


## 01 Master Data Meets AI: How Original Knowledge Becomes Better Advice



Good afternoon. I am Michael Mohr, founder of plugilo and initiator of the Trusted AI Advisory Alliance. The title was chosen deliberately: Master Data Meets AI. Because the master data these two days are about is meeting a new consumer, one that no longer queries the data but talks about it. In the next few minutes I want to show you why reliable AI advice starts with original data, how GDSN data becomes intelligent Cards, and how knowledge travels from the data owner via retail to the customer. Almost every slide ends with a learning. At the end you will have nine sentences to take with you.

### KNOWLEDGE AND BACKGROUND

GDSN synchronizes product master data via certified data pools. GS1 DQX assesses the quality of product master data against defined verification requirements. After successful verification the data is electronically sealed. The data quality service is not a platform for providing or distributing master data; data and verification results are transmitted within the GDSN. "Training" here refers to providing, selecting and updating the knowledge base for pluggy, not necessarily to changing model weights.

### SOURCES AND REFERENCES

GS1 Germany: [GS1 DQX Services](#)

GS1 Germany: [How GS1 DQX works, version 2.7 \(23 Dec 2025\)](#)

## 02 AI is changing the rules of the game.



Imagine a customer who no longer opens ten websites. He tells his AI: find the right product, explain it to me, compare, advise me, buy and take care of the support. This advice already decides which products make it into his selection at all. The first salesperson is an AI.

This changes the rules of the game for everyone in this room: manufacturers and brands just as much as retail,etailers and resellers. Sales channels have become hybrid, the sources of information countless. The customer asks the AI, and the AI asks everyone.

Three numbers show the direction. In an Adobe survey, 39 percent of people in the US had already used AI for online shopping. Traffic from AI sources to US retail sites grew by 393 percent in the first quarter of 2026 compared with the previous year. And with protocols such as Google's UCP, search and checkout are moving into a single conversation. Not every purchase is autonomous yet. But an answer turns into delegated actions, all the way to purchase and service.

So the decisive question is not: are we visible in the AI? Many are visible, even with old copies. The question is: does the AI recognize our original as the original? In a search engine, a brand could hold its ground with well-maintained pages. In an AI answer, what decides is whose data is recognizable as the original. Being in the AI is not enough. Being recognized as the original is what counts.

### KNOWLEDGE AND BACKGROUND

**What the numbers show.** Adobe published the data on April 16, 2026. The accompanying survey covered more than 5,000 people in the US. 39% had used AI for online shopping. The traffic analysis covers visits to US retail websites. +393% refers to Q1 2026 versus Q1 2025 and is a growth rate, not a share of all purchases or visits.

**What agentic commerce means.** An agent can carry out research, comparison and further delegated steps. Google's Universal Commerce Protocol covers, among other things, product discovery, checkout, account linking and order management. Which actions are possible depends on integration, merchant participation and user consent.

**Context for all data owners.** The statement concerns access to the customer and control over one's own information basis. It claims neither complete AI dominance today nor a certain forecast that every future purchase will go through AI.

### SOURCES AND REFERENCES

[Adobe Digital Insights, 16 Apr 2026](#)

[Google: Universal Commerce Protocol](#)

### 03 Frequency is not accuracy



Take a packet soup. The manufacturer changes an additive and corrects the list of ingredients. The original is up to date. But on the web there are three hundred old copies of that ingredient list: at retailers, on marketplaces, in recipe portals, in comparison services. An AI asked about the ingredients finds the old information three hundred times and the new one once. For someone with an allergy, that is not a detail.

The ratio of one to three hundred is an example, not a measurement. It shows what happens in the market: product information is copied, shortened, translated and reused. A correction at the owner never reaches these versions. And frequent occurrence is no proof of accuracy. Three hundred copies are not three hundred confirmations.

On top of that there is an access problem: Adobe measured an average readability score of 66 percent for product pages in US retail. Part of the existing knowledge is simply not well accessible to AI systems. And the economic dimension: Gartner puts the average cost of poor data quality at a minimum of 12.9 million dollars per organization per year. That figure applies to data quality overall, not to AI alone.

For retail, the consequences are concrete: wrong advice, lost demand, service and returns, in individual cases liability. Correct data in your own system is therefore not enough. The AI believes what is described most often. Not what is really true.

#### KNOWLEDGE AND BACKGROUND

**Copies, training and web retrieval.** The 1:300 ratio is illustrative. The slide does not claim a simple majority vote inside the model. Training corpora can contain duplicates. AI search can additionally retrieve web sources without retraining the model weights. Lee et al. study duplicates and their effects on language models.

**Data access and cost.** Adobe's 66% is an average machine-readability score for content on product pages in US retail. It means neither that 34% of all product pages are missing nor that 34% of product data is wrong. For the average of at least USD 12.9 million per year, Gartner refers to research from 2020. It is not a guaranteed minimum loss for every single company.

**Possible consequences for retail.** Incorrect properties or compatibility information can trigger wrong advice, support effort and returns. Missing or incorrect consideration can cost demand. Whether liability arises depends on the individual case. The slide names risks, not a measured overall loss rate.

**The packet soup example.** The example is illustrative. In the EU, ingredient and allergen information for food is governed by the Food Information to Consumers Regulation (Regulation (EU) No 1169/2011). Outdated copies of an ingredient list are therefore not a marginal problem; they can be directly relevant to consumers.

#### SOURCES AND REFERENCES


Lee et al., ACL 2022: Deduplicating Training Data Makes Language Models Better

Adobe Digital Insights, 16 Apr 2026

Gartner: Data Quality

Regulation (EU) No 1169/2011 (Food Information to Consumers)

## 04 Customers ask the AI first.



We have seen this loss of control before. Years ago, review platforms took the authority over opinion away from brands: it was no longer the manufacturer who told people what a product was like, but the reviews, real and invented alike. Now it is about the authority over the product. It is not the manufacturer who tells the AI what is in the soup, but the sum of the copies. As long as nobody marks the original.

The three numbers describe three aspects. First, usage: among people who use AI for research, 42 percent, according to Adobe, rely on it often or always as their main source for advice and purchases. Second, reliability: in the Tow Center's news test, more than 60 percent of the source citations were wrong, and the systems also cited copies instead of originals. That is not an error rate for product advice, but it shows that source attribution does not work by itself. Third, contact with the source: in Google searches with an AI summary, only one percent of visits led to a click on a cited source.

So the answer takes effect without the customer ever seeing your website. And because answers vary, a single test does not show what all customers receive. The quality of your information has to work where the answer is created. The AI advises your customers. With you or without you.

### KNOWLEDGE AND BACKGROUND

**The comparison with review platforms.** The comparison is an analogy: reviews shifted the authority over the product experience from brands to third parties. With AI answers, something similar happens to the product truth when the original is not recognizable as the original. The comparison does not claim an identical mechanism.

**42%: subgroup of AI research users.** The Adobe Consumer Report 2026 is based on 4,000 respondents worldwide in total. The 42% refers to people who already use AI for research. It means frequent or constant use as the main source for advice, shopping or problem solving, not a share of all buyers.

**More than 60%: a limited news test.** The Tow Center tested eight AI search services with 1,600 source-identification queries on 200 news articles from 20 publishers. The report appeared in March 2025. It also observed references to syndicated or unauthorized copies and rare admissions of uncertainty. The results are not a universal current error rate for all AI applications.

**1%: a precise reference value.** Pew analyzed browsing data from 900 US adults in March 2025. In 1% of visits to Google results pages with an AI summary, a click on a link inside the summary followed. A further 8% led to classic search results; without a summary it was 15%. Sources remain important for the subject matter even when the direct visit does not happen.

**Limited predictability.** Answers depend, among other things, on model, query, context, retrieved sources and point in time. Individual tests do not replace systematic observation.


### SOURCES AND REFERENCES

[Adobe Consumer Report 2026](#)

[Tow Center: AI Search Has a Citation Problem](#)

[Pew Research Center, 22 Jul 2025](#)

## 05 One standard instead of a thousand islands.



This leads to the most important insight of this talk: no manufacturer can solve this alone. You can maintain your data perfectly, and the AI still will not know that your version is the original, as long as there is no shared rule for it. One rule for all: the original is marked as the original, its provenance is verifiable, the approval lies with the owner.

That is exactly how this industry solved identity fifty years ago. Not with a thousand house numbers, but with one GTIN. Every isolated solution needs its own interfaces and remains limited. An open standard creates shared rules that every partner fits into.

The Trusted AI Advisory Alliance is meant to create this framework. Data owners and market partners bring knowledge, reach and requirements together and face the large AI systems jointly. Three principles are set: control over the original knowledge stays with the owner. Provenance, openness and competitive neutrality are traceable. And value flows, under transparent rules, to those who contribute knowledge and reach.

A standard does not emerge because someone proclaims it. It emerges because companies use it and shape it. Alone, everyone remains an island. Together, we set the standard.

### KNOWLEDGE AND BACKGROUND

**Why visibility is not enough.** A manufacturer can maintain its data completely and up to date and still be represented in AI answers by older copies, because the systems do not distinguish between original and transcript by themselves. Recognizability of the original requires a shared rule for marking, provenance and approval. The GTIN serves as an analogy for a standard used by everyone, not as a technical template.

**Goal of the Alliance.** An open, source-based advisory standard is meant to connect isolated solutions and give data owners and market partners a shared framework for action. The graphic is a target picture; it does not claim an already established global standard.

**Control and openness.** Owners control original knowledge and approvals. Provenance is meant to remain visible. The network is meant to be open to reputable AI systems and competitively neutral. Revocations govern permitted future use under the rules; they do not guarantee retroactive deletion of all external copies.

**Fair value creation.** Knowledge and reach are to be shared, and contributors are to participate in the success under transparent rules. The charter combines technical cooperation with responsibility and distribution of the resulting benefit.

**Why now.** Urgency arises from standards development already under way, platform moves and the lock-in of companies into isolated solutions. Joint activation increases the chance of shaping the rules; it does not guarantee market success.

### SOURCES AND REFERENCES

[Trusted AI Advisory Alliance, Master Paper \(theses 17, 19, 20\)](#)

## 06 plugilo Trusted AI Advisory Network



Before we go into the building blocks, one picture for the whole: the plugilo Trusted AI Advisory Network. plugilo is a plug. You plug it into every online activity: your website, your shop, your content, all the way into the offline world at the shelf. And wherever it is plugged in, the same real data and the same advice from the original flow through.

You can see it here: web widget, browser extension and shop plugin bring the brands' knowledge to every website, every shop and every browser. In the middle, a manufacturer's website from a neighboring industry with plugilo plugged in: the Card with the original data, the chat with pluggy, the dock. And the plugilo app is the customer's personal Wallet: the same Cards, the same advice, to take along, connected to the brands he follows. Nothing was maintained twice. It is one network, plugged in once.

### KNOWLEDGE AND BACKGROUND


**The building blocks.** Cards and Stacks: structured knowledge objects and their collections. pluggy: the Trusted AI Advisory Agent. App and Wallet: the personal place where customers keep Cards, results and advice. plug-it: the integration into websites and shops. Browser extension: advice wherever the customer is browsing. Partner Portal: approvals, data and analytics for data owners and partners.

**Context.** The screens show plug-it on a manufacturer website (Card, chat, dock) and the plugilo app. The example comes from a neighboring industry and serves as an illustration. Availability of individual functions depends on integration, approvals and rollout.

### SOURCES AND REFERENCES

[Trusted AI Advisory Alliance, Master Paper \(chapters 9 and 12\)](#)

## 07 One Card. Everything that is true.



The good news: a large part of the data foundation already exists. We have identifiers, classifications, master data, media, documents and verifications. What is missing is the last mile from these distributed building blocks into the advice itself.

The Master Data Card bundles this knowledge into one object. People see a uniform, understandable presentation. AI systems use the same structure by machine. Three layers are decisive. The original core: the properties and statements the data owner is responsible for, in the case of our soup the list of ingredients. The evidence: where does the data come from, who approved it, what was verified, how long is a certificate valid. The context: availability, prices, partner knowledge, each with a recognizable sender. On the right you can see what this looks like: the plugilo Card with its layers, from the image and the data to the description, retailers and advice, with the verification seal.

These layers belong together and remain distinguishable. An owner approval is something different from an independent verification. And context never silently overwrites the original core. Existing standards keep their role: GDSN synchronizes the data, GS1 DQX verifies its quality. plugilo makes this knowledge usable for advice as a connected object, permanently linked to its source. The plugilo Card: the digital original. All product knowledge. Approved by the owner.

### KNOWLEDGE AND BACKGROUND

**The existing building blocks.** Identity and structure: GTIN, classifications and attributes. Original data and media: manufacturer knowledge, PIM/ERP, GDSN and documents. Verifications: data quality, certificates, test bodies and registers. Context: prices, availability, tests and partner knowledge.

**Standardized object and source link.** The Master Data Card is a structured data snippet for a defined object. It is readable for people and usable for machines, comparable, versionable and linked to its source. Live status and updates depend on the source connection and maintenance.

**Authority and evidence.** The manufacturer is responsible for its object data. Test bodies are responsible for their verification statements; market participants for their own information. Owner approval, data quality and external certification are separate pieces of evidence. Issuer, scope and validity must be recognizable. This does not make the entire Card automatically independently certified.

**GS1 DQX and GDSN.** GDSN is the network for data synchronization. GS1 DQX assesses product master data against defined verification rules and awards a quality seal after successful verification. GS1 DQX is not a platform for providing or distributing master data.

### SOURCES AND REFERENCES

GS1 Germany: [GS1 DQX Services](#)

GS1 Germany: [How GS1 DQX works, version 2.7 \(23 Dec 2025\)](#)

## 08 Original knowledge becomes first-class advice.



The Card provides the approved knowledge. pluggy, our Trusted AI Advisory Agent, turns it into first-class advice. The connection on the slide is deliberately drawn as fixed: the Card delivers the knowledge, pluggy pulls out the information that fits the specific question.

A customer rarely wants to read a list of attributes. He wants to know: what can the product do? Does it suit me? And later: how do I use it correctly? Understand, decide, apply. For the packet soup that means: is it vegetarian? Does it contain traces of nuts? How long does it need to steep? pluggy answers from the approved data and shows the source. If a piece of information is missing, he says so instead of inventing it.

Binding every answer to its source makes the basis of that answer traceable. It is not a guarantee that every generated answer is free of errors. What matters is the controlled knowledge base and its correct use. And because Card and advisor stay connected, the customer takes both with him: the result and the possibility to ask further questions. Your best brand ambassador is in every plugilo Card.

### KNOWLEDGE AND BACKGROUND

**Knowledge base.** The Master Data Card provides approved original knowledge, including properties, documents, provenance and verification status. pluggy selects suitable sources for the specific customer question and uses them for the answer.

**Advisory tasks.** Understand: explain properties and functions. Decide: compare requirements and suitability. Apply: support use, service and questions after the purchase. The approach can also cover services, expert knowledge or editorial content, provided the respective sources and rights are available.


**Link between Card and agent.** The Card is the anchor for the object and its knowledge. pluggy is the connected advisor. Place of use and approved customer needs determine which knowledge becomes relevant. Actual availability depends on activated integration and the respective usage context.

**Reliability.** Source binding and controlled information are meant to improve the advice. They do not replace quality control and do not guarantee that every generated answer is free of errors.

### SOURCES AND REFERENCES

Trusted AI Advisory Alliance, Master Paper (thesis 9)

## 09 plugy learns from three sources



What makes this advice strong? plugy learns from three sources that remain incomplete on their own. The first are manufacturers and partners, the data owners: original knowledge about properties, applications and evidence. That is the factual foundation.

The second are resellers or publishers: which assortment is available, which services exist, what matters at this location? That is the context. The third is the customer: what is he looking for, which requirements and restrictions does he have, which questions remain open? That is the need.

An example: the manufacturer knows what is in the soup. The retailer knows which variant is on the shelf and what is on offer today. The customer knows that his daughter cannot tolerate nuts. Good advice brings these three pieces of information together. When we say plugy learns from three sources, we mean exactly this interplay. Approved demand signals sharpen the advice but never change the owner's product properties. That is how an answer becomes well-founded and personally relevant. Brand knowledge, retailer context, customer needs: plugy unites all three.

### KNOWLEDGE AND BACKGROUND

**1. Data owner.** Properties, product benefits, applications and evidence flow in via approved Master Data Cards. The owner is responsible for the original core of its data.

**2. Resellers and partners.** Assortment, availability, services, industry knowledge and experience at the place of use complement the foundation. Partner context does not replace the object data the owner is responsible for.


**3. End user and Demand Trainer.** Questions, personal requirements, wishes, saving and comparing provide demand signals within the scope of approvals. These can help improve the relevance and selection of the advice.

**What “learning” means technically.** It refers to building, selecting and updating the knowledge and context base as well as retrieval and grounding during advice. The statement does not claim automatic continuous training of model weights through every use. Source rights, approvals and the separation of original core and context continue to apply.

### SOURCES AND REFERENCES

Trusted AI Advisory Alliance, Master Paper (thesis 9, chapter 12)

## 10 Wherever the customer is.



Now a single piece of advice becomes a network. On the left is the data owner. He maintains his original knowledge and controls the approvals. Card and pluggy carry this knowledge to the places where it is needed: to the AI systems and agents, to the point of sale in the shop, in the store and directly on the product, and to the point of interest, meaning websites, media and service.

Our goal is an open connection via the available interfaces. That does not mean every AI system is connected today. It means we organize the knowledge base so that suitable systems and partners can use it in a controlled way.

The way back is particularly important. The customer can save information, share it or follow it. You can see it on the right: Stacks that customers follow, and every new follower is a direct connection. Via Cards, Stacks and App, a single use becomes a lasting contact, with the object, with the retailer, with the data owner. A mere view does not create this relationship; it arises through voluntary use and approvals. That is where the strategic value lies: your own customer connections instead of dependence on individual gatekeepers. The original at every touchpoint. On the shelf, in the online shop, in the AI.

### KNOWLEDGE AND BACKGROUND

**Reach of the architecture.** LLMs and AI: access for suitable models and agents. POS: resellers, online shop, store and product. POI: websites, media and service. "All LLMs" describes the goal of an open architecture, not existing integrations with every model.

**Card and pluggy stay connected.** The data object and the connected advisor make original knowledge usable at the respective place of use. Details on provenance, verification status, source and updates remain tied to the object. Connection and reach depend on interfaces, integration and approvals.


**Lasting customer connections.** Voluntary saving, sharing and following can connect data owners, partners and end users via Cards, Stacks and App. Every use can create an opportunity for this; mere views do not create an automatic personal relationship. Demand signals are subject to the respective consents and approvals.

**Strategic benefit.** Own connections can reduce dependence on individual gatekeepers. This is not absolute technical independence from infrastructure or platform partners. The approach applies to all owners of their respective data, including brands, resellers and publishers.

### SOURCES AND REFERENCES

Trusted AI Advisory Alliance, Master Paper (chapter 9, network and channels)

## 11 Brand, reseller and customer stay connected



This picture sums up the journey. The brand trains pluggy with its original knowledge. Retail brings it to the points of decision, in the shop, in the store, directly on the product. The customer decides better and takes results, information and the advisor along in his Wallet.

Value is created beyond the individual purchase. The customer keeps access to your knowledge and comes back for advice with the next question. An answer becomes a relationship, for brand, retailer and customer. The next step is concrete: jointly define which original knowledge is activated first, through which partners it reaches customers, and how we make the usage visible. The customer keeps the original. And stays connected with brand and retailer.

### KNOWLEDGE AND BACKGROUND

**The connection.** The closing graphic links brand or data owner, retail and product, customer and Wallet/App. Original knowledge becomes the basis for advice. Partners create points of use. Customers can keep information and results and call up advice again.


**What should be clarified for the start.** First knowledge source and responsible data owner; concrete use case; reachable partners and channels; approvals and integrations; observation of the first usage. A limited, verifiable start can create the basis for further activation.

**About the picture.** The closing graphic is a concept illustration. Individual depictions are not proof of the current state of every product function. Lasting contacts arise through activated partners and voluntary use such as saving and following.

### SOURCES AND REFERENCES

Contact

## 12 plugilo. The new Standard for Trusted AI Advisory.



plugilo is the new standard for Trusted AI Advisory. Plug it, love it, share it: that is the program. The network for it is being built now, and joining the Trusted AI Advisory Alliance is free and without obligation. Register via the QR code, bring in your knowledge, your partners and your requirements, and together we will define the first concrete use case for your company. Thank you very much. I look forward to your questions, here and at our booth.

### KNOWLEDGE AND BACKGROUND

**Michael Mohr.** Phone: +49 177 7776647. Email: mm@plugilo.com. Contact: <https://corporate.plugilo.com/contact/>. The registration QR code is taken unchanged from the Alliance leaflet.

### SOURCES AND REFERENCES

[Register for the Trusted AI Advisory Alliance](#)