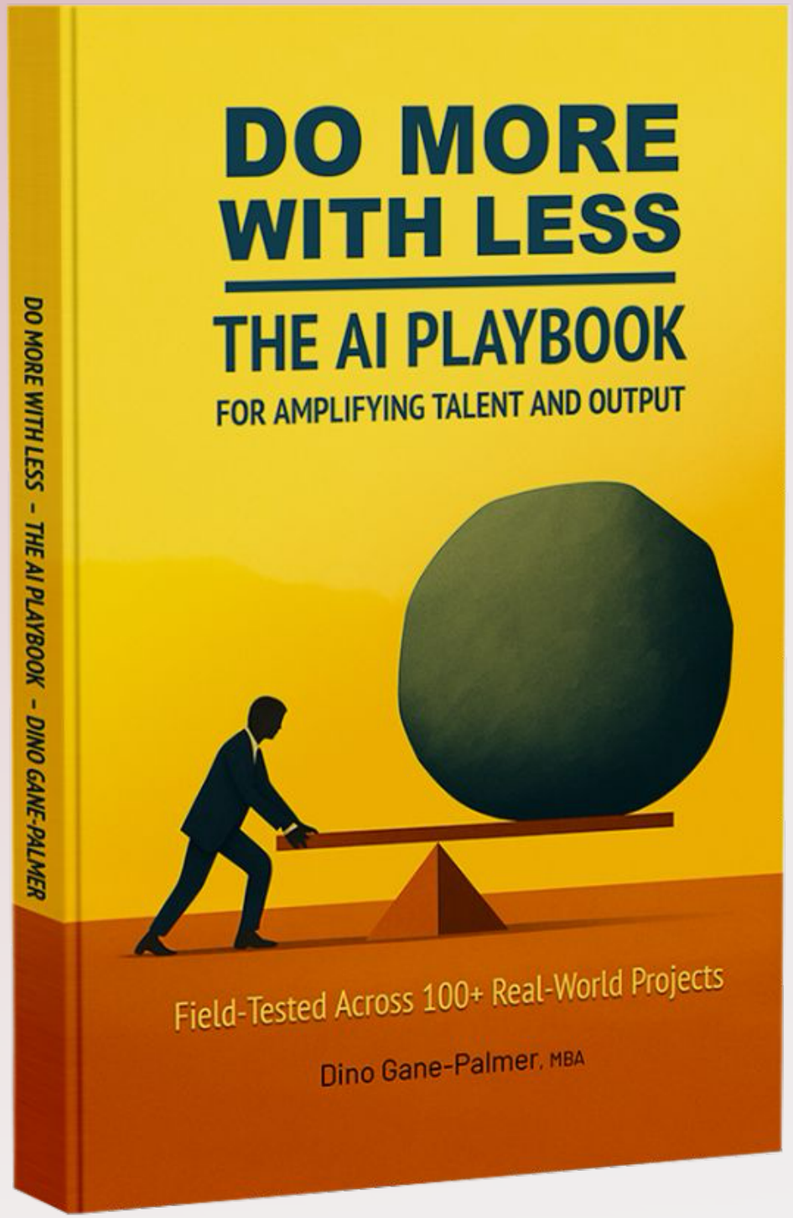

DO MORE WITH LESS

THE AI PLAYBOOK

FOR AMPLIFYING TALENT AND OUTPUT

The Book Summary

Dino Gane-Palmer, MBA

THE UNDISCOVERED OPPORTUNITY

THE FUTURE

HOW TO BUILD THE FUTURE

STEP 1: IDENTIFY

STEP 2: INVESTIGATE

STEP 3: INSTRUCT AI

STEP 4: INTRODUCE

THE UNDISCOVERED OPPORTUNITY

Why organizations have struggled with Generative AI.

Recurring intellectual workflows - such as drafting project reports, monitoring competitor pricing, or triaging customer complaints - consume hours of staff time each week. These mind-numbing activities are exactly the kinds of tasks Generative AI should be able to alleviate. Yet many organizations struggle to get Generative AI working for use cases such as these, where time and money is already being spent and the potential for tangible ROI best exists. Instead, Generative AI use in organizations is largely limited to the use of chatbots.

In the meanwhile, every day, we see headlines about astounding new AI capabilities. From mimicking human conversation to generating code and videos - these new developments seem miraculous and hard to keep up with.

But at its core, all Generative AI capabilities are powered by a simple concept: autocomplete. Think of how Gmail offers you suggested completions, with the press of tab key accepting the completion. That's Generative AI in action.

The underlying technology for these completions - the *Transformer* algorithm discovered in 2017 - has simply been scaled over the past few years to auto-complete a wider variety of data types, in larger quantities at faster speeds.

Generative AI models have gone from auto-completing text conversations, to auto-completing "thinking" steps before providing a response to auto-completing "tool calls" to run web searches and run code before providing a response.

We can expect in the future to have our heads spun by Generative AI models autocompleting large, sophisticated software applications, full movies, 3D worlds, and much, much more.

Nevertheless, **there have been no fundamental breakthroughs since 2017**.

Ultimately, these Generative AI models are taking some input and autocompleting vast amounts of output that traditionally may have taken even months or years to create.

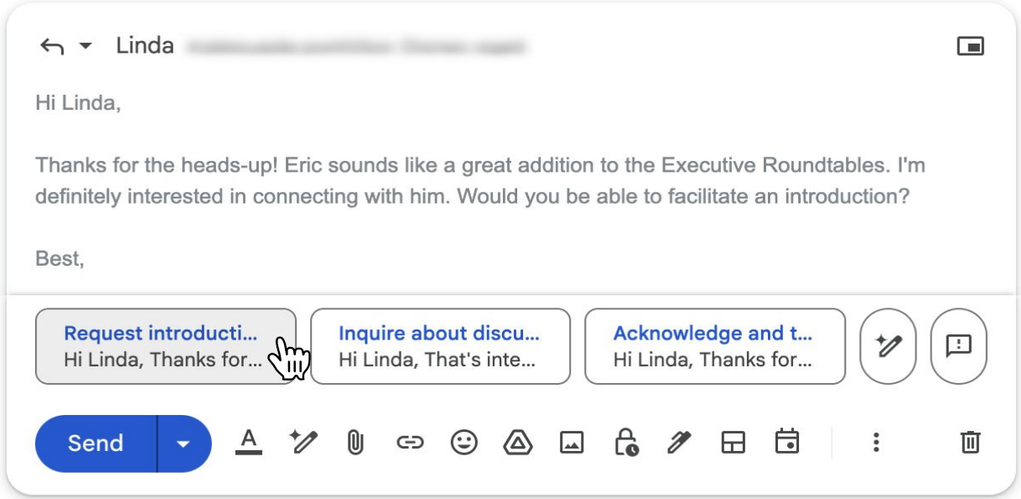

Above: The GMail autocomplete suggestion - is one of the simplest forms of Generative AI in action.

THE UNDISCOVERED OPPORTUNITY

The same Transformer architecture from 2017 continues to be scaled and fine-tuned, but the technology's core limitations - limitations that matter deeply in real-world business settings - remain unaddressed:

- Generative AI models only develop true "understanding" during multi-month training cycles, unlike humans who learn in real time.
- They rely exclusively on public data, leaving them blind to the nuanced, undocumented expertise within organizations.
- Their outputs are non-deterministic, often varying unpredictably even when given the same prompt.

Most demos are amazing because there's no real context required. When you ask a model to generate an image of a cat astronaut riding a unicorn, there are no constraints. It's a creative autofill.

In the real world, business problems don't live in a vacuum. They require deep context: background documents, historical data, constraints, goals, formats, and organizational nuance.

Almost no progress has been made in systematically shaping the context and inputs provided to the Generative AI models for them to easily take on real world work tasks.

Indeed, the pursuit by big tech companies towards shaping this technology as a new species[1] - through chatbots and autonomous agents - has distracted everyone from the truth: the entire value of the output of a Generative AI model is dependent on the inputs provided to the model.

What is needed, to apply this technology to real-world problems, is then to build systems that shape the input, so the model has everything it needs to "autocomplete" that use case successfully. Thus far, building such systems has been left to the domain of software development teams, but this need not be the case. Anyone with Excel can build a system to which they can offload several hours of recurring work each week.

Do More With Less is about how to build such systems.

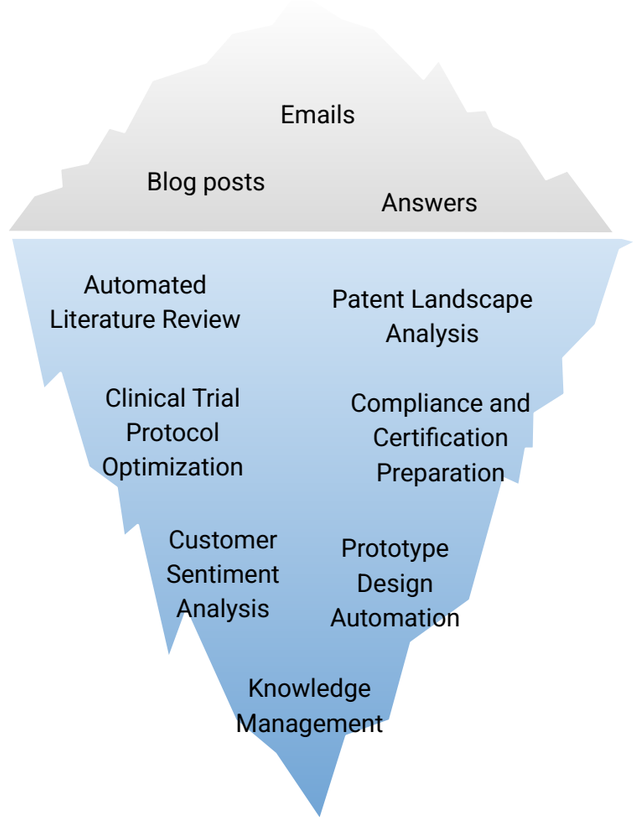

Above: For R&D in large product based companies, there are many "below the surface" applications for Generative AI.

[1] What is an AI anyway? | Mustafa Suleyman
https://www.ted.com/talks/mustafa_suleyman_what_is_an_ai_anyway

THE FUTURE

Designing the work, not doing it.

There are inherent shortcomings to Generative AI as replacements for humans. Despite the ongoing barrage of new releases and announcements, progress by the major AI companies against these shortcomings is either very slow or non-existent. These shortcomings include that Generative AI models:

- do not have access to the 99% of tribal and tacit knowledge humans have in their head and continue to develop offline, ruminating on a topic.
- have limited ways in which they take in information, i.e. just a chat input box, comparative to the always-on rich senses of a human.
- do not reflect and improve on their understanding of the work they do in each task; their "memory" is limited to basic recall of factual information

This means that humans still have a role to play in compensating for the shortcomings of Generative AI, while leveraging its strengths in rapidly creating "completions" - whether they be essays, code or other outputs.

AI, then, provides a way to automate and accelerate more routine aspects of our jobs - in much the same way past technological advancements have.

The invention of spreadsheets shifted jobs away from the mechanical work of summing individual numbers to designing the automations - writing formulas, e.g. SUM(A1:A10). Similarly the challenge, and opportunity, for all of us is to become the designers of work rather than the do-ers of work - delegating the "doing" to AI systems. This shifts jobs towards shaping the instructions - the prompts and other information - that are fed into the AI model, and away from following the instructions ourselves.

Because the AI is not able to learn and improve from task to task, as experts of the processes we delegate to the AI, we then must then focus our attention continuing to improve the how the AI system works, as the world around it continues to evolve.

Rather than be the baker that bakes each bread, jobs now evolve to designing and continuing to evolve the bakery.

Illustrating the impact of "designing the work"

For the following responsibilities of an R&D scientist at a healthcare company we can illustrate the role the human role plays as the expert designer of the work, delegating much of the "doing" to AI:

Responsibility	Responsibility, With AI
Design experiments to explore new compounds, formulations, or biological mechanisms.	AI creates experiment designs, based on details provided in prompts by the R&D scientist using their expertise and knowledge.
Collect, analyze, and interpret complex data from experiments.	AI analyzes and interprets complex data, based on direction of R&D scientist, who learns from the AI's output.
Write detailed lab reports and technical documentation.	AI writes first draft of lab reports and technical documentation, from experimental data. R&D scientist edits and improves the output.
Stay current with scientific advancements by review scientific literature and attending conferences.	AI helps find and synthesise the most relevant content, to facilitate the R&D scientist's learning.

THE FUTURE

Expert in the loop.

In the past, the output of an expert was limited by their time. Now, partnered with an AI system, the expert's output is no longer shackled by their time. As a consequence of this:

- It is no longer necessary for organizations to have as many individuals performing the same role. It is now possible to **scale the work of experts** in that role by partnering them with a AI system that can produce the quantity of output of the entire team.
    - One of the first use-cases to be impacted, at scale, by Generative AI has been customer service. Particularly for online chat based customer service, vendors such as Intercom have allowed organizations to replace teams of customer service staff with AI systems. While the work of handling each customer service conversation has been replaced by an AI, the organizations that have implemented these systems continue to have experts. These experts configure the system for the organization - designing the work of these chatbots - so their knowledge bases continue to be updated. These experts are deeply familiar with the organization's procedures for customer services - better than is documented in the organization's manuals. They also audit the chat logs to check the quality of the work, making adjustments to the system as necessary.
- We can expect an **explosion of experts** as processes and topics that were previously economically unviable become viable with the use of AI systems.
    - We've seen this kind of transformation before: when the Internet drove the cost of content distribution to near zero. It broke the grip of traditional gatekeepers and enabled a wave of niche creators, from MKBHD to Lex Friedman, to reach global audiences. Today, a similar shift is happening with intellectual labor. As AI lowers the cost of tasks like analysis, writing, and decision support, it paves the way for niche super-experts to thrive - individuals who repeatedly solve specific business problems across clients with precision, scale, and minimal overhead. The teams behind AI Meeting Notetaker applications, which largely funnel meeting transcripts into AI models with prompts to create summaries and the like, are one such early version of such "super experts".

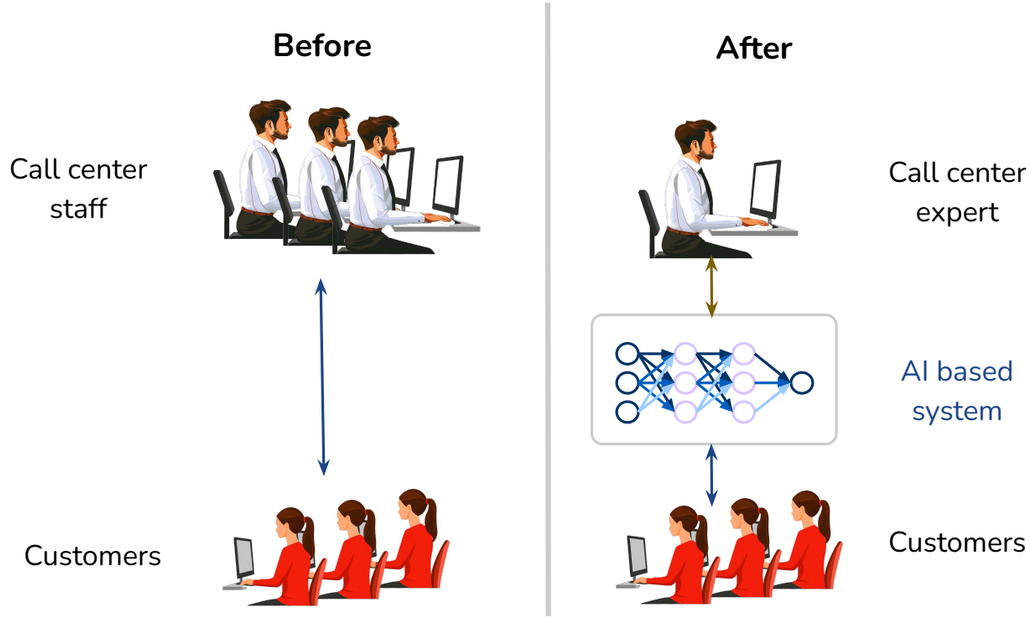

Above: The impact of experts partnered with an AI system.

HOW TO BUILD THE FUTURE

The 8 Hour AI Playbook.

While Generative AI models can produce impressive demos, integrating them into real operational workflows requires serious work. It's not just about the model - it's about the system around it. That system must feed the AI with the right context and knowledge, using prompts refined through iteration, and include infrastructure to guide, govern, and scale its use.

The 8 Hour AI Playbook distills what we, at PreScouter, have learned from working on hundreds of client projects each year. Our cross-industry perspective gives us a unique vantage point to see how Generative AI can be applied to real business use cases. The 8-Hour AI Playbook distills these insights into a step-by-step playbook for successful Generative AI implementation:

- **Identify**: Pinpoint high-friction, recurring workflows where employees spend time on manual research, synthesis, or repetitive work. These are prime candidates for AI augmentation.
- **Investigate**: Understand the task deeply by studying how people currently perform it, what information they use, and what outputs are valued. This insight is critical to guiding how the AI system is built.
- **Instruct AI**: Design prompts and data flows that guide the AI to perform the task. Conduct testing to ensure the AI reliably mimics the thinking process of the expert whose work it is amplifying.
- **Introduce**: Deploy the solution into daily work through onboarding, SOP updates, and training. Ensure teams understand when and how to use the AI tool, and track performance to drive adoption and improvement.

The 8 Hour AI Playbook serves multiple functions: it's a planning tool, a communication device, a progress tracker and a knowledge repository. It also provides a structured path for teams to confidently navigate building use-case specific AI systems together.

As a single page, the Playbook provides simplicity, but you will likely add more elaborate details in other documents. You can find an editable copy at DoMoreWithLess.now.

The 8 Hour AI Playbook

Capture your progress and get everyone on the same page.

ROLES

	Champion	IT	Expert Staff	Prompt Engineer	Other

0h

Identify

Target Process	Potential Impact	Use Case Checklist	Stakeholders
Which processes are worth exploring? (e.g., heavy paperwork, variable work product quality)	Impact of.. • *Automating*: Eliminating the repetitive/boring • *Upgrading*: Work outputs that are "wow"	☐ Routine? ☐ Voluminuous? ☐ Data available? ☐ Exciting ROI?	List key stakeholders who need to buy-in. How can you address concerns they might have?

2h

Investigate

Task Breakdown	What's Complex?	Solution Approach	Data Flow
List out the steps involved in the process.	What is difficult? Are any workarounds or frustrations involved?	☐ Are there off-the-shelf applications available? ☐ Are APIs or integrations available?	Where does the data come from and where does it go? What systems need to integrate with AI?

4h

Instruct AI

Prompt Template	Examples	Checks	Expected Output
Write out the prompt(s), including variables for incoming data and steps the AI needs to take.	Provide examples of the types of inputs the AI will get, as well as outputs it would provide for those.	☐ Passes "New Colleague" Test? ☐ Are there applicable roles ☐ Guardrails for improving accuracy?	How should the output be structured so it can be used in the next step of the business process?

6h

Introduce

Workflow Changes	Human Impact	Monitoring	Success Metrics
How do workflows need to be redesigned?	How will roles be redefined? What resistance to change might you face?	☐ Staff trained? ☐ Staff supported? ☐ Risks managed?	List key success metrics

8h

HOW TO BUILD THE FUTURE

How your AI system will work.

If there is anything you should take away from this book, it this:

How successful the AI is for a use case is entirely dependent on the context it is given - the input information that it is fed for a task.

Because the system essentially passes in messages to an AI model, in theory you can simply build on ChatGPT, Copilot or a similar chatbot for your AI system. For prototyping, this is sufficient. However, you will likely need something more sophisticated to run the prompts on a recurring, bulk basis for the recurring work you target. The simplest version of such a system is one based on Excel, with each row representing each iteration of the task. At DoMoreWithLess.now you can find a spreadsheet ready to plug into an LLM model.

For the use cases we are building a system for, we are constructing input prompts that will likely be several pages long. These pages will be sent as a message to the AI model, with each message corresponding to one task or “run” of the use case. These inputs consist of different parts, shown below.

1. For a specific use case, what is the manual of instructions for how the work should be done by the AI? For example, for a competitive monitoring use case, this could be instructions on how to extract and evaluate information from articles that are fed to the AI model, as part of the task. This manual is maintained by the expert and passed to the AI model with every message.

2. Beyond the general use case level information, how are the messages sent to the AI for each task triggered? Is it from a user input or an event, such as an email arriving in an inbox? Alternatively, it could be that the task runs on a recurring cadence, such as for a weekly report. All the relevant information for the task needs to be gathered and provided in the message.

3. As well as the use case level instructions and task specific information, many tasks also require further information from existing applications. For the use case of drafting a project plan, for example, the project details provided as part of the task could be used to retrieve past similar project plans from a database. These past plans could then be provided to the AI with all the other information, so it shapes its output closer to past organizational norms.

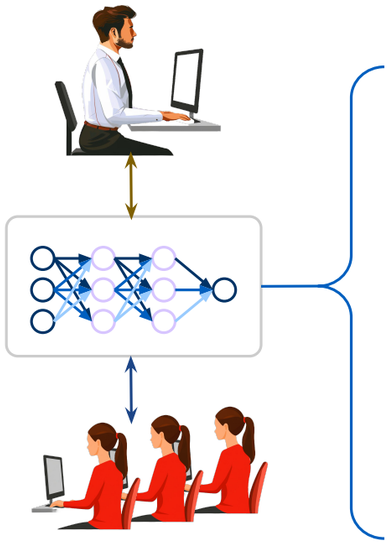

Platforms for building in-house use cases

GPT for Sheets / Excel - plugin for making bulk calls to AI models from spreadsheet cells.

Zapier - connect various data sources to your CRM or other applications

Replit or Loveable - create applications using text instructions

Copilot studio - create rule-based in-house automations and knowledge bases.

Auxee - makes building secure, in-house AI based systems easy for layma

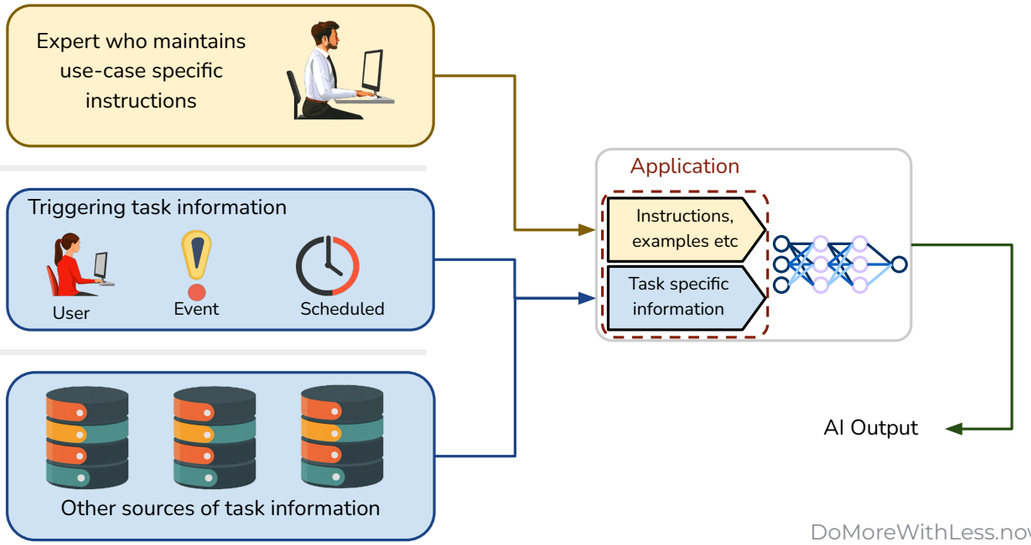

STEP 1: IDENTIFY

Target recurring work.

All job roles consist of ad-hoc and recurring activities. The work, team or department where AI can make the biggest impact is where there is the greatest intensity of recurring work. While AI is beneficial for ad-hoc work, its impact is not as measurable. ChatGPT, Copilot or other chatbots are likely already helping staff with ad-hoc work activities.

Recurring work comes in two flavors:

- Repetitive, routine intellectual work, which can be automated with AI. These use cases drive productivity and have the highest ROI potential.
- Work that produces digital work products, which can be upgraded in quality, by augmenting staff with AI.

Repetitive, routine intellectual work activity examples		
Purchase order processing	Support ticket prioritization	Highest opportunity areas
Data entry validation	Contract compliance checks	
Loan application reviews	Regulatory filing preparation	

Digital work products producing activity examples		
Sales proposals	Marketing content	Secondary opportunity areas
Training materials	Legal document drafting	
Regulatory documents	Business plans	

Ad-hoc work activity examples		
One-off strategic decisions	System outage resolution	Lower priority areas
Special event planning	Stakeholder conflict resolution	
Project crisis management	Exception handling	

Narrow in on the best use case

As you look across your organisation for opportunities for where to apply AI, consider the following criteria:

Voluminous. Think of voluminous tasks as the digital equivalent of a never-ending assembly line. Just when you think you've processed the last item, a new batch appears, like some cruel business version of whack-a-mole. It's the corporate world's version of Groundhog Day, minus Bill Murray's charm.

Routine. If you can write a detailed how-to guide for the task, chances are it's routine enough for AI to handle. And if you find yourself thinking, "I've explained this a hundred times," - well, it might be time to explain it once to the AI and let it continue doing the work for you.

“Data Available” means your information isn't just lying around in dusty file cabinets or living rent-free in your top performer's head. It's digitized, organized, and ready for AI consumption - or can be. Once the AI has received the data, where will the completed work be sent? Will it be a database, social media website or will the AI perform some action - such as sending an invoice?

Wow Factor. If explaining the potential of your AI use case doesn't make your eyes light up like a kid on Christmas morning, it might not be the right fit. The right AI application should make you feel like you've just discovered electricity - and you can't wait to light up your entire business. Why? Because the ROI is self-evident in a 5x to 10x improvement in speed, quality, cost or other characteristics. AI makes this level of step change possible, and that is what we are looking for. Once you find a use case for which AI makes a step change possible, precise ROI calculations typically become unnecessary.

STEP 2: INVESTIGATE

Understand how the work is done.

So you can understand how to model the work using AI, conduct user research sessions with the staff currently performing for work activity for the use case chosen from *Step 1: Identify*.

1. Screen Sharing & Task Walkthrough

Ask the staff member to share their screen and walk through each task on the task list related to the use case. As they proceed:

- Observe closely: Watch not just what they do, but how they do it - mouse movements, hesitations, window switching, etc.
- Ask questions: Probe beyond the surface. Don't settle for "I just do it this way." Ask why, when, and what happens if they did it differently.

What to say:

"Can you narrate what you're doing and thinking as you go along?"

This "think aloud" method reveals not just the how, but the why - surfacing mental models, goals, and workarounds that might otherwise stay hidden.

2. Self-Check: Could You Do It Yourself?

Continuously ask yourself:

"If I had to do this task, could I do it based on what I'm seeing?"

If the answer is no, dig deeper. This is a sign that critical tacit knowledge is being used - important insights that need to be documented for designing AI systems.

3. Document Tools, Platforms and Data Flows

Make a complete list of all software applications, platforms, and tools the participant uses.

- Take note if the user switches between tools to complete a task - this often signals **an integration opportunity**.
- Later, these tools will help IT assess how an AI system could plug into or span across the current ecosystem.
- For each step in the workflow, document:
    - Where data comes from
    - Where it goes
    - What validation or quality checks are applied
    - Authentication steps or security protocols involved (e.g., logins, VPNs)

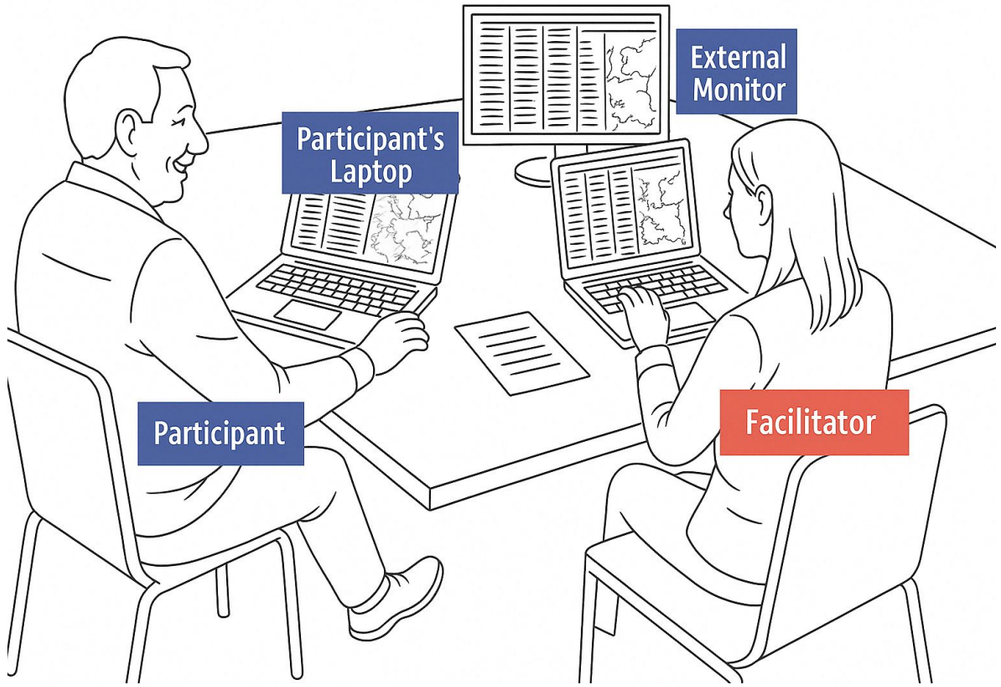

Above: Example in-person user research session setup. Use of an external monitor to mirror the participant's screen may make screen-sharing easier to observe. Adapted from nngroup.com.

STEP 3: INSTRUCT AI

Model the work activity for the AI.

This is a worked example of taking meeting transcripts (the task specific data) and turning them into work documents that might otherwise take hours to draft, using AI prompting strategies. What the AI produces might not be perfect, but it will save 70%+ of the time many people might spend on these documents.

Template
Separate the fixed part of the prompt - the core instructions - from the variable components, or template variables, which change with each new dataset. This separation makes it easy to update the inputs without altering the proven prompt logic, ensuring repeatability and reliability.

Explain everything the AI needs
Imagine sharing your AI instructions with a newly hired colleague and asking that person to execute the task based solely on those instructions. They have some understanding from general industry, but not the nuances of your business. If they struggle to understand or produce what you're looking for, your AI system will likely face similar challenges.

Provide examples
By studying these examples, the AI can infer many subtleties that would be difficult to articulate explicitly.

```
{{meeting transcript}}

---

You are a professional procurement analyst writing an RFP document for vendors. The above is a meeting transcript with the internal client who is seeking a new vendor.

(1) Write up the above meeting transcript into a 2000 word RFP narrative that covers the following elements of our standard RFP template:

    Background and Objectives: Provide context for the RFP and describe the overarching goals the internal client hopes to achieve by engaging a new vendor.

    Scope of Work: Outline the specific services, deliverables, or products that the vendor is expected to provide.

    Technical and Functional Requirements: Detail the required capabilities or performance criteria that the solution must meet.

    Project Timeline: Include key milestones, deadlines, and desired go-live date.

    Evaluation Criteria: Describe how vendor proposals will be assessed (e.g., cost, experience, methodology, compliance).

    Submission Requirements: List all required components of the vendor's proposal (e.g., pricing breakdown, references, project plan).

    Contractual and Legal Considerations: Include any standard terms, SLAs, or data privacy requirements.

    Budget Constraints (if any): State the available budget range or expectations.

(2) Determine any questions that vendors may have when reading the RFP, and provide these.
```

This is an abbreviated version of a real RFP draft created with the shown prompt.

```
RFP for High-Throughput Screening (HTS) Platform Infrastructure and Automation Equipment

1. Background and Objectives

Our organization seeks to establish a state-of-the-art high-throughput screening (HTS) facility to accelerate drug discovery through the rapid screening of small molecule compound libraries against a broad array of biological targets. …

2. Scope of Work

Vendors are requested to provide proposals for one or more of the following components, as applicable:

    • Small molecule compound libraries
    • Deep freezer units for long-term storage
    …

3. Technical and Functional Requirements

    3.1 Compound Library
        • Size: 150,000–200,000 compounds
        • Type: Small molecules (Mw < 300 Da), protein-class agnostic
    …

4. Project Timeline

RFP Issue Date: [Insert Date]

Vendor Questions Due: [Insert Date]

    …

5. Evaluation Criteria

Proposals will be evaluated based on the following criteria:

Compliance with technical and functional requirements

    …

Anticipated Vendor Questions

Vendors may have the following questions:

    • Are there any preferred or incumbent vendors?
    • Will site visits be allowed before proposal submission?
    • Is there a preferred vendor for LIMS or BMS integration?
    • Can the client provide floor plans or utility access points?
```

STEP 4: INTRODUCE

Measure how well the AI system works.

We now have an operational AI system for the target use case. First pilot the system with a small set of tasks or group of users, before scaling to a wider group. This AI system will operate with oversight from the expert for the use case

In the Identify stage, we identified the use case based on the “Wow factor” - where the ROI is self-evident in a 5x to 10x improvement in speed, quality, cost or other characteristics. We forwent precise ROI calculations for the numerous use cases considered to focus on what is evidently most valuable.

Nevertheless, the “wow factor” will already have given you an intuition for the one or two key measures of success. You will ultimately want to translate this measure to a monetary metric, to inform wider rollout of the AI system.

For example, for an automation related use case, the monetary metric may be:

```
Number of tasks processed by AI
x Time typically taken for a human to conduct the task
x Hourly cost of human who conducts the task
= Cost saved per task
```

For an upgrade related use case, an example monetary metric is:

```
% Improvement in task quality or task speed
x Number of tasks processed by AI
x Monetary value of each task
= Value created per task
```

In the case of an upgrade use case, for example, the quality differential could be the number of regulatory updates that would otherwise be missed, for example. The monetary value could be the expected mean value, historically, of fines or reputational damage from missing an update.

In some cases, it is not easy to define a clear monetary metric. In these cases, it may still be possible to quantify time savings or a similar metric. Staff could be asked to record the time they have saved by querying the AI system, for example, rather than seeking the information manually, along with the number of instances of this activity in a typical week. In this way, an approximate monetary metric can still be measured.

The system you have built may justify an expanded implementation with more custom software development, which can now be justified with these metrics.

AI System Pilot – Preparatory Checklist

Pilot Planning

- ☐ Define size and composition of the pilot group
- ☐ Set duration of the pilot and key milestones
- ☐ Agree on to record before launch
- ☐ Establish baseline metrics and success criteria

Technical Setup

- ☐ Confirm required tools/software are installed
- ☐ Set up access permissions for pilot group
- ☐ Validate data connections (e.g., real-time/live data feeds)

Staff Enablement & Support

- ☐ Finalize training materials and plan
- ☐ Assign trainer(s) or onboard facilitator(s)
- ☐ Define point of contact for pilot-related questions
- ☐ Establish issue escalation process during the pilot

Communication Plan

- ☐ Prepare champion(s) to communicate goals and expectations

Risk & Mitigation Planning

- ☐ List known limitations or bugs in the current system
- ☐ Define fallback processes or manual overrides

Feedback & Iteration Process

- ☐ Choose tools/channels for collecting staff feedback
- ☐ Set frequency for check-ins during the pilot
- ☐ Establish update/iteration plan during pilot

Post-Pilot Evaluation

- ☐ Plan how results will be synthesized and presented
- ☐ Outline decision process for potential full rollout

ABOUT THE AUTHOR

Dino Gane-Palmer is the founder and CEO of PreScouter, an Inc. 5000–recognized innovation consultancy that helps Fortune 500 companies and global organizations capitalize on new markets and emerging technologies. An entrepreneur and technology strategist, Dino launched PreScouter while earning his MBA at Northwestern University's Kellogg School of Management. He also holds a Master of Engineering in Computing from Imperial College London.

With a passion for turning new technologies into practical business tools, Dino has spent almost two decades helping industry leaders navigate innovation. He is also the creator of Auxee, an enterprise AI platform that enables teams to codify, automate, and scale high-value workflows using generative AI - transforming how knowledge work gets done.

Dino is a frequent speaker, podcast host, and advisor on the future of work and the real-world application of AI in business. His insights blend hands-on leadership experience with a no-nonsense approach to technology adoption.

He lives in Chicago with his wife and son.

You can connect with Dino on LinkedIn, sign up for his newsletter at DoMoreWithLess.now, or tune into his podcast on YouTube, Spotify, or Apple Podcasts.