



Simon Hefti

DATA DRIVEN VALUE CREATION

WORKFLOW

Wednesday

Posts 1-8  
January-March  
2025

THEORY

Thursday

TREND


Tuesday



**D ONE**

WE MAKE SENSE.

d-one.ai



## Content

### Trend Tuesday

Disruptive Startups Shaping the Future Today	3
Blood Setpoints: The Overlooked Path to Personalized Medicine	6
Trend Tuesday: Could AI Catch a Fly?	9

### Theory Thursday


When the Unexplainable Becomes Explained	5
The Future of AI: Here, but Unevenly Distributed	8

### Workflow Wednesday

The Modern Data & AI Platform – A Game-Changer	4
Did you learn to swim without water?	7
Do You Still Sketch Your Data Pipelines?	10

Contacts / About D ONE	11
------------------------	----

*The wonderful art as always  
by @basilonmypizza  
<https://basilhefti.ch/>*



## Trend Tuesday: Disruptive Startups Shaping the Future Today

TechCrunch's list of disruptive startups for 2024 puts a spotlight on innovation. Here are topics which particularly resonate with me - proof of what's possible when data, AI, and brilliant engineering come together:

### Healthcare

Abridge transforms doctor-patient conversations into structured notes, helping reduce IT overload for doctors and improve care. As Eric Topol says, freeing doctors from keyboards can slash healthcare costs and boost outcomes. A simple but powerful use of AI in medicine!

### Energy Innovations

Sila is revolutionizing battery chemistry for a greener future, while KoBold Metals (I like the witty name) uses AI to discover critical minerals powering renewables. Two centuries after Volta's first battery, there's still untapped potential for better energy density and stability. We need this progress quickly, and I am happy and grateful to see these technologies advance.

### Logistics & Trucking

Waabi's AI platform aims to solve labor shortages in trucking, a global challenge. As the Economist noted, industries are struggling to fill roles as older workers retire and younger ones avoid tough conditions. By the way, many other industries find themselves in similar situations and turn to AI to address these challenges. Waabi's multidisciplinary AI approach offers a smarter, safer path forward.

### Ultra-Fast Delivery

Zepto delivers groceries in under 10 minutes in India—blending speed, convenience, and sustainability. Localized supply chains, electric bikes, and fewer consumer trips highlight the environmental promise of fast delivery.

As William Gibson said: “The future is already here—it's just not evenly distributed.” TechCrunch's list reminds us what's possible today when bold thinking meets technology. Which startup on this list excites you the most? Let me know below!

See [https://lnkd.in/edTC\\_DC](https://lnkd.in/edTC_DC)










#2

## Workflow Wednesday: The Modern Data & AI Platform – A Game-Changer

Companies recognize the pivotal role of data and AI in driving success. This realization challenges data systems to process more sources, meet rising expectations, and deliver faster decisions.


The key? Balancing stability for consistent, reliable outcomes with agility to adapt and innovate.

The solution lies in a platform built on four focus areas:

-  Lab – Where data & ML pipelines take shape.
-  Fab – Where they run reliably, 24/7, at scale.
-  AI/ML Service Layer – Providing AI and ML functions to both Lab and Fab.
-  Governance Layer – Guaranteeing structure, security, and quality.
-  Think pharma: the Lab develops new drugs, while the Fab ensures their safe production.

What were once separate systems, split by technological limitations, now merge into a single platform. Here, data engineers and scientists collaborate seamlessly on both unstructured data (e.g., documents, images, audio) and structured data (e.g., tables).

This isn't just a new system—it's a new way of working, requiring shifts in workflows, governance, and team structures. While it may seem like a recipe for chaos, it's not: the modern data & AI platform is a highly governed, scalable system designed to drive results.

 To design for success, focus on four dimensions:

- Stability & Agility
- Cutting-edge AI, including Generative AI
- Extensibility
- Interconnectable data—whether structured or not-yet-structured

The payoff? Faster onboarding, innovation at scale, and a competitive edge.

How is your data platform evolving to meet these demands? Let's discuss below!



## Theory Thursday: When the Unexplainable Becomes Explained

For centuries, many believed the Earth was fixed, continents never shifted, and diseases came from “bad air.” Then science stepped in—with simple observations like falling apples, fossil records, and microscopic bacteria—to set the record straight.

A modern example is NASA’s Parker Solar Probe. It’s flying closer to the Sun than ever (unlike Icarus, it’s doing just fine) to uncover hidden truths.

I’ve long been captivated by the solar wind. In 1998, Thomas and I used ion charge states as thermometers to estimate temperatures. It’s fascinating how straightforward tools can reveal so much:

- \* Charge states “freeze in” as the solar wind expands, preserving a snapshot of solar temperatures.
- ☕ The first solar wind measurements? Just a metal cup collecting ions. The resulting current indicated ion density.
- 🤖 In space, bandwidth and data quality are limited, so we apply AI and machine learning to detect hidden structures—just as we do in business, finance, and healthcare.

With each new dataset from the Parker Solar Probe, we gain deeper insight into how the solar wind forms. I’m eager to see the breakthroughs and real-world applications that follow.

💡 The big lesson?  
What feels impossible today often becomes routine tomorrow—when we keep asking questions.

What’s a once-“impossible” breakthrough that now drives your industry? Let’s discuss. 📌

See <https://lnkd.in/d5FwyerP> for details.






## Trend Tuesday

### Blood Setpoints: The Overlooked Path to Personalized Medicine




Recently, I strapped on the same glucose monitor people with diabetes rely on, all to learn more about my food patterns. What surprised me: how stable my blood sugar remained—even overnight. I assumed it would steadily drop in the wee hours, but our body is great at tapping into reserves.

It turns out blood sugar and body temperature aren't the only strictly regulated markers. A new study in Nature (Nature volume 637, pages 430–438, 2025, <https://lnkd.in/eUA9f9VC>) and insights from Eric Topol (<https://lnkd.in/ePyJkSBT>) highlight how blood composition also runs on individual setpoints—stable levels that bounce back even after infections or injuries.

#### Why does this matter?

-  Blood tests are among the most frequent medical analyses, often one per person per year in many countries.
-  Individual setpoints are more precise than the one-size-fits-all ranges currently used in diagnostics.
-  By relying on population-based reference values, we're missing crucial personal baselines.


#### Next steps:

-  Define individual setpoints from existing lab data to improve accuracy.
-  Use AI to flag early deviations from your normal, much like wearables do with heart rate changes.
-  Integrate this into clinical workflows for more precise monitoring—no extra tests needed

Of course, population averages have their place, and we need secure, interoperable systems to handle all this personalized data without causing undue alarm. Implemented thoughtfully, these individual setpoints will enhance everyday healthcare.

As someone who's thrilled to work with this kind of data, I see a practical path toward individualized medicine that's both simple and cost-effective. Next time you have blood work done, consider asking about your personal baseline—it could be a game-changer for your health.

Now, if only my body weight would be regulated as strictly...

-  What trend in health and AI is inspiring you the most right now?



## Workflow Wednesday: Did you learn to swim without water? 🏊

Building data solutions without real data is just as impossible!

For years, IT has relied on the Dev → Test → Preprod → Prod model. While this works for classic software, it falls short for data and AI. Test data rarely captures the real-world complexity created daily in fast-paced, unpredictable business processes. Only live data reveals the true patterns needed to build robust pipelines and train accurate models. 📊

The solution is simple: give data teams a safe, agile environment - the Lab - to work directly with live data without jeopardizing production. The Lab mirrors the “real” environment, the Fab, but is optimized for rapid experimentation. 🚀

### In a Nutshell:

Lab:

- Enables quick onboarding of new data, libraries, and AI models 🔍.
- Think of it as the research phase—where experiments run, failures are expected, and learning is constant. 💡

Fab:

- Operates stable, scalable pipelines across the organization 🗝️.
- Changes here undergo tight controls and 4-eye reviews for production-grade stability.

Both environments share the same underlying structure (imagine two identical Databricks workspaces) but have different access controls:

- Data Visibility: The Lab can view Fab data, but not vice versa 👁️.
- Security: Both enforce fine-grained measures like row-level security and masking 🛡️.
- Consumption: Only the Fab serves production data products; Lab data remains isolated from live processes 🔗.

Once a pipeline or model proves successful in the Lab, it transitions to the Fab via CI/CD. Keeping Lab and Fab structurally similar (often differing only by schema name) makes this handoff seamless. With robust modeling practices (e.g., Data Vault), updates roll out continuously without breaking existing processes. 🔄

We follow the “Up or Out” principle: if a pipeline or model succeeds, it moves up to Fab 🚀; if not, it’s cleaned out and deleted ❌. And no production process depends on Lab data.

Lab & Fab offer the best of both worlds: fast, real-data experimentation coupled with a stable production environment. That’s how we unlock the true potential of data and AI - and that’s how I like my swimming! 🏊

How does your team balance innovation and stability? Let’s discuss!

More on the modern data and AI platform:  
<https://lnkd.in/eunKKhgK>



More on AI model deployment:  
<https://lnkd.in/eFJedckv>



## Theory Thursday - The Future of AI: Here, but Unevenly Distributed

“The future is already here—it’s just not evenly distributed.” – William Gibson

Everett Rogers’ Diffusion of Innovations Theory (1962) reminds us that breakthroughs spread in stages—early adopters get the benefits first while many lag behind. Today, three vivid examples highlight this uneven progress:

### DeepSeek’s Spike-Based Processing (Mixture of Experts)

DeepSeek takes inspiration from biology by activating only the “experts” needed for a task—similar to how neurons fire selectively. This approach, echoing S.B. Laughlin’s 1981 findings in Vision Research, maximizes efficiency and cuts costs. It’s a reminder that sometimes, focused intelligence beats sheer power.

### Battery Innovation for a Solar-Powered Future

Utility-scale battery farms like those at Moss Landing (California) and Hornsdale (Australia) store solar energy to stabilize grids during peak demand. While these projects deliver cleaner air and enhanced energy security, they’re still confined to a few pioneering regions—leaving many areas dependent on fossil fuels.

### ETH Zurich’s Breakthrough in Breast Cancer Metastases




Researchers at ETH Zurich have discovered a way to eliminate circulating breast cancer cells before they form metastases by “tricking” them into self-destruction. Although this promising therapy faces regulatory and R&D challenges, it could revolutionize cancer treatment

—showing how life-saving innovations take long before reaching the masses.

### What Does This Mean for AI & Innovation?

Innovation is accelerating - and its benefits often arrive in pockets. While cutting-edge labs like DeepSeek push the boundaries, many organizations still grapple with basic AI adoption due to mindset, limited resources, outdated infrastructure, or regulatory hurdles.

#### Key Takeaways:

-  Collaboration is key: Cross-disciplinary efforts spark breakthroughs
-  Efficiency wins: Moving from “more compute” to “smarter compute” is essential
-  Act now: With top institutions like ETH Zurich, Switzerland leads—if we bridge the gap from lab to market.

Not everyone sees the future at the same time—but together, we can close the gap. How do you see these innovations reshaping our world? Drop a comment or send me a message—I’d love to hear your thoughts!

The D ONE x-mas card talks about how big ideas take time: <https://lnkd.in/eJbXVbNE>





## Trend Tuesday: Could AI Catch a Fly?

Data-driven banking describes the strategic use of data, machine learning and artificial intelligence for driving sales and marketing, improving customer experience, streamlining operational efforts, managing risks more effectively, and boosting product and service innovation.

### Are we building intelligence the wrong way?

I remember the first time I discovered SQLite—a tiny, self-contained database that runs in under 1MB. It wasn't just small; it was powerful. It made me question why so much software is bloated. More is not always better.

The same principle applies to AI. DeepSeek's Mixture of Experts is already proving that activating only the necessary parameters saves energy without sacrificing performance—much like how (biological) neurons fire in bursts rather than all at once.

### Tiny but Mighty: Efficiency Wins




Smarter design has outperformed brute force before:

- ✓ Micro-Max – A chess engine with 2000 ELO in just 1,337 bytes.
- ✓ smollm – A pip-installable tiny LLM that gets the job done.
- ✓ LZ4 Compression – Super-fast, super-lightweight.

Also musicians know: Minimalism isn't weakness—it's mastery.


### AI's Future: Smarter, Not Bigger

Scaling up has worked, but we're already hitting energy constraints. The future belongs to models that achieve more with less. We don't need ever-larger models; we need more efficient architectures:

-  Sparse, modular AI that activates only what's needed.
-  Neural architectures inspired by biological brains.
-  Specialized hardware optimized for efficiency (like the experimental neuromorphic chips).

Of course, AI and a fly's brain serve different purposes, but nature proves intelligence doesn't require massive energy consumption. Instead of chasing scale, let's chase elegance.

The AI race will soon shift from "who's bigger?" to "who's smarter?"

-  Where else do you see efficiency making the difference?

*The fly brain has been mapped, btw, see Nature volume 634, pages 139–152 (2024)*  
<https://lnkd.in/egmpVcxt>

*I met R. D. Hipp of sqlite at the Tcl conference 2008 in Vancouver,*  
<https://lnkd.in/ebDr2YXD>



## Workflow Wednesday Do You Still Sketch Your Data Pipelines?

### *Code-Driven Workflows Scale, Diagrams Don't*

A photo of a crowded event—how many people are in it? A phone call transcript—who spoke when? Not long ago, these were “unstructured” data points. Today, AI extracts faces from images and timestamps from conversations, turning them into structured datasets. The line between structured and unstructured data has blurred, yet many teams still draw their data transformations in visual tools—then jump to a “free-hand SQL” section when complexity strikes.

Instead of drawing workflows, modern platforms treat transformations as code, often generated dynamically to handle ever-changing data. Turning photos into lists of people or transcripts into structured records isn't a one-off task—it's ongoing processing for all incoming data, so insights flow continuously into analytics and AI. Ten years ago, we saw that dependency management was key to scaling. With no off-the-shelf solution, we built our own code-based pipelines. Today, tools like dbt provide these features out of the box: transformations as code, a clear dependency graph, and support for both SQL and Python.

Code-driven pipelines also integrate seamlessly with CI/CD, data testing, and lineage tracking, ensuring every change is tested, documented, and deployed. Mini-batches keep data fresh without the overhead of full-streaming, while automated tests guard against silent failures. It all adds up to a unified, AI-ready data environment that's engineered rather than hand-sketched.

So, are you drawing your data workflows—or engineering them? Drop a comment and share what excites—or worries—you most about this shift. 🙌



## Simon Hefti

Physicist, Data Scientist,  
Co-Founder of D ONE.  
Lecturer and Bassist.

Simon is a distinguished expert in data-driven value creation with a robust background in statistics, data analysis, and software development across various sectors. Known as a pioneer and an actual rocket scientist, Simon's career has been marked by leading high-impact projects, including groundbreaking mobile banking solutions and the largest Point of Sale systems in Switzerland. His journey began with a PhD in Physics, followed by a postdoctoral position at the University of Michigan focused on space research.

As the co-founder of D ONE, the leading Swiss company in data-driven value creation, Simon serves a broad range of national and international clients. His innovative approach and leadership also led to the launch of herlock.ai, an AI tool transforming the legal industry.

## About D ONE

D ONE is the leading Swiss company for data, machine learning and artificial intelligence with national and international customers and a team of experts for data driven value creation. D ONE designs and implements projects which create value from data, acts as a guide on the journey to a data driven enterprise, and helps companies to shape processes, organizational structure, and company culture, leveraging the hands-on know-how along the entire value chain.



Contact:

mail: [simon.hefti@d-one.ai](mailto:simon.hefti@d-one.ai)

web: [simonhefti.com](http://simonhefti.com)

LinkedIn: [www.linkedin.com/in/simon-hefti](https://www.linkedin.com/in/simon-hefti)



**D ONE**  
WE MAKE SENSE.

[d-one.ai](http://d-one.ai)