available in seconds or in less than an hour, and, as described above, it can be worth millions. Opening up this layer of the iceberg can trigger an effect in those who experience it for the first time similar to the one the release of ChatGPT had at the time.

8 · Transferability

The pattern is not tied to family offices. It fits any domain in which four conditions come together: the decisive knowledge is not public and will not become public; it lies with a few people with decades of access; the decisions at stake are large enough that access has a price; and there is demand, from outside or within the organization, that pays this price.

A few examples make this concrete. A film producer who has been among the best-connected people in his industry for forty years knows which director will never work with which actor again, which distributor buys which material, how contracts are actually negotiated and who decides at which broadcaster. His LHNM would consist of thousands of files and transactions and of his human relationships, none of which is part of a language model. A music producer knows how artists are made, how to move promoters, which contracts made or ended which careers. A top-tier venture capitalist knows more about hundreds of investments, founders and co-investors than any database; worldwide perhaps twenty such firms shape the market, and their knowledge is the reason.

The plastic surgeon is the example that shows opportunities and limits at the same time. One of the best in his field could build a model from thousands of cases, consultations and images that tells a patient before the first appointment what is possible with his face, what he should not do and what the result would look like, and that, if the procedure does not suit this surgeon, refers him to the colleague who can perform it better. That is an LHNM in the full sense: knowledge, judgment, network and door. Depending on its intended purpose, it may be a medical device and thus, in the European Union, a high-risk system under the AI Act,[56] with everything that comes with that. Regulated domains need an additional layer of approval, liability and final medical responsibility. That does not argue against the model; it determines who may build it, and it is a signal to policymakers to regulate innovation with a sense of proportion.

Three variants: the leading expert, the family, the company

An LHNM can be run in three variants, and the architecture is the same in all of them. The first is the external variant: a leading expert makes his knowledge and network available to paying third parties, as the case described here does. The second is internal use: a firm or family builds its own model and is the only one to query it. For business families and family offices, this may be the more important variant. Their knowledge is scattered: with the patriarch, with two long-standing advisors, in the family office in another location, in investments on three continents, in the minds of

people who will no longer speak with the next generation. With every change of generation, part of it is lost. An internal LHNM keeps this knowledge within the family and makes it queryable and transferable, without it having to go to a service provider.

The third variant is the company, and it is the largest. What companies today call "internal AI" is almost always a language model on top of their own servers: inventory management, drives, contracts, tickets, emails. That is useful, but it is a language model turned inward. It opens up what the company has written, not what it knows. A public example makes this concrete. The family business STIHL, founded in 1926 and today present in more than 160 countries,[57] does not owe its lead to its ERP system. In an interview in October 2026, its CEO described what the lead consists of instead: as a family business, the company does "not constantly have the capital market breathing down its neck" and does not have to correct its strategy when things go less well for a while; it is not the technology that decides but the customer; with specialist dealers, the company is "almost religious"; and the head of the US business was replaced because he wanted to "play a different game."[58] None of these sentences is in a system. Each of them decides what the company does. The lead lies with the chief engineer who knows why a tolerance applies and when it may be broken; with the head of sales who has known the customer's head of purchasing for fifteen years; with the owning family that works with the supplier in Asia on a handshake; with the auditor who knows how the deals are actually structured; with the manager who led the expansion into a country and knows the intermediaries on the ground. In an organization with a thousand employees, that may be thirty people, and their knowledge is at once the most valuable and the least secured asset of the firm. A company LHNM is the memory of these thirty: their knowledge and their network, structured, with provenance and time, selectively served by rooms, kept current daily by themselves. It is the company's brain trust,[59] except that it can be queried day and night and does not retire with the founder.

Three objections suggest themselves, and they differ in weight. The first is the serious one: in companies, knowledge is power. Whoever puts his knowledge into a system gives up part of his standing, and that, not the technology, is what made the knowledge management of the 1990s fail. The answer can only lie in ownership or other forms of participation: whoever contributes gets a stake in the model, for instance through a separate share class that grows with the value of the model and is tied to participation. For family businesses with employees who have been with the company for twenty or thirty years, this is close to what they already live. The second objection is codetermination and the protection of employee data, which is stricter in Europe than elsewhere; that is an architecture question from the start, not a reason not to design the model. The third objection is that a model which bundles the crown jewels creates the most attractive target for espionage the firm has ever had. In the author's view, it carries the least weight. For decades the world has safeguarded formulas, design plans and weapons codes that are not stolen every day. An LHNM that exports nothing, works in rooms, logs every access and raises an alarm on mass retrieval is easier to protect than the thousand laptops across which the same knowledge is scattered today. That does not mean security is secondary; on the contrary: attackers are getting better with language models, and quantum computers are about to arrive. But it is solvable, and it should not keep anyone from building.

Values as parameters

The STIHL example shows something that goes beyond the knowledge of individual people and that no system has been able to capture so far: the values of a family or

firm. Over generations or decades, every business family has developed an understanding of how it does business and how it does not, with whom, on what terms, what it would never do even if it pays off, and what it always does even if it does not pay off. This is the core DNA of most families, and it is in no ERP system, no CRM and no other software, because these systems manage transactions, not convictions. It is written down, if anywhere, in the family constitution, in the charter of a foundation or in the minds of the shareholders.[60] An LHNM is the first system into which these values can be built directly, namely as parameters in the sense of section 3: as set weights that determine which counterpartyies come into question at all, which paths the model suggests and which it does not, in what order people are approached, and when a request is declined even though it would bring in money. A family that for three generations has not sold to financial investors, does not trade with certain countries or does not play its suppliers off against each other can tell the model so, and the model acts accordingly, on every request, without an employee having to explain it anew each time. The same applies to other industries and to government actors with their respective canon of values.

This makes an LHNM the bridge that has so far been missing in the family office world: between the family, its family office, the operating business, the employees and, where the family wants it, the public. These spheres follow different logics and are run by different people; what connects them are the family's values, and until now those could only be passed on through people. A model that contains these values as parameters, together with the knowledge and the network of all spheres, is the first shared memory of these parties, and it is the place where a family can determine what it uses machines for and what it does not. The author's paper "Old Money, New Machines" addresses this question from the families' perspective; this paper supplies the tool for it.[60]

There is also a point that carries weight in Europe. The large language models carry the values of their builders within them. A 2024 study showed that their answers on the value scales of the World Values Survey come closest to the English-speaking and Protestant European countries,[61] and the Chinese models are shaped in a different direction. Whoever rents a language model rents this imprint along with it. For a family business in Upper Austria or Nagano, on many questions that is not the stance with which it wants to do business. An LHNM changes nothing about the language model, but it puts a layer on top of it in which the organization's own values are decisive, not those of the provider. As far as the author can see, this is the only way an organization today can get its own stance into an AI system without training a model.

Highly sensitive domains

It is tempting to see LHNMs where knowledge is traded openly. The greater benefit, however, probably lies where it never will be: in domains that are discreet, protected and mission-critical. Medicine, defense, diplomacy, ministries, large private fortunes. In all these areas there is unique knowledge about people, connections and access that is neither published nor handed on to successors, and that is lost with every retirement, every departure and every death. Whoever runs a defense company, heads an embassy or has known a ministry for twenty years has a network that no organization owns, only the person. An LHNM is the first form in which such a network can stay within the organization without anyone having to hand it over. That it is especially suited to these areas follows from its architecture: it releases nothing.

Who the model is for

An LHNM is a tool for the top of a field, and it is important to say so without arrogance. It does not train junior staff and does not empower the third tier of a firm; there are other tools for that. It encodes the knowledge of those who make the difference in their field: decision-makers, masters of their craft, people whose name opens doors. There are surprisingly many of them overall, but only a few in any single industry, and each of these people is a node through which others become able to act in the first place. The model is meant for them, and for those who want to do business with them or improve their structures.

9 · Impact

What changes when experts capture their networks in models? The answer has several sides, and they are set out here one after another.

For the knowledge holder

For the knowledge holder, an asset that until now died with him can for the first time be valued and transferred. The knowledge and the network of a leading expert are today his greatest possession and at the same time his most fleeting: they cannot be transferred, shared, sold or inherited; a serious illness sets their value to zero. An LHNM preserves this asset, makes it scalable, analyzable and expandable, and it makes it ready to be handed over. For family businesses and firms facing succession, this is not a side effect but the core. For the expert himself, it means that his time goes only into judgment and that his experience grows with every request instead of wearing down with every futile piece of groundwork. All of this at the beginning of a development in which AI, robotics and quantum computers are still young; what such a model can do in ten years cannot be seriously described today.

For the user

What changes for the user is what he receives. Databases sell numbers, expert networks sell conversations, language models sell text. An LHNM leads to the result: to the capital, to the buyer, to the partner, to the doctor who makes the right decision. Contacts one would otherwise not reach in a lifetime, in days instead of years. That is the reason an LHNM can have a price that grows with the value of the decision and not with the number of queries. Internally as well as externally, this calls for a new kind of pricing that no longer depends on the cost per token and whose consequences cannot be fully foreseen today.

For the AI economy

For the AI economy, a layer emerges that has been missing so far. The investments in computing power, models and research are looking for applications that prove their value directly. An LHNM does this on both sides: the owner feels the benefit in his business, the user in his result. Where this happens, trust in the technology grows faster than where it summarizes texts. For the agents that are now moving into companies, an LHNM is also the knowledge source they lack in closed worlds: an agent that is allowed to query an LHNM can act where today it stops at the edge of the plateau. It is foreseeable that a market of its own will arise from this: firms that support experts in building their LHNM, off the shelf for the standard case and tailor-made for the top. That the build takes months and requires knowledge most knowledge holders do not have is the barrier to entry and at the same time the moat

of whoever overcomes it first. Whether a product of its own can be developed from this, a model detached from the first case and sold to others, the way NVIDIA sells its platform for Physical AI to every robot manufacturer, is open and deserves critical examination. For consulting firms and the large audit firms, such a tool would be worth a great deal, as it would for technology groups that could offer it to the Mittelstand. Against this stands what this paper itself says: the schema is transferable, the content is not, and without the expert who sets the values and opens the door, what remains of an LHNM is a database with a good interface. A product would therefore be a toolkit plus hands-on support rather than a finished model.

For organizations and regions

For companies and organizations that have built up knowledge, networks and intellectual property rights over a long time, the model therefore means more than one more tool. It changes processes (the groundwork moves to the machine, the judgment stays with the human), productivity (the key people are asked, no longer searched for) and in the end the competitive position, because for the first time an organization can compete with everything it knows, owns, applies and controls, and not only with what happens to be sitting in the room. This applies especially to regions whose companies are decades or centuries old. Japan counts around 46,700 companies older than a hundred years;[62] Germany has its hidden champions and world market leaders. The advantage, however, is latent, not automatic, and it has an expiry date. In Germany, 57 percent of Mittelstand owners are 55 or older, around 545,000 successions are due by 2029, and up to 569,000 companies could close without a successor by the end of 2029.[63] The same regions that own the largest part of the iceberg below the waterline face the largest handover wave in their history and are at the same time among the most hesitant adopters of the new technology. On top of that, production itself is moving away. According to a special analysis by the DIHK from April 2026, 43 percent of German industrial companies plan investments abroad, 41 percent explicitly for cost reasons, the highest figure since 2003.[64] Family businesses are no exception: Miele, in family hands since 1899, is moving its washing machine production from Gütersloh to Poland,[65] and STIHL makes its batteries in Romania and develops its robotic lawnmowers in China because, in the words of its CEO, Germany has "a productivity problem" and he "currently lacks the imagination" for new manufacturing in the country.[58] What goes with it appears on no balance sheet: the people who know how the plant really runs, and the relationships a factory has built around itself over decades. For these regions, an LHNM is less an opportunity to be seized than a window that is closing: what does not get into a memory in the coming years leaves with the generation. A country that owns this knowledge and does not capture it, anchor it and keep it alive ends up with less than one that owns less and captures it. The same applies to every single organization: whoever does not dare, out of concern about data protection and regulators, for lack of risk capital or simply out of habit, and lets the window pass, loses not only a lead. He is pushed aside by those who have captured the same knowledge in their domain, whether in their own country or elsewhere. Knowledge protects whoever has it only as long as no one else makes it retrievable.

In case the place is lost

Another effect lies outside competition. Knowledge that exists only in heads, filing cabinets and local systems is as vulnerable as the people and buildings that hold it. Ukraine has been experiencing what this means since 2022: according to the count of the UN refugee agency, around 5.9 million Ukrainians were living abroad as refugees

in February 2026,[66] and with them the country lost engineers, doctors, entrepreneurs and administrative specialists whose knowledge and connections are now missing there; what lay in destroyed businesses, archives and public authorities is in part irretrievable. Many consider the situation in Europe less secure than at any time since the end of the Soviet Union. In such a case, an LHNM with backups at several locations in different countries is the only thing that remains of a firm, a public authority or a family when the place is lost: not the raw data alone, but the linked memory of who knew what, who knew whom and who decided what. This applies most strongly to the lowest layer of the iceberg: critical infrastructure, defense, the supply chains of semiconductor manufacturing. That the lithography systems used to expose the newest chips work with lasers from the family business TRUMPF in Ditzingen,[67] or that Japanese manufacturers such as Advantest and DISCO can hardly be replaced in testing and dicing wafers, is in every industry report.[68] What is not in those reports is the knowledge of the people who have been building, adjusting and supplying these systems for decades, and who in the chain calls whom when something fails. In Europe and Japan, a considerable part of this knowledge is at home in family businesses. The point on which most experts agree, however, lies in Taiwan: an attack on the island in which the manufacturing plants were destroyed and the people who run them were killed or fled would not only bring the production of the most advanced chips to a standstill worldwide, but would destroy knowledge that hardly exists anywhere else. This knowledge is at once what a state can least afford to lose and what is least secured today. That is not an argument for fear but for precaution, and it applies equally to family businesses and ministries.

Ownership, balance sheet and participation

With the third variant, an asset emerges that did not exist before, and with it questions that have not yet been answered. Who owns a company LHNM? In a family business the answer is simple: the family. It becomes more difficult as soon as a financial investor is involved, the company is sold or individual contributors leave. What appears on the balance sheet is unaffected by this, and that is the real point: under international accounting rules, internally generated customer lists and similar items may not be recognized as assets.[69] A model that in a few years may be among a company's most valuable assets would be invisible on its balance sheet and would appear only on a sale, as part of the purchase price. How such an asset is to be valued, transferred and taxed will occupy courts, auditors and legislators. The paper raises these questions and passes them on to the specialists as a task.

The third variant also touches on a question that is larger than this paper and becomes more pressing for society year by year. If agents take over a growing share of work, what happens to the people whose value lies in what they know? A company LHNM gives an answer for some of them: whoever has learned for a lifetime, whoever knows the customers, the suppliers, the mistakes and the ways out, is not replaced but given a stake. His knowledge becomes a share in an asset the company would not have without him, and he keeps the incentive to keep it current instead of taking it with him when he leaves. That does not answer the question of what becomes of the many whose work can be described and can therefore be automated. But it shows that there is knowledge that automation does not devalue but makes visible in the first place.

Why enterprise AI has achieved so little so far

Along the way, the third variant answers a question that has accompanied enterprise AI for two years. Use is broad: in Germany in 2026, for example, for the first time

more than half of companies with twenty or more employees use AI,[70, 71] while in the Mittelstand the most recent figure (data for 2022 to 2024) was one fifth.[72] By their own account, 3 percent of users exploit the potential strongly, 59 percent not at all, and one company in three uses AI mainly out of fear of falling behind.[70] The effect is flat: worldwide, 44 percent of companies are scaling AI across the organization, but only 37 percent see any effect on results at all, and this share has not moved since 2025.[73] Gartner expects more than 40 percent of agent projects to be canceled by the end of 2027 because of costs, unclear value or inadequate risk controls.[74] An analysis of more than 12,000 European companies by economists at the European Investment Bank and the Bank for International Settlements puts the productivity gain from AI at around 4 percent, a one-time level effect that is considerably larger for medium-sized and large companies than for small ones and grows more strongly with each additional percentage point of spending on training than with spending on software.[75] The usual explanation is data quality, skills, culture; BCG calculates that the successful companies put 70 percent of their effort into people and processes and only 10 percent into algorithms.[76] Economics knows the pattern as the productivity paradox: Robert Solow remarked in 1987 that one could see the computer age everywhere except in the productivity statistics,[77] and Brynjolfsson, Rock and Syverson have shown that general purpose technologies go through a J-curve, because the invisible investments in processes and skills precede the measurable effect by years.[78] In the author's community, one finding is often heard this fall: companies are having difficulty bringing the capabilities of the models into their processes, and adoption remains far behind what is generally expected and suggested; families and family offices are hardly ever mentioned in this debate. This paper names what is missing in this J-curve: an agent that knows the company's documents but not its judgment produces generic work. It can summarize the supplier contract and does not know that this supplier was the only one to deliver during the last shortage. The models have their effect only once they are connected to the human layer, and this layer has so far hardly been built anywhere.

10 · Counter-thesis

The previous sections rest on practical experience and theses. This section takes the opposing side and sets out its arguments. There are six in all.

First: it is not a category but a combination. Delphi plus Boardy plus Affinity makes an LHNM; the parts exist, only the connection is new. That is true, and it applies to most categories in the history of technology. A smartphone was a telephone plus a computer plus a camera. The question is not whether the parts exist, but whether the connection can do something none of the parts can. A clone cannot call anyone, a matchmaking service has no judgment, relationship intelligence software has no network from outside. The LHNM unites all three, and only through this does it answer the question with which this paper begins.

Second: large firms build this internally in weeks. Every bank, every law firm, every corporation has the data and the money. That is also true, and none of these firms will ever offer the result to outsiders, because the relationships are their business. In that sense, the LHNM is the first model of its kind that can be queried from outside. The lead lies not in the technology but in the business model and in the willingness to operate it.

Third: the expert remains the bottleneck. The judgment and the door depend on one or a few central people, and they do not scale. This is the most serious weakness, and

it can only be measured, not argued away. The relevant metric is the share of a request that the machine handles without the expert, and the number of requests he can assess per week once the groundwork is gone. In the case described here, both are being measured on the first cases. The author expects productivity to increase many times over, in connections, in time and in money; this has not yet been proven. In addition, an expert can authorize who opens doors in his name; an LHNM is also the tool with which he can take responsibility for this for the first time. For productivity and scaling this means: the groundwork scales with the machine, the judgment scales with the number of people the expert trusts. A firm that has three such people can move many times what an advisor with ten employees achieves today; one that has only one person gains time, but no second door.

Fourth: relationship knowledge decays. What was true in 2015 is no longer true in 2026; people move on, firms disappear. True, and that is exactly why every entry has a date, and exactly why the model learns from its use. A network that is not used decays. A network that is queried and confirmed daily ages more slowly than its owner. Expert systems failed not only because of knowledge acquisition but because of maintenance; an LHNM that is filled once and then left to itself would go the same way. That is why the feedback loop from section 7 is not an extra but the condition under which the model lives.

Fifth: discretion. Users grant access because they expect discretion, and an LHNM must meet this expectation at every step. How that works is set out in section 5: an LHNM that hands out data is not one. Only an architecture that hands out nothing and creates no lists deserves the trust it lives on. A second element comes in that is often missing from the debate: the benefit that implies consent. The supplier whom a music producer asks whether he can equip a tour has never explicitly consented to this request, and he will still always want it. The same goes for the colleague to whom a surgeon refers a patient he should not operate on himself. An LHNM that knows the probability is high that the other side will welcome the request, the offer or the tip, and that the information would, if it came to it, be worth money to them, does not ask against their interest but in it. Whether that is enough is decided in the end by those involved, not by the author.

Sixth: the knowledge is being harvested anyway. An industry pays specialists to feed their judgment into the large models, and in a few years the foundation model will know what today only the expert knows. For specialist knowledge this is partly true, and the paper does not contradict it. But it does not apply to what an LHNM contains at its core: who in a particular world owns, wants and decides what, whom one can trust and who opens the door. This knowledge cannot be sold as a training dataset, because it is not generic and because it remains tied to the relationship it comes from. What a foundation model can learn, an LHNM loses. What it cannot learn remains.

11 · Limits and testable hypotheses

This paper describes a system in internal operation, before it is opened to the outside; the author has already been able to test its function and effect there. It is a growing system, and the effort in data and knowledge that goes into it should not be underestimated; many organizations bring knowledge from decades that first has to be brought to the surface. Since the system works in the lower two layers of the iceberg, the details remain unpublished. What can be said: it already produces connections, opens doors and surfaces information that the author did not have in this form in twenty years of practice. The limits of the paper are thus named: it rests on a

single case, the author's own, and on a prior art search, which can never prove absence; internal systems of banks and family offices are not public. The figures on the size of the corpus are orders of magnitude, not measurements, because the holdings keep growing. The effect on the user side has been shown in individual cases, not in a series.

Five testable hypotheses follow from this. First: in at least nine out of ten cases, an LHNM finds the same counterpartys that the expert would have found without the machine, and in a measurable share additional ones he would have overlooked. Second: after the build, the share of a request that the machine handles without the expert is above 80 percent, and the groundwork corresponds to the work of 20 to 30 full-time employees of a traditional consulting firm (section 6); measured against the value chain of searching, checking, approaching and following up, which today ties up weeks of working time, that would be a jump of an order of magnitude. Third: the time from request to introduction, solution or rejection falls from weeks to days or even hours. Fourth: at least one second domain (film, music, medicine, public administration) can be mapped with the same architecture without changing the layers; this is also decisive for firms active in several industries at once, and for public authorities and ministries that have to work across departmental boundaries. Fifth: users pay prices for results that are orders of magnitude above those for data, which, as described in section 9, requires a new form of pricing should LHNMs become established. Each of these hypotheses can fail. The paper is a testable hypothesis, not proven superiority, and it explicitly invites criticism and further development.

Annex

Glossary

Large Language Model (LLM): a neural network trained on text that predicts language and derives knowledge and conclusions from it. *World Model*: a model that learns the behavior of a physical environment in order to act or plan in it. *Vertical AI*: applications of large models to one industry, with its data and processes. *Retrieval*: picking out suitable passages from a knowledge base before a language model answers. *Skill*: an executable rulebook that prescribes a procedure to a language model. *Connector*: an interface to an external data source. *Data network effect*: the effect that a system gets better for the next user with every user.

Footnotes

1 · Michael Polanyi, The Tacit Dimension, Doubleday, Garden City (NY) 1966; new edition University of Chicago Press, Chicago 2009. Quotation p. 4: "we can know more than we can tell".

2 · Peter Thiel with Blake Masters, Zero to One: Notes on Startups, or How to Build the Future, Crown Business, New York 2014.

3 · Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, Illia Polosukhin, "Attention Is All You Need", Advances in Neural Information Processing Systems 30 (NeurIPS 2017), pp. 5998 to 6008.

4 · Tom B. Brown et al., "Language Models are Few-Shot Learners", Advances in Neural Information Processing Systems 33 (NeurIPS 2020), pp. 1877 to 1901.

5 · David Ha, Jürgen Schmidhuber, "World Models", arXiv:1803.10122, March 27, 2018, https://arxiv.org/abs/1803.10122

6 · Yann LeCun, "A Path Towards Autonomous Machine Intelligence", Version 0.9.2, June 27, 2022, OpenReview, https://openreview.net/pdf?id=BZ5a1r-kVsf

7 · NVIDIA, "NVIDIA Launches Cosmos World Foundation Model Platform to Accelerate Physical AI Development", press release, January 6, 2025,

https://investor.nvidia.com/news/press-release-details/2025/NVIDIA-Launches-Cosmos-World-Foundation-Model-Platform-to-Accelerate-Physical-AI-Development/default.aspx

8 · SDxCentral, "AMD acquires AI research firm World Labs in all-stock deal worth $8.2B", September 29, 2026, https://www.sdxcentral.com/news/amd-acquires-ai-research-firm-world-labs-in-all-stock-deal-worth-82b/ (agreement announced on September 28, 2026; closing expected by the end of 2026, subject to regulatory approvals).

9 · Robotics & Automation News, "AgiBot claims top spot in global humanoid robot shipments in 2025 with more than 5,100 units delivered" (Omdia figures), January 21, 2026, https://roboticsandautomationnews.com/2026/01/21/agibot-claims-top-spot-in-global-humanoid-robot-shipments-in-2025-with-more-than-5100-units-delivered/98223/

10 · Fortune, "Tesla is hiring workers for $48 an hour to wear motion-capture suits to train its humanoid robots", August 19, 2024, https://www.fortune.com/2024/08/19/tesla-robot-hiring-workers-optimus-training-ai

11 · Michelle Kim, "The gig workers who are training humanoid robots at home", MIT Technology Review, April 1, 2026 (readers' choice as the "11th breakthrough" on the list 10 Breakthrough Technologies 2026), https://www.technologyreview.com/2026/04/01/1134863/humanoid-data-training-gig-economy-2026-breakthrough-technology/

12 · Google DeepMind, "Scaling up learning across many different robot types" (Open X-Embodiment), October 3, 2023, https://deepmind.google/discover/blog/scaling-up-learning-across-many-different-robot-types/

13 · Rodney Brooks, "Why Today's Humanoids Won't Learn Dexterity", September 26, 2025, https://rodneybrooks.com/why-todays-humanoids-wont-learn-dexterity/

14 · Business Today, "OpenAI introduces Dots: always-on AI agents that work 24/7 and act on your behalf", September 30, 2026, https://www.businesstoday.in/technology/artificial-intelligence/story/openai-introduces-dots-always-on-ai-agents-that-work-247-and-act-on-your-behalf-558688-2026-09-30

15 · Techeconomy, "Meta AI Grows Hands: Muse Spark 1.1 Turns Chatbot into an Agent That Plans, Books and Follows Through", July 28, 2026, https://techeconomy.ng/meta-ai-grows-hands-muse-spark-1-1-turns-chatbot-into-an-agent

16 · Google DeepMind, "AlphaEvolve: A Gemini-powered coding agent for designing advanced algorithms", May 14, 2025, https://deepmind.google/blog/alphaevolve-a-gemini-powered-coding-agent-for-designing-advanced-algorithms/

17 · Dario Amodei, "We Must Pace the Frontier", September 2026, https://darioamodei.com/post/we-must-pace-the-frontier

18 · Mark E. J. Newman, Networks: An Introduction, Oxford University Press, Oxford 2010, especially chapters 6 and 7 (weighted and directed graphs, centrality measures).

19 · Sergey Brin, Lawrence Page, "The Anatomy of a Large-Scale Hypertextual Web Search Engine", Computer Networks and ISDN Systems 30 (1998), pp. 107 to 117.

20 · Peter McCullagh, "What is a statistical model?", The Annals of Statistics 30 (2002), no. 5, pp. 1225 to 1310. There, a statistical model is a set of parameterized probability distributions over the data; whether the parameters are estimated or set is irrelevant to the definition.

21 · microagi, "microagi to build future of AI robotics on Google Cloud", press release (PR Newswire), July 22, 2026, https://www.prnewswire.com/news-releases/microagi-to-build-future-of-ai-robotics-on-google-cloud-302830752.html (company statement).

22 · SiliconANGLE, "Microagi nabs $55M to teach factory robots how to work", July 16, 2026, https://siliconangle.com/2026/07/16/microagi-nabs-55m-teach-factory-robots-work/ ; Sifted, "Munich robotics startup microagi raises $55m in Germany's largest ever seed round", July 2026, https://sifted.eu/articles/munich-robotics-startup-microagi-raises-55m-germanys-largest-ever-seed-round

23 · Bessemer Venture Partners, "Part I: The future of AI is vertical", September 3, 2024, https://www.bvp.com/atlas/part-i-the-future-of-ai-is-vertical

24 · Harvey, "Harvey raises $550M at a $15.5B valuation to help legal teams own their intelligence", September 9, 2026, https://www.harvey.ai/blog/harvey-raises-dollar550m-at-a-dollar155b-valuation-to-help-legal-teams-own-their-intelligence

25 · Axios Pro, "Abridge raises $300M at $5.3B valuation", June 24, 2025, https://www.axios.com/pro/health-tech-deals/2025/06/24/abridge-raise-300m-5b-valuation-billing-scribe

26 · Bloomberg News, "The Hottest AI Companies Right Now Are 'Apps'", March 6, 2025, https://www.bloomberg.com/news/articles/2025-03-06/the-hottest-ai-companies-right-now-are-apps

27 · Bessemer Venture Partners, "Part IV: Ten principles for building strong vertical AI businesses", January 28, 2025, https://www.bvp.com/atlas/part-iv-ten-principles-for-building-strong-vertical-ai-businesses ("models will cease to be a moat").

28 · Marc Andrusko, Alex Rampell (a16z), "Fruits of the Walled Garden", October 9, 2025, https://a16z.com/fruits-of-the-walled-garden/ ("the painstakingly assembled dataset that no one else has").

29 · Affinity, "Warm Intro Agent", help page, accessed October 24, 2026, https://support.affinity.co/s/article/Warm-Intro-Agent

30 · Louisa AI, "Investment Banking", product page, accessed October 24, 2026, https://louisa.ai/investment-banking/

31 · Pearl, "Real-Time Expert Matching: How Pearl's Intelligent Routing Delivers Fast, Accurate Answers", July 18, 2025, https://www.pearl.com/post/real-time-expert-matching-how-pearl-s-intelligent-routing-delivers-fast-accurate-answers

32 · Fast Company, "Meet Delphi, the AI startup that lets experts turn themselves into chatbots", June 24, 2025, https://www.fastcompany.com/91356476/delphi-ai-digital-mind

33 · Creandum, "Backing Boardy", https://creandum.com/commitments/boardy (seed round of 8 million US dollars led by Creandum, announced on January 14, 2025).

34 · FinTech Global, "AYU launches AI capital matching for family offices", February 23, 2026, https://fintech.global/2026/02/23/ayu-launches-ai-capital-matching-for-family-offices/

35 · AYU Capital, "Submit your deck", product page, accessed October 30, 2026, https://cupid.thisisayu.com/ ("Powered by Cupid, our proprietary AI capital concierge").

36 · Y Combinator, "PartnerHQ: Buy warm intros to your prospects", launch page, March 14, 2024, https://www.ycombinator.com/launches/Kee-partnerhq-buy-warm-intros-to-your-prospects

37 · Edward A. Feigenbaum, "The Art of Artificial Intelligence: Themes and Case Studies of Knowledge Engineering", Proceedings of the 5th International Joint Conference on Artificial Intelligence (IJCAI 1977), pp. 1014 to 1029; full text as Stanford report STAN-CS-77-621, https://apps.dtic.mil/sti/tr/pdf/ADA046289.pdf

38 · Elliot Soloway, Judy Bachant, Keith Jensen, "Assessing the Maintainability of XCON-in-RIME: Coping with the Problems of a VERY Large Rule-Base", Proceedings of AAAI-87, pp. 824 to 829.

39 · Ikujiro Nonaka, Hirotaka Takeuchi, The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation, Oxford University Press, New York 1995.

40 · Dealroom, "Mercor doubles to $2B gross revenue run rate as AI labs buy expert data", July 6, 2026, https://dealroom.co/news/137121-mercor-doubles-to-2b-gross-revenue-run-rate-as-ai-labs-buy-expert-data/ ; TechCrunch, "Mercor quintuples valuation to $10B with $350M Series C", October 27, 2025, https://techcrunch.com/2025/10/27/mercor-quintuples-valuation-to-10b-with-350m-series-c/

41 · The Decoder (citing Bloomberg), "OpenAI trains AI to take on junior banking tasks", October 21, 2025, https://the-decoder.com/bloomberg-openai-trains-ai-to-take-on-junior-banking-tasks/

42 · Harvey, "Introducing Workflow Builder", June 24, 2025, https://www.harvey.ai/blog/introducing-workflow-builder

43 · Kowatek, "MasterClass On Call gives you on-demand access to AI facsimiles of its experts", December 11, 2024, https://blog.kowatek.com/2024/12/11/masterclass-on-call-gives-you-on-demand-access-to-ai-facsimiles-of-its-experts/

44 · Gianlucca Zuin, Saulo Mastelini, Túlio Loures, Adriano Veloso, "Leveraging Large Language Models for Tacit Knowledge Discovery in Organizational Contexts", arXiv:2507.03811, July 4, 2025, https://arxiv.org/abs/2507.03811 (820 of 864 simulation runs with complete reconstruction; synthetic organization, not a real workforce).

45 · Mina Cho, Russell Funk, Alok Gupta, Mochen Yang, "Large Language Models Explain Experts Better Than Experts Themselves", Carlson School of Management, University of Minnesota, draft June 2026, arXiv:2608.07488, https://arxiv.org/abs/2608.07488

46 · Akimoto Yusuke, Yamazaki Seiya (ARISE analytics), "Large Network Model: 大規模言語モデルを応用した通信ネットワークシミュレーションのための生成モデルの検討", 39th Annual Conference of the Japanese Society for Artificial Intelligence (JSAI 2025), paper 2Win5-96, https://doi.org/10.11517/pjsai.JSAI2025.0_2Win596

47 · Srinivas Kilambi, Mahesh Banavar, Large Numerical Models from a Business Perspective: LNM, a Parallel Universe to LLM, Springer, Cham 2026, chapter 3, https://link.springer.com/chapter/10.1007/978-3-032-14869-8_3

48 · Inavate APAC, "Refik Anadol announces first open-source generative AI model based on nature", January 15, 2024, https://www.inavateapac.com/news/article/refik-anadol-announces-first-opensource-generative-ai-model-based-on-nature

49 · Shomila Malik, "Large Relationship Models: The AI That's Been Hiding in Plain Sight", Medium, June 17, 2025, https://shomila.medium.com/large-relationship-models-the-ai-thats-been-hiding-in-plain-sight-d9277d266627

50 · Altrata, World Ultra Wealth Report 2025, September 30, 2025, https://altrata.com/reports/world-ultra-wealth-report-2025 (510,810 individuals with a net worth of more than 30 million US dollars).

51 · Mark S. Granovetter, "The Strength of Weak Ties", American Journal of Sociology 78 (1973), no. 6, pp. 1360 to 1380.

52 · Ronald S. Burt, Structural Holes: The Social Structure of Competition, Harvard University Press, Cambridge (MA) 1992.

53 · Toby Stuart, Anointed: The Extraordinary Effects of Social Status in a Winner-Take-Most World, Simon & Schuster, New York 2025.

54 · hiQ Labs, Inc. v. LinkedIn Corp., 31 F.4th 1180 (9th Cir. 2022); after remand, judgment of the U.S. District Court for the Northern District of California of November 4, 2022, largely in favor of LinkedIn; the litigation ended in December 2022 with a settlement. LinkedIn has for years taken legal action against automated scraping of its profiles and restricts access through interfaces to a few partners.

55 · Regulation (EU) 2016/679 (General Data Protection Regulation), especially Article 6(1)(f), Article 9, Article 14 and Article 15.

56 · Regulation (EU) 2024/1689 of the European Parliament and of the Council of June 13, 2024 (AI Act), OJ L of July 12, 2024, especially Article 6(1) and Annex I.

57 · STIHL, "Production and Sales", company website, accessed October 30, 2026, https://corporate.stihl.com/en/about-stihl/production-sales ("available in more than 160 countries across five continents").

58 · Martin-W. Buchenau, „Stihl-Chef: Neue Fertigung in Deutschland? ‚Da fehlt mir derzeit die Fantasie'", interview with Michael Traub, Handelsblatt, October 25, 2026. Quotations: „Stihl hat als Familienunternehmen nicht ständig den Kapitalmarkt im Nacken"; „Nicht die Technologie bestimmt, sondern der Kunde"; on specialist dealers: „Da sind wir fast religiös"; on the replaced US head: „Er wollte ein anderes Spiel spielen als wir."

59 · The term "Brain Trust" was coined in 1932 for the circle of advisors of presidential candidate Franklin D. Roosevelt, according to Raymond Moley by James M. Kieran, reporter for the New York Times; Samuel Rosenman attributed it to Louis Howe (Raymond Moley, Rexford Tugwell, Adolf Berle, all Columbia University). See Raymond Moley, After Seven Years, Harper & Brothers, New York 1939.

60 · David A. Pieper, Old Money, New Machines (working title), white paper in preparation, 2026; see chapters 1.4 (the two halves of the family office), 5.2 (legacy, family constitution) and 6 (the family office as control center).

61 · Yan Tao, Olga Viberg, Ryan S. Baker, René F. Kizilcec, "Cultural bias and cultural alignment of large language models", PNAS Nexus 3 (2024), no. 9, pgae346. On the value scales of the World Values Survey, the answers of the models studied were closest to English-speaking and Protestant European countries.

62 · Teikoku Databank, "全国「老舗企業」分析調査（2025年）" (analysis of long-established companies 2025), March 27, 2026: 46,708 companies with more than a hundred years of history, as of December 2025; announcement reproduced at https://www.zaikei.co.jp/releases/3365307/

63 · KfW Research, „Standort Zukunft: Wie Deutschland zu neuer Stärke und Wachstum findet", chapter 3.4 Nachfolge, 2026, https://www.kfw.de/PDF/Download-Center/Konzernthemen/Research/PDF-Standort-Zukunft/3.4-Nachfolge.pdf

64 · DIHK, „Industrie investiert stärker im Ausland: Kostendruck erreicht historischen Höchststand", special analysis of the DIHK business survey at the start of 2026 (around 1,700 industrial companies), April 16, 2026, https://www.dihk.de/de/newsroom/industrie-investiert-staerker-im-ausland-kostendruck-erreicht-historischen-hoechststand-176268 („In diesem Jahr planen 43 Prozent der Industriebetriebe Investitionen außerhalb von Deutschland"; 41 percent cite cost savings as a motive, the highest figure since 2003).

65 · IG Metall, „Einigung bei Miele: Standorte bis 2028 sicher", June 10, 2024, https://www.igmetall.de/im-betrieb/einigung-bei-miele-standorte-bis-2028-sicher (relocation of washing machine production from Gütersloh to Poland, 700 jobs affected in Gütersloh, 1,300 in Germany overall).

66 · Interfax, "Ukrainian State Border Guard Service provides border-crossing statistics" (with UNHCR data as of February 19, 2026: 5.924 million), April 20, 2026, https://interfax.com/newsroom/top-stories/117217/

67 · TRUMPF, "Minister-President Winfried Kretschmann visits TRUMPF: regional innovation set to power future industrial growth", press release, Ditzingen, January 19, 2022, https://www.trumpf.com/pl_PL/newsroom/globalne-informacje-prasowe/komunikaty-prasowe/release/minister-president-winfried-kretschmann-visits-trumpf-regional-innovation-set-to-power-

future-industrial-growth/ ("ASML uses EUV lasers from TRUMPF to manufacture lithography systems"); ownership structure (Leibinger family) according to Wikipedia, "Trumpf", https://en.wikipedia.org/wiki/Trumpf (secondary source).

68 · Yashvardhan Sharma, "Test and Assembly", chapter 13 in: Chip Supply Chain: How AI Chips Are Made, chipsupplychain.org, edition of October 21, 2026, https://chipsupplychain.org/pdf/test-and-assembly.pdf (according to a CSET analysis of TechInsights data for 2025: Advantest 62.7 percent of the market for chip testers, DISCO 87 percent of the market for dicing and grinding equipment).

69 · International Accounting Standard 38, Intangible Assets, paragraph 63: "Internally generated brands, mastheads, publishing titles, customer lists and items similar in substance shall not be recognised as intangible assets."

70 · Bitkom, „Erstmals nutzt Mehrheit der Unternehmen KI“, press release, September 14, 2026 (603 companies with 20 or more employees; 57 percent use after 36 percent in 2025 and 20 percent in 2024; potential exploited „eher stark“ by 3 percent, „eher wenig“ by 28 percent, „gar nicht“ by 59 percent; 33 percent out of fear of falling behind), https://www.bitkom.org/Presse/Presseinformation/Erstmals-nutzt-Mehrheit-Unternehmen-KI

71 · ifo Institut, „Mehr als die Hälfte der Unternehmen nutzt Künstliche Intelligenz“, June 5, 2026 (54.5 percent, all company sizes), https://www.ifo.de/fakten/2026-06-05/mehr-als-die-haelfte-der-unternehmen-nutzt-kuenstliche-intelligenz

72 · KfW Research, „Künstliche Intelligenz im Mittelstand: einfach zu handhabende Technologien mit breiten Einsatzgebieten überwiegen“, Fokus Volkswirtschaft no. 554, July 29, 2026 (KfW SME Panel, use in 2022 to 2024, surveyed in spring 2025), https://www.kfw.de/PDF/Download-Center/Konzernthemen/Research/PDF-Dokumente-Fokus-Volkswirtschaft/Fokus-2026/Fokus-Nr.-554-Juli-2026-KI_Mittelstand.pdf

73 · McKinsey & Company, "The State of AI in 2026: On the road to ROI", August 2026 (1,719 respondents in 97 countries), https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

74 · Gartner, "Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027", press release, June 25, 2025, https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027

75 · Iñaki Aldasoro, Leonardo Gambacorta, Rozália Pál, Debora Revoltella, Christoph Weiss, Marcin Wolski, "AI adoption, productivity and employment: evidence from European firms", BIS Working Paper no. 1325 and CEPR Discussion Paper no. 21082, 2026, https://www.bis.org/publ/work1325.pdf (EIB Investment Survey 2019 to 2024, more than 12,000 companies; summary: VoxEU, "How AI is affecting productivity and jobs in Europe", https://cepr.org/voxeu/columns/how-ai-affecting-productivity-and-jobs-europe).

76 · Boston Consulting Group, "From Potential to Profit: Closing the AI Impact Gap", January 15, 2025 (more than 1,800 executives), https://www.bcg.com/publications/2025/closing-the-ai-impact-gap

77 · Robert M. Solow, "We'd Better Watch Out", review in: The New York Times Book Review, July 12, 1987, p. 36 ("You can see the computer age everywhere but in the productivity statistics").

78 · Erik Brynjolfsson, Daniel Rock, Chad Syverson, "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies", American Economic Journal: Macroeconomics 13 (2021), no. 1, pp. 333 to 372.

Note on method

All ideas, theses and assessments in this white paper are the author's and have grown over two decades of practice. Artificial intelligence (Claude, Gemini) was used to check prior publications, for referencing, for researching sources, for form and spelling and for text drafts under the author's editorial control. The paper describes a system that itself rests on this collaboration; the author considers it appropriate to say so openly. The AI did not write the content of the paper.

About the author

David A. Pieper is an investor and entrepreneur and has worked in the world of large family fortunes for more than twenty years, at the intersection of family offices, family businesses, capital markets and real estate transactions. He holds an Executive MBA (IREBS) and an LL.M. in Private Wealth Management, has published on foundations and succession, and is the author of the white paper series on momentum, tail risk and

quality investing. He is currently building the Large Human Network Model described in this paper for the world of families, family offices and foundations.