APPLYING ARTIFICIAL INTELLIGENCE

THE PRACTITIONER'S HANDBOOK

DAN ROSE JOHANSEN

Biased Publications

APPLYING ARTIFICIAL INTELLIGENCE

APPLYING ARTIFICIAL INTELLIGENCE

THE PRACTITIONER'S HANDBOOK

DAN ROSE JOHANSEN

BIASED PUBLICATIONS

Published by Biased Publications and distributed by TBK Publishing®

ISBN 978-87-93116-51-1 (softcover version)

Created with Vellum

CONTENTS

Foreword by Lars Tvede xiii
Introduction xix

1. MAKING BOOK-KEEPING MORE PRODUCTIVE 1
Chaotic first year 2
First and confusing feedback 3
Beta-launch 5
Humans in the loops – what we did right 6
Growth 7
Research and development don't always pay off 8
Going global 9
Happy ending 10

2. DEFINING AND DEMYSTIFYING AI 12
Machine learning 13
Artificial intelligence – a misleading name 13
Defining AI is essential 14
Cat or dog? 15
Data, information and predictions 16
Relations and casualties 16
Three (or so) ways for AI to learn from data 17
Generative versus predictive AI 21
Generative AI can also predict 22

3. AI BUILDING BLOCKS 24
No homemade pasta 25
Building blocks 26
Labels and features are fundamental to AI building blocks. 28
Tabular building blocks 29
Classification 29
Regression 30
Forecasting 32
Vision building blocks 35
Classification 35
Object detection 36
Image segmentation 38
Image similarity 39

Language Building blocks 40
Classification 41
Named entity recognition 42
Intent analysis 44
Sentiment analysis 44
Summarising text 48
Sound 48
Audio classification 48
Multimodal building blocks 49

4. PREPARING YOUR ORGANISATION FOR REAPING THE BENEFITS OF AI 51
Domain experts 54
Knowledge drives demand 54
It's not about hiring technicians first 54
Why the encapsulated and technical approach fails 55
AI is bottom-up 57
How do you teach AI? 58
Bridge the gap of ignorance. 58
To adopt or not to adopt isn't a choice 59
A technology for solving business problems 59
Where is your organisation today? 60
Ignorant 62
Three adoption strategies 63
Summary 66

5. SOLVING PROBLEMS USING AI 68
Methods aren't new 69
The Todai Method 70
Method steps 70
Using the method 70
1. Inspiration phase 71
2. Discovery 72
3. Data handling 73
4. Development 73
5. Implementation 73
6. Monitoring 74
Milestone funding 75
Mind the mindset 75
Building AI is like developing a product 76
How to be product-focused 78
Measure user adoption and make it your responsibility 78
Think like a product manager 79

Build a product vision 79
Find product market fit 80
User experience design 80
Summary 80

6. INSPIRATION PHASE 82
Objectives 82
Activities 83
Training people to identify opportunities 84
Make AI accessible 85
Idea workshop 85
Add a pretotype 86
Design sprints 86
Pre-analysis 87
Decision models 88
Ask for data 88
Build models 89
Interview domain experts 90
Look what others have done 90
Inspiration phase checklist 91
Building the business case 91
Moving forward 93
Set goals and calculate backwards 93
Assume everything will go wrong 93
Secure stakeholder support 94
Summary 95

7. DISCOVERY – UNDERSTANDING THE PROBLEM 96
Discovery phase 97
Discovery phase activities 98
Solving the right problem 98
What are you trying to achieve? 98
Selecting a business outcome 99
The Paperflow problem 100
The SundAI problem 101
Connecting with strategy 103
Understand the problem domain 105
Problem components 105
Understanding the Paperflow problem 107
Use well-known frameworks 109
Has the problem been solved before? 109
Do you agree on the problem? 110
Use a workshop tool 111

Acquiring domain knowledge through subject matter experts 111
Data is only half the story 112
Domain expert interviews 113
Exploring feature engineering 113
Implicit knowledge 114
Beware of bias 115
Interview more experts 116
Expertise-induced blindness 116
Interview guide 117
Testing before building 120
Summary 120

8. DISCOVERY – DECISION MAKING AND ACCURACY 121
The effort in AI is skewed 121
Decision model 123
Generative models and decision models 124
Confidence thresholds 127
Decision models incite new models 128
You probably want more info 129
Other types of models 129
Decision levels 130
Level 1: Warn/inform 130
Level 2: Suggest 130
Level 3: Conditional automation 131
Level 4: Automatic 131
Level 5: Autonomous 132
Confidence and accuracy 132
Accuracy and confidence are not the same 133
Accuracy value 134
Agree on how you measure accuracy 134
Not all information is equal in value 135
Perceived accuracy 136
Who are you trying to beat? 137
The value is in the decision 138
Decisions first – always 139
Build the business case on decisions 141
Decision flows 143
Element of luck 144
Map the flow 145
Decision-making is complicated 146
Decisions are feelings 146
Kahneman has a point 147

Decisions are affected by politics, compliance and strategy 148
Summary 149

9. DATA HANDLING 150
Activities in the data handling phase 151
Goal of the data handling phase 151
Data terminology 151
The goal of data is to represent the world in which the AI should work 153
Perfection is impossible 153
How to handle it 154
Data quality 158
User-generated data 160
How much data and at what cost? 160
Building blocks matter the most for collection costs 161
Generative AI data costs 162
Adoption strategies and data costs 162
Foundation models 163
The desired accuracy 164
Optimal cost 165
Faster and better data labelling 165
Model-assisted labelling 165
Active learning 166
Understanding and analysing data 167
Look at the data 167
Black swans 169
Data bias 170
Mitigating bias 171
Data as a competitive advantage 172
Personal data 173
Challenges 174
Handling PII 174
Synthetic data 175
Add data to labels with no or little data 176
Remove personal data 176
Whether data is synthetic depends on the use 177
Synthetic data will become a big deal 177
Types of synthetic data 177
Synthetic texts 177
Synthetic images 178
Synthetic tabular data 179
Models of the world 180
Summary 180

10. DEVELOPING AI 181
The development phase activities 182
Team roles in AI projects and their challenges 182
Still a need for data scientists? 183
Anatomy of AI solutions 184
Model(s) and prediction pipeline 186
Generative models 189
Mindset for building AI 191
Experiment efficiently 191
End-to-end first 192
Be ready to start over within twenty-four hours 193
Speed is of the essence 194
Simplest approach possible 194
Building and understanding models 195
Understanding models 196
Explainable AI 201
Testing results 204
Monitoring phase: Building AI operations 207
Monitoring AI activities 207
Operations, feedback, and retraining 208

11. IMPLEMENTATION PHASE – AI AND HUMANS 213
What makes implementation so important? 213
Starting the implementation phase 215
Before the implementation phase 216
The implementation phase reveals discovery failures 216
Achieving adoption 218
You get what you measure 218
Adoption is a marketing task 220
Humans find patterns 220
Gaining trust or treating fears and insecurities 222
The problem 222
Solution 224
Sense of control 228
The Kasparov perspective on AI adoption 229
Social contract 229
Summary 232

Afterword 233
Acknowledgments 237
About the Author 239
Index 241
Notes 247

FOREWORD BY LARS TVEDE

AI is huge - we all know that. However, until you have worked quite a lot with it, you may not know what it's good and bad at. I have personally made thousands of AI prompts and queries, and it works very differently from what I had expected beforehand.

As I see it – now - a Large Language Model (LLM) is like a genius with a huge talent for fantasy and creativity. However, it is mostly not as good at doing arithmetic as the simple calculator that I used in the 1980s. In fact, half the time, it gets basic calculations wrong.

AI tends to get other facts right, but not entirely. For instance, if I ask it about a historical event, it may give a brilliant summary and explanation and describe the year and place it happened, plus the people involved. Fantastic!

Except that two of the people it mentions were not involved, and it happened 16 years later than it stated.

I figured out how to deal with this: Ask several AI models the same questions and compare the results. If they are different, ask each "are you sure that …?"

Sometimes, when working with LMMs, I feel like a police interrogator, if not a psychologist trying to get the truth out of a somewhat unreliable person. Eventually, it tends to work very well.

However, one area where LLMs tend to be brilliant the first time, is when you ask it to come up with product names, brand names, logos, and slogans. This is because generative AI is creative, which is the opposite of what some people who have never used it often believe. And that brings me to the point that you can use AI to automate tasks and thus save a ton of money and speed up stuff. For instance, in figuring out how proteins fold in 3D, or tagging, illustrating, and summarising huge amounts of text in a jiffy. But you can also use for hyper-creative tasks, such as composing music, creating poems and being visually artistic. These are very different tasks.

So, there is stuff to do, but just comprehending the options is a challenge. For instance, as I began writing this foreword, there were 484,932 AI models uploaded to the Hugging Face website. Each of these was trained for a specific task, and the number is doubling approximately every 3.4 months. This corresponds to 11.5 times more per year.

Get out the calculator!

This sounds like there will be approx. 5.3 million different AI models at the end of 2024.

Incidentally, if we roll the rule forward to early 2030, we will have well over 1 trillion different AI models by then, hooray. But, wait-a-minute. That sounds crazy, so, perhaps not. But on the other hand, perhaps yes, I'll get to that. But in general, I believe that there are three applications of AI that will be particularly interesting:

1. AI at scale is the use of AI to solve large problems that require the processing of massive amounts of data. This is an area of AI that is growing rapidly, and it has the potential to revolutionize a wide range of industries, from healthcare to finance. AI at scale could be used to analyse large amounts of medical data to develop new treatments for diseases. It could also be used to improve the efficiency of supply chains and manufacturing processes, for instance.
2. AI orchestration is the combined use of multiple AI models. This is important because AI models can often be better at solving problems if they work together. AI orchestration could be used to coordinate the work of multiple AI models to detect and prevent fraud in financial transactions. It could also be used to improve the accuracy of weather forecasts.
3. Personal AI is the use of AI to deliver personalized services to individuals. This has the potential to have a major impact on our lives. Personal AI can be used to give us tailored recommendations for products, services, and experiences. It can also be used to help us manage our health and finances. Personal AI could be used to give us personalized recommendations for movies, music, and books. It could also be used to help us learn new languages and skills, to be our education coach (à la Khan's Academy), our fitness instructor, food mentor, etc.

I believe that all three of these AI areas have the potential to be transformative. AI at scale has the potential to solve some of the world's most pressing problems. AI orchestration has the potential to make AI more powerful and efficient. And personal AI has the potential to make our lives easier and more enjoyable.

Now, going back to the previously mentioned ridiculous sounding 1 trillion different AI models by 2030. This number is not my forecast, but I don’t think it is impossible either. If most people by then have a personal AI trained on them as individuals, and if each of these personal AI models is in fact not really one model but based on AI orchestration with a heterogeneous symphony of models working in

concert on helping each individual, then yes, we might get there, or close to it. I mean – we might get close to the trillion AI models in 2030, or something thereabouts.

A factor in the growth of Ai is that it is becoming exponentially cheaper. Partly due to ever-better dedicated chip types, we have been able to drive down the costs of training AI models with a psychedelic half-life that is sometimes 5-6 weeks. In fact, there is an expectation that the training of an AI model that cost over $4 million in 2020 will be reduced to just between $30 and $300 in 2030. By which time, it should be said, AI may also have passed a Touring Test and thus be as least as smart as a typical human for all tasks. Two years after that, it will probably be smarter on all tasks than the smartest human that ever lived.

The field of AI has long been shrouded in technical jargon and algorithmic complexity, often alienating those it aims to benefit. Johansen's work diverges from this path, emphasising that the application of AI is fundamentally different from traditional IT. Where traditional IT focuses on systems, infrastructure and code, AI is an explorative journey into the realms of human behaviour, decision-making, and business processes. This differentiation is critical; it shifts our perspective from viewing AI as a mere tool to understanding it as a transformative agent in business and society.

The Todai Method, as meticulously detailed in this handbook, is built upon this very premise. In a world where AI is rapidly evolving, the Todai Method acts as a compass, guiding practitioners towards a user-centric and discovery-driven approach. This is especially pivotal as we stand on the brink of a generative AI revolution, where the ability of AI to create, predict, and decide will significantly influence every aspect of our lives.

Johansen makes a compelling case for an approach which ensures that AI solutions are not just technically sound but are also aligned with the real-world needs of businesses and their customers. In doing so, it bridges the gap between the theoretical potential of AI and its practical, value-driven application in the business world. And as we venture deeper into the era of generative AI, where machines can create

content, the importance of a human-centric approach becomes more pronounced.

In conclusion, "Applying AI - The Practitioner's Handbook" is a timely and invaluable resource for anyone looking to harness the power of AI in their business.

I hope you will enjoy reading it as much as I did!

Lars Tvede

P.S. Hey GPT: "When Lars finished writing this foreword, Hugging Face had 485,008 models. How long did it take Lars to write the foreword?"

INTRODUCTION

I've been working with artificial intelligence for almost ten years and am impressed by what the technology can do.

Recent breakthroughs like OpenAI's ChatGPT and Stable Diffusion from Stability have garnered global attention and showcased remarkable results. However, confusion around the words "artificial" and "intelligence" mystify and obscure their practical use. From a distance, AI appears to combine information technology tools to develop systems capable of producing predictions, estimates, forecasts, projections, and other types of assessments that come with a probability. Let's stay with the statement that AI is a label for IT tools that make predictions.

It can be extremely difficult to determine which problems will benefit from AI and how these projects can be steered to successful completion (or abandoned before they run amok). Between 75% and 80% of AI projects fail because inadequate tools are applied to problems, or the projects are mismanaged.

This book is for anyone involved with identifying processes or areas where AI can improve performance. It will help you to build business cases that justify your efforts (or lead to an early and inexpensive rejec-

tion) and help you to organise projects for its implementation. Anyone involved with improving organisational performance can benefit. Rookie or pro. Entrepreneur or civil servant. Architect or engineer. Executive or operational agent. Anyone.

The secret behind any successful IT project is sharing a standard method, understanding its vocabulary and the steps involved, and early user engagement. The method described is *specifically* designed for AI projects, and it works best when the entire team follows its recommendations.

Applying the following principles can bring your AI project's success rate far beyond 90%. This requires abandoning hopeless projects early, completing projects with a promising business case fast, and getting users to embrace the outcome produced by the systems.

As well as learning about AI, its fundamental components, and its practical aspects, you'll discover how to:

- Spot areas in your organisation where AI can provide value.
- Develop a systematic approach to applying and developing new AI solutions. This includes creating a business case, managing projects effectively, and ensuring user adoption.
- Effectively collect data for your AI solutions.
- Build AI that users love and are eager to adopt.

Adopting new technology too early can be expensive and require significant investment to achieve marginal value functionality. On the other hand, if you adopt it too late, you risk missing out on potential savings, better customer service, and other benefits that could improve your operation and competitive position. However, the pace of AI development is accelerating rapidly. Off-the-shelf solutions can solve mundane problems at very reasonable cost points. In fact, the cost of using AI is dropping at an astonishing rate.

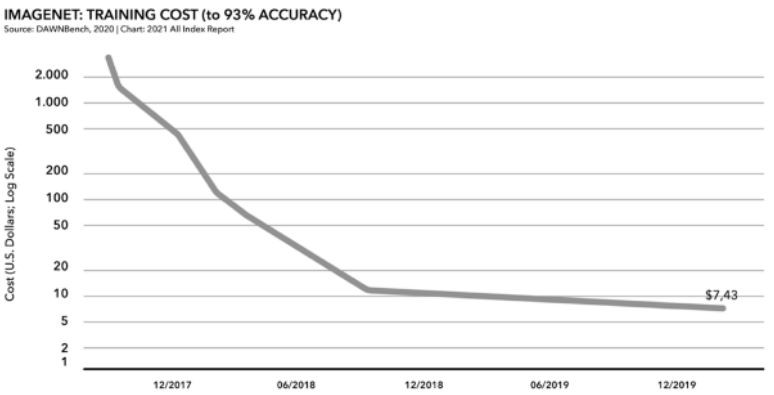

Figure 1: The cost of training an AI-model to 93% accuracy

The graph above shows that in 2017, it cost over $2,000 in computer power to train a model to 93% accuracy. In 2020, the same model cost as little as $7.43 – a 300-fold cost reduction in three years. Training time also dramatically improved. Thirteen days in 2017 dropped to as few as 2 minutes and 38 seconds in 2020. In 2023, the time and costs are so low that measuring them no longer makes sense.

AI is becoming more affordable, and its performance is significantly improving. For example, ImageNet's image recognition accuracy improved from 83% in 2013 to over 99% in 2021.

The graph might indicate that progress has slowed, but a one percentage point improvement from 98 to 99 corresponds to a 50% reduction in the error rate. Such a reduction can be a game-changer for a business case, turning a no-go into a no-brainer.

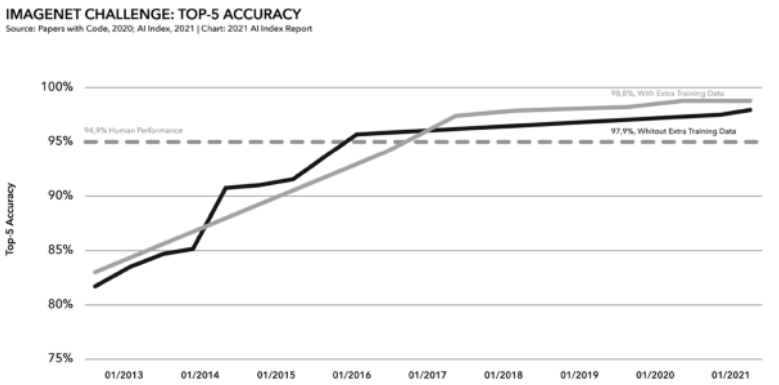

Figure 2: ImageNet challenge: top-5 accuracy

People make approximately 6% errors when assessing an image. This means that in 2016 AI assessments surpassed those of humans. Since 2021, AI has become exceptionally proficient, and it can now outperform people – even experts – in more and more areas.

Investment in AI is rapidly growing and showing an almost exponential increase. In 2015, global investment was around $13 billion. By 2021, the figure had surged to almost $100 billion.

Although AI is part of *information technology* (IT), your implementation will fail if you attempt to prepare and run an AI project like a traditional IT one. AI requires everyone within and around the project to understand what AI is, how it works, what outcomes it can produce, and how those outcomes can be translated into operational value. My method breaks down an AI project into a step-by-step approach. This includes the:

1. inspiration phase, where you can identify what AI can do for your business
2. discovery phase, where you set your objectives and design your solution
3. data phase and the challenges associated with inaccurate or low-quality data, as well as data-related ethical and legal issues
4. development phase, where you build your solution

5. implementation phase, where you deploy your AI-driven solution into a human context
6. monitoring phase, where you make sure that the solution works and produces value

Remember that the reason to adopt AI is that it provides specific business advantages. For example, it could save you time or improve product quality. As AI evolves and matures, it will bring countless new opportunities. And there's no need to wait – there are already countless practical ways to use it.

I have written this book to help you do just that.

CHAPTER 1
MAKING BOOK-KEEPING MORE PRODUCTIVE

In 2013, I joined a startup called Billy, which developed user-friendly accounting software for small businesses and independent contractors. While the software simplified accounting, users had to input invoice data, a tedious and time-consuming process manually. This sparked the idea for a new startup, Paperflow, to provide an inexpensive and easy-to-use service that could automatically read and interpret payable invoices and enter the data into the accounting software.

Existing solutions were geared towards big companies and needed a customised set up for each creditor. This made them too expensive and complex for small businesses. People spent countless hours reading and entering invoice data, causing unnecessary costs and errors in financial reports. We aimed to make the process more affordable and productive by offering a simple solution that didn't require custom implementation. In other words, we wanted to offer small business owners a book-keeping solution that worked at the click of a button.

We believed there was a vast global market for such a service and were confident we could develop an easy-to-use solution. Small business owners could say goodbye to tedious manual accounting operations and devote more time to growing their businesses.

How could this not become an instant success?

CHAOTIC FIRST YEAR

We started with a small team of three: the Billy founder, a chief technology officer (CTO), and me as marketing manager. I prepared for our launch while the CTO was coding. The CTO jumped ship five months later, and I took over development and product management. That meant I had to find developers and figure out how to proceed.

First, we tried the traditional if-then-what-else approach. That failed, so we introduced AI tools. At that time, we were testing an open-source coding library called Tesseract, an off-the-shelf software that used optical character recognition (OCR). It was the go-to library for reading alphanumeric text on images and had been around for over twenty years. In theory, the software should have been able to read documents out of the box. We only needed to input images and to get the results: words, numbers, and their positions in the document.

Or so we thought. Tesseract was an AI or machine learning-based system, and we didn't fully comprehend what that meant. I'm not sure that the words "machine learning" or "artificial intelligence" were ever spoken in the office.

As we studied the results, we realised we were working with something different from traditional IT. We couldn't say if the OCR was correct when we scanned a new invoice. Parsing values such as amounts and dates from invoices was challenging if we couldn't be sure the letters and numbers the OCR read were correct.

It dawned on us that the OCR's accuracy was contingent on the fonts and layouts used, and we'd been unaware of the volume of invoice variations prospective clients would bring. We had limited knowledge

of the performance we could expect from our solution before we saw it work in production.

OCR technology has existed since the mid-1960s and was more affordable, with lower prices and more options available by 2015. While working on Paperflow, Google launched a high-quality cloud-based API for OCR that could be completed at a much lower cost. We decided to use that to deliver the values, allowing us to focus on identifying relevant content in the invoices. However, interpreting text and numbers was difficult, especially when identifying the data elements necessary to determine the creditor's identity, invoice and due date, quantity, product description, unit price, discount, order number, total amount, VAT, and other relevant information. This task is known as *entity recognition* (NER) in AI terminology. While the term existed in academic circles, it wasn't widely recognised or applied in the industry or available in off-the-shelf AI tools.

We were dependent on immature technology. While the algorithms now referred to as AI had been around since the 1960s in academia, they were rarely applied in an operational context. Most of the material on the subject was either promotional fluff from software vendors or highly technical papers from academia.

As we approached the end of the first year, we needed to speed up product development, and we hired our first developer with AI experience. He started working on the problem and testing different machine-learning approaches. We were finally making progress.

FIRST AND CONFUSING FEEDBACK

Next to the data recognition engine, we also needed a user-facing application enabling book- keepers to approve the results. This would submit invoices by email to the AI engine, which presented the results on the interface within 30-60 seconds. The user could approve or correct the data, transferring and committing it to the accounting software.

After six months of hard work, we finally had a web application to present to potential customers. This was a significant milestone, as the product was now becoming tangible. Such solutions weren't readily available or a standard feature in any accounting software, and our potential users were unfamiliar with their potential. We were excited to showcase something new that could substantially improve productivity.

However, although people were curious, they had numerous objections.

We were trying to develop a seamless "one-click" invoice workflow that was inexpensive, easy to use, and timesaving. But bookkeepers weren't looking for an easier way of doing things. They didn't want the flow to be automated and hands-free. Some said that manually entering invoices "gave them a feel for the state of the businesses," and others said that unless the results were 100% correct, it was useless.

We were surprised that the possibility of saving time and avoiding typing mistakes wasn't met with more enthusiasm. Also, the book-keepers wanted our solution embedded in their accounting software before considering using it. Some recognised our solution's improved efficiency in processing invoices but didn't want to bother with an additional application. We had to find a way forward.

Our assumption that extracting invoice data was challenging and that bookkeepers would independently determine the optimal way to use it was wrong. They wanted access to the functionality within their existing software, making things much more complicated.

It didn't help that the book-keepers had differing opinions about interpreting identical invoices. While one identified a specific field as a "project number," another considered it an "invoice number," and a third labelled it an "order number."

We also had problems making the AI engine produce high-quality output fast enough, lacking the required expertise. Our hiring approach had been to recruit anyone with machine-learning experience. However, AI is a vast field with numerous specialisations, and

our team's AI skills weren't an ideal fit for the product we were developing.

As a temporary solution, we opted to avoid AI for a while. Instead, we developed a simple rule-based approach. Surprisingly, we discovered that such rules could yield an accuracy rate of up to 60%. Although not competitive, it was a starting point.

Our product could now extract data from invoices, but the accompanying web application still didn't align with the users' requirements. A stand-alone solution meant we had to construct and maintain all the integrations for differing accounting software systems — a demanding job for our small team.

We decided to streamline our product scope. That entailed eliminating the web application, making us solely an API-based platform. That way, we could concentrate on a specific aspect of the book- keeping process: identifying data on payable invoices. And it distinguished us from our competitors, who all relied on web or desktop applications.

BETA-LAUNCH

We believed that we now had a *minimal viable product*. The accuracy of invoice parsing was insufficient to revolutionise accounting, but it was enough to launch a functional product and generate revenue. We could have invested more time crafting a superior solution, but we needed to prove we could generate revenue to keep the funding coming in.

We worked with Paperflow's co-founder Billy (which had more than 10,000 users at that time) to kick off the beta launch. But one thing is to convince a SaaS vendor to include our service in their offering, and another is to win over users and make them happy.

For Billy to use Paperflow's invoice reader, they had to integrate with our API. Invoices were sent to our service automatically, the information parsed and then returned. They had no problem implementing the API in a technical sense, but dealing with the uncertainty of AI was new to them. When Billy sent us an invoice and received the results, they didn't know if they could trust them. The results could

be perfect with all the invoice data read correctly by our software, or it could be wrong. Most cases would be correct and few incorrect. That would never happen in traditional IT, but explaining the difference was tricky. The room for interpretation of results also made testing harder. Implementing AI with this first customer was supposed to be an easy and sure thing but dragged on and took a toll on the team.

HUMANS IN THE LOOPS – WHAT WE DID RIGHT

Following the release of our initial platform version, I attended a machine learning Meetup in Copenhagen hosted by Tradeshift (a Danish startup that primarily dealt with electronic invoices) where they presented their solution for reading invoices from images or PDFs. Tradeshift was working on the same problem, and their presentation gave me valuable insight.

Tradeshift's approach involved collecting data by analysing previously entered invoices and using a program to identify the position of values on the invoices. This required less effort and achieved more data than Paperflow's manual marking of each word. The downside was lower data quality.

The reason is as follows: The book-keeper entered "100" as the total amount. The program then searched for "100" and marked it as the total amount. This position was used as training data to teach a machine-learning model to read other invoices. The benefit of this approach was that you can process many invoices at minimal expense and quickly obtain a significant amount of training data. However, the program's identification of "100" as the total amount might be inaccurate, as "100" could also appear elsewhere on the invoice. Additionally, users make more errors when entering data without AI forming part of the training data.

Tradeshift had achieved an 82% accuracy rate, on par with the market benchmark at that time. However, the poor data quality used to train the models curtailed improvement. At Paperflow, we wanted to surpass this accuracy rate. One solution was to obtain more high-

quality training data, but doing that as a startup with limited funds was problematic.

We decided to offer a premium product version with a human verifying the AI output before passing the data to the customer's system. The verification step would allow us to gather high-quality data and get the AI accuracy past the 82% standard while having the customer cover the costs.

Our mistake was in announcing this version would yield 100% accuracy. Even a four-eye principle doesn't give absolute accuracy. Eliminating errors required other processes and controls and would be extremely expensive. Even if we had no errors, the book-keepers would still disagree about correct values. Our 100% claim led to negative feedback from users who found errors. We changed the claim to "equal to the average bookkeeper."

This also failed.

Like anyone, book-keepers tend to believe they perform better than average. Although we stumbled with our marketing messages, the human-in-the-loop product development approach succeeded. We found customers who recognised the benefits and were willing to pay a premium. The feature allowed us to gather significant amounts of high-quality data, which our competitors lacked. Data quality became our strategic advantage.

At Paperflow's inception, we'd assumed that acquiring "some training data" would be a onetime task that could be completed within a few weeks or months. In fact, data acquisition became a continuous operation that at one point involved almost fifty full-time employees—a vast difference from our initial expectations.

GROWTH

New and improved data meant our product's accuracy increased, leading to an increase in our client base. But we still suffered from having launched prematurely. The system was complex to maintain and keep online, consuming valuable time and resources we could

have used for product improvement. As a result, we were slow in introducing new features and accommodating customised requests. Despite this, we managed to sustain our growth.

After limiting the product to an API-based engine, our go-to-market approach shifted. We now had to convince accounting system providers (such as local providers like Billy and Uniconta and international providers like Microsoft Business Central) to incorporate our service into their offerings. That caused us to lose direct contact with end users. We also lost control of how the predicted values were presented in the different accounting software solutions.

For example, our solution predicted the document type, such as an invoice or credit note. Regardless of the document type, we'd still be able to predict most of the values found. However, the prediction appeared empty if the AI couldn't find a value. Sometimes, the payment date wasn't on the invoice, or the AI couldn't find it – even when it did find the payment date, the accounting software might not present it to the user. This was frustrating as it was hard for the user to identify whether Paperflow, the accounting software, or the invoice was the problem.

Relinquishing control over how the results were presented to users was more significant than we'd anticipated. To prevent users from having a negative experience, we had to refine the onboarding process for accounting software providers. Convincing them to disclose how they applied Paperflow's product wasn't easy. This challenge is now recognised for AI solutions as they tend to be in the backend, but it took us by surprise at the time.

RESEARCH AND DEVELOPMENT DON'T ALWAYS PAY OFF

We aimed to create and provide the most accurate invoice scanner possible at an affordable price. By accumulating more and better data, our service quality continued to improve. We then took a more research and development-focused approach to enhance the AI algorithms and collaborated with the University of Copenhagen. Our goal

was to outperform our competitors by being at the forefront of technology.

We focused on probabilistic programming, a branch of AI modelling that wasn't widely used in an operational environment at the time – it had only recently started gaining momentum. Even in academia, it wasn't a significant field, although the basis for probabilistic models has been around for a few hundred years. We invested significant effort in it as a potential avenue for improving accuracy, but despite its initial promise, it didn't yield significant results. The technology was still too immature, and the resulting service was unstable and prone to crashes, requiring even more time and resources for maintenance.

In retrospect, we should have directed this effort towards improving the core product and enhancing user satisfaction.

GOING GLOBAL

We envisioned Paperflow to be easily scalable internationally. Once we had a functional solution and had gained traction in our local market, we decided to invest in expanding abroad. Again, our efforts didn't succeed like we'd hoped for. With Danish companies receiving more than 20% of their invoices from abroad, we were confident about having ample training data. However, we'd wrongly assumed that our AI could read invoices from other countries effectively. Moreover, believing that European countries less digitally advanced than Denmark would find AI-powered invoice reading appealing was also wrong. While extracting data from a document and inputting it into an accounting software system appears identical worldwide, the willingness to automate it varies significantly.

Our go-to-market approach also posed a challenge to international expansion. We had to identify accounting software providers, convince them to cooperate, get them onboard and integrate Paperflow into their offerings. However, cloud-based services like Paperflow weren't widely accepted in many countries. Furthermore, we lacked a position and network abroad that would enable us to reach potential partners.

For those we attempted to get on board, we met an even broader range of accounting workflows, which made it complicated.

The prevalence of paper-based business documents in many countries posed another challenge. Even in countries such as Sweden and the UK, geographically close to Denmark, around 50% of invoices were still paper-based and delivered by mail. This starkly contrasted with Denmark, where less than 10% of invoices were paper-based, and 90% were emailed as PDFs. As a result, the task for our AI was more complex, as it had to read a more significant number of physically printed invoices that varied in appearance.

HAPPY ENDING

Paperflow ended up in a good place. It didn't become the unicorn all entrepreneurs dream of, but it did well despite the many challenges. In August 2022, a large SAP provider acquired it, and its revenue continues to grow.

The main things we learned from the Paperflow experience were:

- Despite proof of value, users are generally reluctant to embrace new technology, which hinders technical progress. Implementing innovations requires identifying the technology enthusiasts and visionaries prepared to work with seemingly premature solutions and iron out inconveniences.
- Scaling innovative solutions requires substantial user involvement, education and coaching.
- It is hard to predict how users will respond to output that is not 100% accurate. The book- keepers with whom we communicated expected absolute accuracy. They overestimated or ignored their own accuracy rate and demanded much more from an automated system.
- Users aren't necessarily motivated by improvements in productivity. If such improvements seem to reduce their level of control, they resist the change.

- Sticking with tried and tested AI technology will likely give faster and less expensive results. Trying to get ahead while at the forefront of technology is risky. There are ample opportunities for improvements elsewhere in most projects.
- Data quality plays a significant role in the outcome of any AI-based system. The cost of generating enough quality data can become the most significant issue.
- Because it's hard to foresee what it will take to produce a certain level of AI accuracy, users must be involved early in the project to validate accuracy needs. If they eventually reject the output accuracy, the project will fail. Even when performance improves, a bad reputation can be hard to rectify.
- Certain problems can only be addressed using AI-based systems. Identifying these problems and understanding the process of preparing a business case and what it takes to run a project requires that all stakeholders have a foundational understanding of AI technology. This makes AI projects different from other IT projects where the stakeholders don't need a common foundational understanding of the technologies.
- Using an AI-specific project management method, including a shared vocabulary, improves any aspect of a project compared to a trial-and-error happy-go-lucky approach.

Building an AI-based company from scratch paid for my AI education, taught me how to apply it, and showed me all the pitfalls. I'd like to pass on that experience to you.

AFTERWORD

In the three years I've been writing this book (and my eight years in AI), AI has changed from a niche technology to a powerful tool every business wants and needs to figure out and handle. It's no longer simply a field for research but an applied technology where what used to involve hard work and loads of code can now be done in minutes. During the last year or so the expansion of generative AI has also changed the playing field dramatically. Still, the approach that works the best for applying the technology and building solutions is no different. The same goes for the tools and techniques for understanding and aligning a problem space with business needs.

Teaching people what AI can do and what it is and giving them a common language and best practice matters even more now that the technology is unavoidable.

I've had the luxury of cherry-picking from best practices that originated as far back as the 1960s, and I find that the user-focused mindset is the one that works. AI is nothing without its users. It doesn't matter what kind of AI or how advanced a solution you build, the same basic human dynamics must be accounted for and handled.

That's why learning the discipline of applying AI and using this method is a long-term strategy. Many people spend a lot of time choosing the right data platform, LLM or cloud provider to give them a strategic advantage. But when you truly understand the discipline, the technology parts will fall into place. Practice will – of course – be your most valuable teacher.

My method is more elaborate in the discovery phase than many other AI methods. That is especially valuable when applying Generative AI, for which there will be many more use cases in the future. Compared to Predictive AI, Generative AI takes away much of the code, retraining flows and prediction pipelines. At the same time Generative AI also comes with infinite output space and much greater complexity. This means that the business understanding, and the softer skills and techniques, often found with the humanist, are becoming ever more important and demand a shift of effort away from technology and towards AI users.

The demand and use of data in AI has also changed. We used to need a lot of structured data to be able to complete AI projects. Generative AI can use our unstructured data and companies can now unlock value that previously seemed unattainable.

Sceptics might say that AI still isn't useful due to its limitations and the fact that it can make "stupid" mistakes. But it's not either/or. There's a time and place for AI; the bottleneck isn't the technology itself but our ability to utilise it. That situation won't change for many years ahead as advances in technology greatly outperform the speed of which societal and business structures can change. This is also where we will see the largest clashes and conflicts. For example, work on the first EU regulation on AI that was agreed upon in December 2023 started in 2015 – during the same eight years I've seen the technology go from niche to powerful and invasive.

Much of the text, reasoning and motivation behind the legislation hasn't changed. It's likely that the gap between technology and the political system will continue to increase. For example, women have had equal rights in the labour market for decades in the Western

world, but the number of female CEOs in Denmark (a progressive country in most ways) is still much lower than male CEOs. In fact, more CEOs are named "Lars" than there are female CEOs.

People are conservative when it comes to changing business structures, but technological advances have the opposite trend.

Going forward, the next big thing will be the rise of AI agents – big LLMs that can access the internet, emails, and other documents will be let loose. Chances are it will be like the wild west for a few years, but it will also be extremely productive and valuable.

We're only beginning to learn how to use AI agents. At the time of writing, my company is working on several projects that include agents with access to both the internet and company documents. The biggest issue is that considering the vast opportunities this provides is almost paralysing.

Before long, we will see sophisticated AI models similar to Large Language Models (LLMs) that are multimodal. These models will be capable of processing and generating various types of data, including text, audio, images, and video. This development process is already underway. The output can also be in different data types depending on what the model decides is better. We'll also see a new kind of model for reasoning that in contrast to LLM doesn't use text as the basis.

But remember that AI has no motivations or feelings, so it won't take over the world as some fear!

You might wonder if you should bother building AI solutions now it is being built into all kinds of software. It will soon be in Word, Excel, your browser, whatever you use daily at work and elsewhere. But the field will keep changing. It's similar to the situation with website development. It used to take an experienced developer to code websites, but now you can build your own with Wix, Squarespace or similar. However, there's still a need for web developers, even a lack of them. The same move towards less need for expertise will happen in the AI space. Everyone will be able to make simple AI solutions with a

few clicks but the demand for more complex solutions will grow even faster.

Some businesses feel threatened by AI. They've been scrambling to hold on to their market as AI can substitute their services. I've seen cases where clients told a company supplying a service that they'd soon no longer need their solution as AI would be able to do what their product does. The clients failed to consider that software products come with specialised user interfaces that still hold a large part of the value. Such companies have started building AI into their products and will likely be fine.

I am convinced that AI will change the world for the better. Yes, it comes with some ethical problems and inconveniences that will be hard to deal with. Deep fakes and clever LLMs in the hands of criminals can scam people on a huge scale. This will take some effort to fix.

But when we look at the big picture, AI has huge advantages. It will help us diagnose and cure diseases faster, find new vaccines, and eliminate tedious and boring work. It will also mean a huge increase in wealth for everyone – not just the rich. Usually, Western economies grow at 2-3% a year, but recently a study found that AI will most likely make all sectors improve productivity to 7% a year, making it more likely to be the new normal for economic growth.[1]

AI isn't about killer robots going on a killing spree. It isn't about computers conspiring to take over the world. It's about using the intelligence and information we already have in a better, more efficient, more productive, and more beneficial way. It will help us learn more easily, have easier and more equitable access to information, and be more productive.

Change is always worrying, and some people fear the changes AI is bringing. This is often based on misunderstanding what AI actually is and what it can do. We are always in control of AI, and with it we can build a better world.

ACKNOWLEDGMENTS

I want to extend my heartfelt appreciation to everyone who contributed to the creation of this book.

Firstly, I am deeply grateful to Hans Peter Bech for his invaluable role as the editor, providing guidance and expertise throughout the editorial process.

Special thanks to Kim Farnell for her diligent efforts in ensuring the content adheres to British English standards, enhancing the readability and accessibility of the text for all our readers. Jelena Galkina's talent in designing the captivating cover art has indeed brought the book's essence to life, captivating readers from the moment they lay eyes on it. Koen Campman's skilful illustrations have enriched the pages of this book, adding depth and visual appeal to the content within. Emma Crabtree's meticulous proofreading has been instrumental in ensuring the accuracy and clarity of the text and contributed to a polished final product.

I also extend my gratitude to my colleagues at Todai for their unwavering support, dedication, and willingness to engage in experimentation, challenges, and testing. Your input and feedback have been invaluable in refining the methods presented in this book, leading to significant improvements and advancements in our field.

This book stands as a testament to the collaborative efforts of a talented team, and I am genuinely grateful for the contributions of each and every individual involved.

Thank you all for your passion, expertise, and commitment to excellence.

ABOUT THE AUTHOR

In 2015, I co-founded Paperflow, a SaaS company with an AI-based product. I worked there as a product manager and later as CEO. I now run Todai, an AI consultancy in Copenhagen. We have delivered AI solutions to some of Denmark's largest companies and trained hundreds of individuals in the technology.

Over the years, I have taken on numerous roles, including marketing manager and developer. And I have made every conceivable mistake and detour possible in AI. This has given me a comprehensive understanding of AI's technical and business aspects. It has also helped me identify common problems and develop effective ways to address them.

Over the last few years, I learned much about how to – and how not to – work with it. Now more and more people are beginning to work with AI for the first time, I want to share what I know. Hopefully, it will make your AI experience less painful.

I run my company AI consultancy Todai in Copenhagen. We build AI solutions, offer advice and teach. If you'd like to get in touch, contact me at dan.rose@todai.ai - https://todai.ai/ .

INDEX

A

Accuracy 7, 10, 111, 120, 170, 171, 172, 174

Active learning 9, 146

Adoption 8, 10, 143, 186, 189

AI agent 113, 165, 176

AI strategy 67

Auto AI 67

Autonomous 7, 118

B

Bias15 0, 151

C

ChatGPT 15, 30, 34, 53, 54, 55, 58, 64, 68, 71, 82, 90, 113, 154, 155, 160, 164, 176, 182, 198

Classification 3, 4, 39, 45, 49

Concept drift 180

Confidence thresholds 7, 115

Confusion matrix 10, 173

CRISP-DM 72

D

Data collection 141, 150

Data engineering 135

Data quality 8, 23, 26, 139

Data scientists 159

Dataset 135

Decision levels 7, 117

Decision Making 7, 111

Decision model 7, 112

Desired state 102

Discovery 5, 6, 7, 73, 92, 93, 108, 111

Distribution 135

Domain expert 7, 104

Domain knowledge 103

E

Ethics 152, 182

Explainable 10, 174

F

Feature importance 174, 175

Feedback 10, 179, 187

Forecasting 4, 42, 43, 44, 142

G

Generative AI 3, 8, 34, 90, 113, 143, 179, 200

Governance 10, 181, 182

Guardrails 165

H

Hallucinations 164

I

Image similarity 4, 48

ImageNet 16, 17

Implicit knowledge 7, 105

Intent analysis 4, 52

Interview 6, 7, 87, 107, 108

K

Kahneman 8, 105, 132

Kasparov 10, 196, 197

L

Large Language Model (LLM) 12

Law of diffusion in innovation 193

Loss of control192

M

Maintenance 88

Measure 78, 189

Model-assisted labelling 9, 145

Monitoring 5, 10, 75, 164, 179, 181, 198

Multimodal 4, 56

N

Named entity recognition (NER) 51

Noise 136, 138

O

Object detection 4, 46, 49, 142

Off-the-shelf AI 5, 67, 68

OpenAI 15, 34, 53, 113

P

Personal data 9, 152

Precision 172

Predictive AI 34, 200

Prompt 165

Prompt-engineering 165

R

Recall 172

Regression 4, 41, 42, 117, 142

Retraining 163, 179

S

Sense of control 10, 195

Sentiment analysis 4, 52

Skewed data 171

Stability 15

Stable Diffusion 15, 56, 155, 156

Synthetic data 9, 153, 154, 155

T

Tabular 3, 39, 57, 70

Temperature 165, 166

Testing 7, 10, 110, 177, 178

Trust 191

U

Upstream drift 181

User 5, 8, 79, 122, 141, 177

V

Vision 4, 45

Research shows that 85% of AI projects fail to deliver the expected outcomes. Although AI technology is now robust and promises immense business value, implementations often fail with broken budgets and dispirited staff in its wake.

According to the author, the problem is always the same. The project methods we use for designing and implementing traditional IT do not work for AI. Thus, the high failure rate is not caused by AI technology but by the way we run the projects.

In the book, the author explains how AI differs from traditional IT. He then provides an easy-to-follow step-by-step method to write the business case, plan the project, design the solution, get the user on board, and deliver tangible business value.

Dan Rose Johansen (33) is the founder and CEO of todai.ai, a consulting company helping clients apply artificial intelligence to solve everyday business problems.

From 2016 to 2020, he was co-founder and later CEO of Paperflow, a company that applied artificial intelligence to lift data off printed documents at high speed and low cost.

Dan Rose Johansen holds a BA in Economics, IT, and project management from the Copenhagen Business School.

Finally, a book that shows you how to apply artificial intelligence to everyday problems. The author has deep practical experience with AI and generously shares it with his readers.

★★★★★ Martin Boberg, CEO, Diction

Dan's book brings us a practical approach to reap the potential of AI. Although the technology is sophisticated, he explains it in simple terms and shows how it can be applied to solve critical business problems. I can recommend the book to anyone involved with business optimisation.

★★★★★ Nicolai Hofsø, Director of P2P, Vertical, Mobile & Fintech, Visma e-conomics

Over the years, I have read many books on AI, studied articles, taken academic courses, and spoken to countless fellow AI practitioners. This book sums it up neatly and gives valuable answers on implementing AI into your organisation practically.

★★★★★ Christian M. Prip, Head of Data Consulting, Netcompany

BIASED PUBLICATIONS