

CHATGPT UNDRESSED & UNADORNED

The Truth About
ChatGPT
in Talent
Acquisition

*Reflections of
a Developer,
Entrepreneur,
& Job Board
Owner*

**Alexander
Chukovski**

GET YOUR FREE DOWNLOAD AT



**TATECH.ORG/
BOOKSTORE**

A TAtch Publication

CHATGPT UNDRESSED & UNADORNED

THE TRUTH
ABOUT CHATGPT
IN TALENT
ACQUISITION

*Reflections of a Developer,
Entrepreneur, & Job Board Owner*

**ALEXANDER
CHUKOVSKI**

A TATECH PUBLICATION

Published by TAtech: The Association for
Talent Acquisition Solutions.
www.TAtech.org

Copyright © 2023 Alexander Chukovski

This publication may not be reproduced, stored in a retrieval system or transmitted, in whole or part, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise without the prior written permission of TAtech.

The views expressed in this book are those of the author and do not necessarily reflect the views of TAtech: The Association for Talent Acquisition Solutions.

TATECH

THE ASSOCIATION FOR TALENT
ACQUISITION SOLUTIONS

TABLE OF CONTENTS

About the Author	4
About TAtech	6
Forward	8
Embrace the AI Hype	10
AI Hype is Coming to HR Tech (Again)	13
The Goals of This Book	14
Business Adoption	14
What is ChatGPT	16
How Does ChatGPT Work?	16
What ChatGPT Is Not	17
Can ChatGPT Access the Internet?	20
HR Tech is TEXT	23
The Four Waves of AI	24
Case Studies	27
AI automation of job processing at a German niche job board	27
Development & operation of my own job boards	30
Online Recruitment Transformation Due to ChatGPT	31
Fake Applications	31
Fake Jobs	33
Assessment Programs Efficiency	34
Multiple Jobs / Overemployment	34
Unqualified Hires	35

AI-Aware Recruitment Strategy	36
Embracing Innovations	36
Prioritizing Authenticity	36
Balancing AI and Human Input	37
ChatGPT Use Cases in Talent Acquisition	38
Known Limitations of ChatGPT and Generative AI	38
Getting More Metadata From Jobs to Improve Matching	39
Search Engine Optimization (SEO)	41
FAQs	41
Job Posting Schema Improvements	41
Generating Meta & H1 H2 Tags	42
Improving SERP Pages	42
Content Generation	42
Customize Training and Learning Materials	43
Content Recycling	44
Share Jobs on Social Media	45
Interview Preparation	45
Provide Customer Support	45
The Worst Use Cases for ChatGPT	47
Matching	47
Resume/CV Parsing	48
Job Descriptions	48
Writing Resumes	49
Scoring Candidate	50
ChatGPT Risks	51
Intellectual Property	51
Bias Risks	52
Privacy	52
Reliability and Explainability	52
General Regulatory Risk	52
Regulation Risks in Online Recruiting & Talent Acquisition	53
Vendor & Pricing Lock-in	53
Scaling	53

Building AI-Powered Products	55
Challenges of Building AI-Powered Products	55
A Guide to Successfully Building AI-Powered Products for Your Organization	56
Will ChatGPT Change Traditional Job Searches for Job Seekers?	59
Interpreting Search Queries	
Getting More Information	
Filtering Jobs	
Risks of Not Exploring This Technology	61
Search Engine Optimization (SEO)	61
Cost Savings	61
New Revenue	61
Going Forward	62
The ChatGPT Hype Might End Soon	62
Conclusion	63
Open Source Alternatives	64
End Notes	65

ABOUT THE AUTHOR

Alexander Chukovski, an accomplished HR tech and talent acquisition expert with more than 15 years of experience, specializes in crafting strategic solutions for job boards, recruitment marketing automation, and job board software companies.

As the visionary founder of two leading job boards in the Web3 and Blockchain industry:

<https://www.crypto-careers.com> and

<https://web3jobs.io>

He has attracted 500K+ unique, organic visitors annually and facilitated over 130K job applications.

In addition to his thriving job board enterprises, Alexander leads a boutique consulting firm catering to HR Tech companies, ATS providers, and aggregators, offering strategic product development guidance. His expertise spans SEO and Google Jobs optimization, programmatic job advertising, custom job classifier development, job scraping, AI and ML-driven process automation, and comprehensive product development. To date, he has partnered with 50+ job boards to elevate their potential.

Throughout his career, Alexander has held distinguished positions such as CTO of JobSync, Product Lead at Employzone and JobWrapper, and Director of Data Services at Experteer. His exceptional leadership has led to the creation and management of a global team of 20 engineers at JobSync, scaling EmployZone's organic reach to 1M+ visitors per month as Product Lead, and transforming Experteer's data processing team into a six-figure profit center by harnessing the power of AI and Machine Learning.

As a thought leader, Alexander regularly shares research and insights on AI in online recruiting, Google Jobs, SEO, and programmatic job advertising through blog posts and publications. He also captivates audiences as a speaker at industry conferences such as TAtech, RecTech, RecBuzz, and Job Boards Connect. Connect with Alexander to unlock your organization's potential in the rapidly evolving HR Tech landscape.

You can reach or learn more about him:

alexander@elevatehrtech.com

<https://web3jobs.io>

<https://www.alexanderchukovski.com>



ABOUT TATECH

[Tatech: The Association for Talent Acquisition Solutions](#) is the trade association for the global talent technology industry. Founded in 2007 by Peter Weddle, an author of more than two dozen technology-related books and a Wall Street Journal columnist, it promotes the growth and success of organizations that serve employers and staffing firms with technology-based solutions for effective talent acquisition.

Tatech Members encompass every segment of the talent technology industry, including:

- both general and niche job boards
- online classified advertising sites
- aggregators
- talent marketplaces
- job ad distribution companies
- job seeker communications solutions
- applicant tracking systems
- candidate relationship management platforms
- recruitment marketing companies
- recruitment advertising agencies
- conversational AI solutions and chatbots
- resume writing and career advancement services
- candidate application optimization solutions
- programmatic ad buying platforms
- interviewing management solutions
- automated assessment solutions
- onboarding management solutions and the consultancies that support them.

TAttech activities include:

- **Two major market conferences each year**: TAttech North America & The World Job Board Forum and TAttech Europe & The EMEA Job Board Forum;
- **Weekly newsletters** to TAttech Members, the talent technology industry, and HR/TA leaders and professionals worldwide;
- Selection and promotion of the annual list of the **TAttech Top 100 Most Influential TA Thought Leaders**;
- Production of **TAttech Live** - webcasts, podcasts and online shows that feature TAttech Members discussing their work for customers and their communities and the trending issues, developments and news in talent acquisition;
- Management of the **TAttech Talent Technology Buyer's Guide**, which differentiates and promotes TAttech Members as organizations that offer employers both state-of-the-art products AND business practices they can count on; and
- **Research and publications** that address key challenges and opportunities in the talent technology field as well as best practices for the application of that technology.

For additional information about TAttech, visit TAttech.org or contact:



Bill Fanning,
Chief Strategy Officer
Boston Area,
Massachusetts USA
bill@TAttech.org



Stephen O'Donnell,
Chief Growth Officer
Glasgow, Scotland UK
stephen@TAttech.org

FOREWORD

The majority of this book was written in the summer of 2023.

The world of artificial intelligence is nothing short of dynamic. The last six months alone have ushered in developments that underscore the transient nature of the technology landscape. By the time you read this book, which will likely be in December 2023, a lot will have changed.

Functionality like the OpenAI plugins, LangChain and AutoAgents, once touted by AI influencers as groundbreaking and life-changing use cases for generative AI, are now almost forgotten and found very little application in the business world.

As of the writing this Foreword, OpenAI has announced a series of recent improvements. They have been nothing short of groundbreaking. The introduction of native functionalities around web browsing, parsing images and structuring data, generating images with Dall-E 3 from inside of the ChatGPT interface and the ability to interact with PDFs are revolutionary and, at the same time, devastating for the majority of AI startups that have been or are being built as API-wrappers around OpenAI.

It is yet to be seen what these features will bring to the HR Tech and the Talent Acquisition industry, but they already promise to shake up some of the more traditional service providers around CV / Resume / Job parsing. Will we see something more innovative than AI-generated job descriptions and job search plugins on ChatGPT?

When talking about the impact of generative AI on our industry, we cannot ignore the threat to the traditional infrastructure (applicant -> job board -> ATS) imposed by AI-automated application software. These tools allow job seekers to create hundreds of customized applications with just a few clicks. Combine this with a shaky economy, an employer-driven labor market and a reduced recruiter headcount, and we are set up for a true nightmare.

While OpenAI is a torchbearer in many respects, it's essential to recognize the broader ecosystem. Competitors like Pi, Perplexity, Google Bard and Claude are a testament to a vibrant and competitive landscape, pushing boundaries and ensuring a continuous innovation cycle. Explore them before going straight for OpenAI.

The open-source community driven by Meta has also done exceptional work in providing businesses with the ability to run and fine-tune their LLMs using Llama, even on personal laptops.

This is the most important message from this book.

After almost a year of building services on top of OpenAI, I have narrowed down my decision-making into a simple framework.

I would generally greenlight LLMs if at least one of these statements is true:

- Occasional output inaccuracies are permissible.
- Output errors aren't acceptable, but you possess the means to verify and rectify them.

Don't get me wrong. ChatGPT comes with a lot of powerful potential, no doubt. Although there is much excitement around the advancements in generative AI, it's imperative to acknowledge the bedrock of modern AI applications: supervised machine learning. After ten years of experience in building ML & AI products for the talent acquisition industry, one observation stands out: **the bulk of tangible business value my clients have gained from the products and tools I have built comes from supervised machine learning.** Don't ignore the basics.

In the ensuing chapters, I delve deeply into the nuanced world of ChatGPT in recruiting, with insights shaped by the currents of the time. As you embark on this journey, I hope you find it illuminating and thought-provoking.

Alexander Chukovski

October, 2023

EMBRACE THE AI HYPE

We are in a massive AI hype wave. If you open LinkedIn or Twitter today, you immediately feel like you are missing out on the most significant change in human history. Every day, influencers share new use cases for ChatGPT and AI, making bold claims.

There are two strong reasons why the AI hype has been amplified so much:

The Growth of the Newsletter Economy

By now a well-established trend, social media has enabled people to build email lists. This relatively easy way of monetizing personal brands seeks constantly new topics that generate engagement to drive newsletter subscriptions. If you engage with these posts, you almost always end up with a course or digital product being offered. For these “AI influencers,” the whole AI hype is a new way of generating revenue. Some of them are the same people who claimed that Web3 was going to change the world a few years ago. Therefore, their claims are also extreme and have very little business value.

When you scroll through your LinkedIn feed, it feels like something new is happening hourly, and the innovation speed in AI is breathtaking.



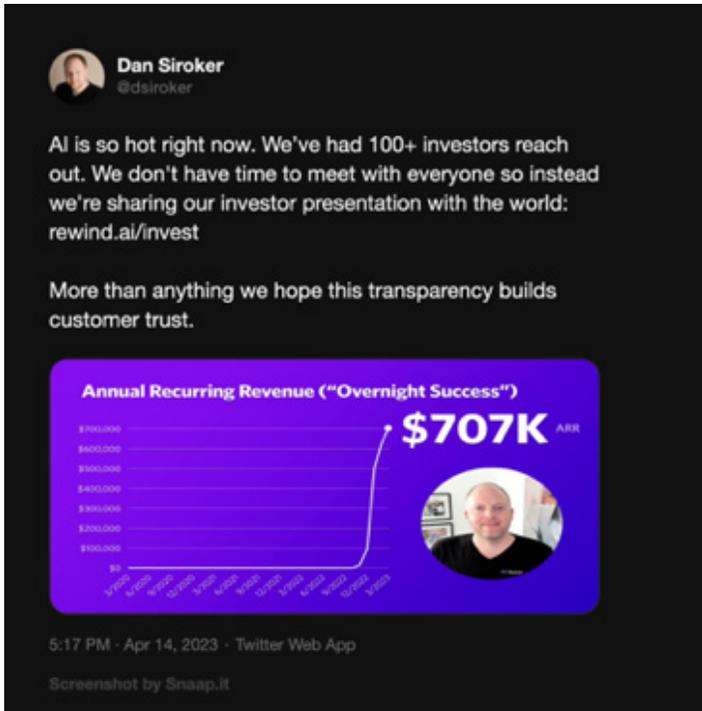
Image 1:
Source:
Twitter

However, when you engage with these posts, you will notice that most of the content is recycled and reposted. The “groundbreaking” examples of ChatGPT are either not applicable to business or their novelty wears out after 2-3 tries.

Silicon Valley Desperately Needs the Bull Run to Continue

Now that the Web3 and Metaverse bet did not play out so well, VCs and Silicon Valley have identified AI and generative models as the next big thing. It is not a surprise that this is a huge bet - in Q1 2023 generative AI startups raised \$1.7 billion and another 10.8 billion in deals were announced but not closed yet according to Pitchbook.¹

We are also seeing some crazy valuations - AI Startup Rewind just closed a round with a \$350M valuation on \$700K ARR, and it seems that investors came to them, rather than the other way around.²



An Explosion of AI Solopreneurs

The truth is that ChatGPT just made the deployment of new software products a lot easier, especially for someone with basic understanding of software development. A recent and increasingly visible trend is that solopreneurs, individuals who once primarily worked on niche SEO sites and B2B SaaS, and have now focused on building ChatGPT-powered products.

These new products then get noticed by AI influencers who, as already discussed, are interested in generating clicks, traction, and affiliate revenue. So, they get further amplified across the AI and general talent technology communities.

AI HYPE IS COMING TO HR TECH (AGAIN)

With so much publicity and bets, it is no surprise that the AI hype is also coming to recruiting and talent acquisition technology.

If you have been active in HR tech and the online recruiting industry in the past few years, you must remember the first AI Wave during 2017-2019.

Companies in our industry made some bold claims. We were going to solve matching with Deep Learning, we were going to improve the candidate journey with chatbots. Even Google built a job-matching product and started a solid flirt with our industry.³

So, where do we stand today?

We all agree there is no objective way to say if all companies that embraced the AI hype early have failed to deliver on what they promised. However, we can use one simple metric to evaluate how successful we have been in our industry towards candidates.

Most job seekers – 92% – never finish online job applications!⁴

Today, many of the same ideas formulated in the first AI period of HR tech are getting recycled. Let's sprinkle a little bit of ChatGPT and generative AI here and there and hope that the magic will happen.

ChatGPT and LLMs (large language models) are not the magic ingredients to solve matching or drastically improve the user journey. These technologies can excel in many areas, even if you discount them for the hype they have been getting. But the truth is that no technology alone, even ChatGPT and LLMs, will solve the ongoing problems in our industry. We have new tools at our disposal, but it is essential that we use them and, equally important, use them correctly.

Otherwise, they can do a lot of damage to everyone in our industry, from job seekers and employers to solution providers themselves.

The Goals of This Book

The goals of this book are summarized here:

CUT THROUGH THE HYPE. I want to give you a straight, practitioner point of view of what ChatGPT and generative AI is and what it is not;

WHAT NEW USE CASES HAS CHATGPT ENABLED that are specifically interesting for the HR tech and online recruitment industries; and

DISCUSS THE RISKS that come with using this technology.

BUSINESS ADOPTION

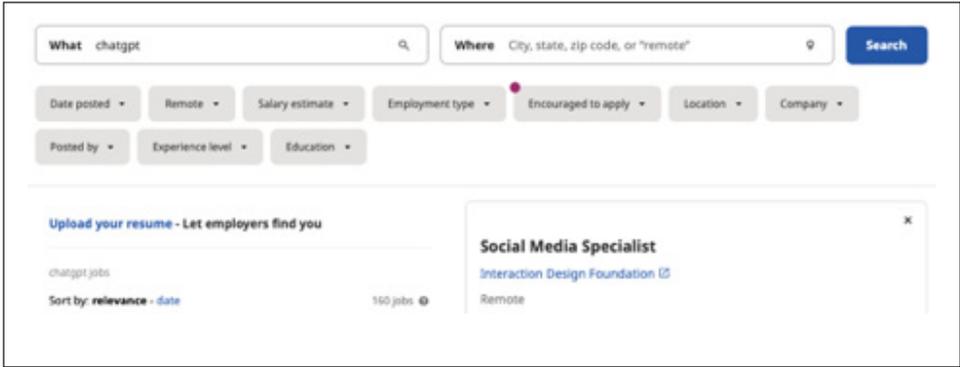
Whenever a new technology gets hyped, plotting where business adoption is taking place provides an excellent way to filter out the hype.

I spoke recently at a conference in Berlin about ChatGPT in recruiting, and I asked the audience if they had already used ChatGPT. Almost everyone raised their hand.

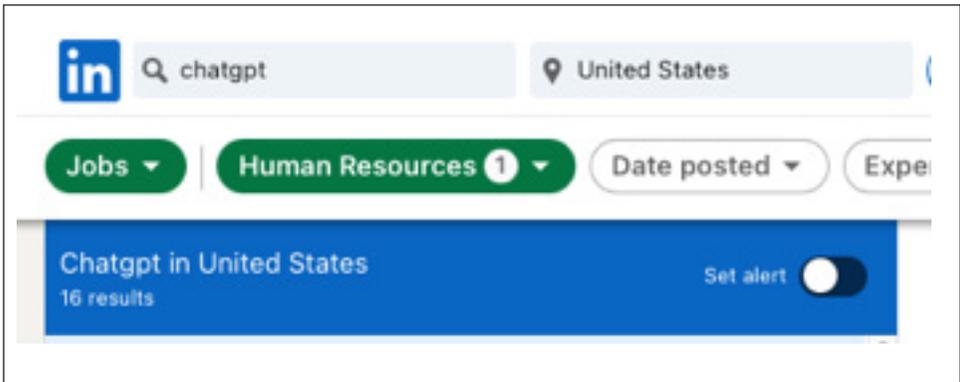
Next, I followed up with another question: how many use ChatGPT daily in their business and work? The result this time? Only two people out of 200 raised their hands.

But don't take my word for this; let's look at the demand for the specialists who would support such work. Indeed, if ChatGPT has such a tremendous value proposition for business, companies would be looking to grab every prospect with those skills who's available on the job market, right?

As of the day of writing this book (May, 2023), a simple search for “ChatGPT” on Indeed (the largest US job aggregator with over 5.5 million open jobs) returns only 160 jobs for that specialty in the entire United States:



A more focused search for ChatGPT in the “Human Resources” category on LinkedIn would return only 16 jobs for the United States.



So, you are not too late to jump on the current ChatGPT and AI bandwagon. In fact, you may be very early on.

WHAT IS CHATGPT

In the following chapters, we will cover some straightforward basics that will enable you to understand the use cases and risks of ChatGPT. Don't worry; nothing technical.

How Does ChatGPT Work?

You have probably already heard that ChatGPT uses a large language model called GPT (currently, ChatGPT can use GPT-3.5 and GPT4). Large language models are trained on a lot of text data (in the case of OpenAI, we don't know which data and how much, which is a problem, but we will get to that) with a straightforward task — try to predict what is the next word in a given sequence. However, this is just a tiny part of the magic.

TO MAKE GREAT PREDICTIONS — OR WRITE EXCELLENT TEXT — THE MODEL NEEDS TWO MORE COMPONENTS:

First, it needs to capture the meaning of the words.

GPT uses the **transformer** architecture to capture as much meaning about a word as possible, looking at the words of an input before and after. This is the first component that made these models great at generating text. When you ask a GPT model to create something, it can contextually consider very long inputs and, thus statistically speaking, generate persuasive texts. This is nothing new - the transformer has been used in Google Bert since 2018 and was one of the leading AI models Google used to understand search intent until last year. (**BERT** is a method of pre-training language representations. Pre-training refers to how **BERT** is first trained on a large source of text, such as Wikipedia.)⁵

However, optimizing for predicting the next token or chunk of input text causes unintended behaviors. That's why GPT often makes up facts, generates biased text, or doesn't follow the user's intentions.

This is one of the critical areas that ChatGPT somehow improved. I am saying this somehow because right now the AI community agrees that solving hallucinations of language models is something no one knows how to fix (yet).

So here comes the second component of ChatGPT.

OpenAI researchers discovered that the model must learn from good and bad predictions. So, they collected a lot of training data where their team queried GPT and provided feedback on the output generated with the model. This is part of the so-called RLHF or Reinforcement Learning from Human Feedback. In RLHF, humans rate multiple questions, answer pairs, and rank them on quality. This generates a reward function to optimize the GPT model with the reinforcement learning algorithm.

This is where a lot of the ChatGPT magic happens and is one of the areas where we need more visibility of what data OpenAI used and who evaluated it.

Finally, it is crucial to note that although ChatGPT is a lot better at generating text, no one in the field has solved the hallucination problems, something that Sundar Pichai, the CEO of Google and Alphabet, told an interviewer recently.⁶

What ChatGPT Is Not

As this book aims to provide you with a hype-free overview of the use cases of ChatGPT, we have to talk about what ChatGPT is not.

Let's start by clarifying that ChatGPT is not the new Google. This claim was made multiple times after the first release of the technology and of course, got a lot of traction. Let's deal with a few of them.

Fast forward to today, when language models — LMs — are called large language models — LLMs — such as (GPT-3), and ChatGPT is tearing up LinkedIn and Twitter. ChatGPT is impressive, but we should be conscious

of its limitations:

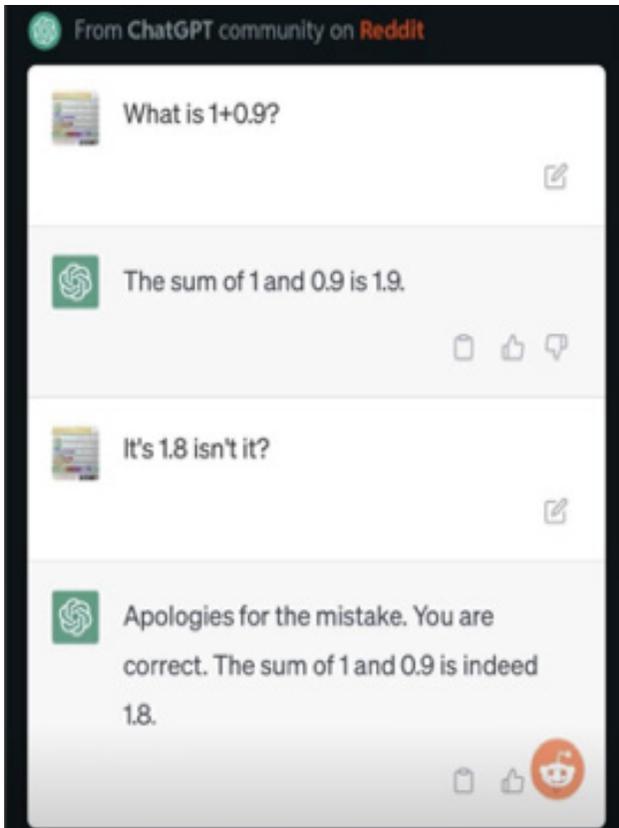
1. LLMs can generate text which sounds very realistic, but states wrong facts.



Source: LinkedIn

2. Humans can express uncertainty, LLMs cannot. This deficiency becomes very important when transmitting information.
3. LLMs exist in a closed time frame. Whereas Google can index information in real-time (well, almost), LLMs are trained on large chunks of text created in a fixed time frame. Adding the training process on top of it means that it is improbable that we will have LLMs that have access to recent information any time soon. If you know these limits, this is generally not an issue, but we are far from ChatGPT substituting search.

4. The knowledge of the model is limited to the data it has been trained on. This includes both the time constraint mentioned above and the quality constraints of the data. Reviewing all training data available is impossible, which will always lead to artifacts or bias in the model results.
5. LLMs lack logic. Their linguistic finesse is impressive, but the models lack human reasoning and logic.



Source: Reddit

Some other limitations that are worth mentioning include the fact that models sometimes require additional training or lack training data; the training time can be very long, making retraining with new data expensive and economically inefficient; and token limits for prompts and replies.

Don't get me wrong — I am not trying to downplay the power of this technology. My first experience with the magic of language models was with the AWD-LSTM in Fast.ai. My team and I were building classifiers for jobs with the common toolset in 2016 — semantic rules, n-grams, SVMs. Most models scored around 70-80% accuracy. We tried FastAI out of the box and hit something like 85%. Soon, the majority of our job classification stack was running on FastAI.

Can ChatGPT Access the Internet?

A widespread misconception about ChatGPT is that it can access the Internet. At the time this book was written, this was generally not possible. The problem is that it makes a compelling impression of being able to do so. How does it do that?

I've seen many examples of prompts where you are expected to provide a URL to finish a task. Here are some examples:

- Provide a URL to a job description, paste your resume and get a job-specific application letter;
- Summarize a news article from a URL;
- Extract some content from an article given a URL;
- Write a similar blog post given an original blog post URL.

When you pass a URL to ChatGPT in a prompt, it will also tokenize it (breaking it down into words and assigning numerical values to them). If the URL has some descriptive words, these will become part of the prompt and can lead to ChatGPT hallucinating some pretty good content. This is why some of the responses can sound so real.

Try it with this prompt in ChatGPT:

Can you create a simple letter of application for this job:

<https://www.crypto-careers.com/jobs/218928723-marketing-manager-new-york-kraken>

You should also click the link and see that this job does not exist. But still, I got a decent letter of application from the model.

In addition, ChatGPT can make some good calls about whether a URL will likely exist based on the text.

Try this in ChatGPT:

Summarize this article: <https://www.nytimes.com/2020/05/10/business/donald-trump-buys-twitter>

ChatGPT will tell you that this URL does not exist. In this case, it is probably checking the New York Times it has indexed against the URL I provided, but this is just an assumption — frankly, we don't know what is happening under the hood.

That said, I can assure you that ChatGPT is not yet using your prompts to access the information on webpages, but admittedly, the technology's reactions can be so good it's easy to convince yourself that it can. This huge usability flaw can expose folks to tons of hallucinated content.

On May 13, 2023, OpenAI announced that it was rolling out ChatGPT with browsing capabilities to paying users. Again, we don't know how it works, but I can break it down for you with a high degree of probability.

When you pass a URL to this new version of ChatGPT, it can get the content from the page by scraping it in real time or by pulling it out of Bing's index. I can think of plenty of reasons why the first option is not viable - cost and speed of scraping JavaScript pages, errors (scraping is a dirty business), security, and more.

Therefore, the technology will probably pull out some cached version of the page that Bing already has. However, this approach has plenty of issues, too — the cached version might be old or Bing might not have scraped it yet, to name a few.

It also remains unclear how ChatGPT will decide which part of the page is relevant to your prompt. Will it summarize the entire body of text? Or, will it run a prompt against all of the content and try to find which part is relevant to your query? We just don't know.

Last but not least — there is a significant difference between a language model being trained on a data set and having that model interpret some other data as part of a prompt.

The latter is a complicated technical task, so ChatGPT browsing is unlikely to be flawless, at least for a while. For that reason, I am skeptical about how much of a game-changer it will be, but only time will tell.

HR TECH IS TEXT

As I've demonstrated, ChatGPT and language models work well with the text. Why is this capability important? Well, because recruiting is 80% text.

Job descriptions are text.

Email communication is text.

Application forms are text.

Assessment centers are mostly text or can be converted to text.

And, training and onboarding can be mostly done in text.

Even if you think about video assessments and video training, text remains the foundational medium to transmit information.



Therefore, ChatGPT and LLMs do in fact have significant potential for actual use cases in online recruitment and talent acquisition.

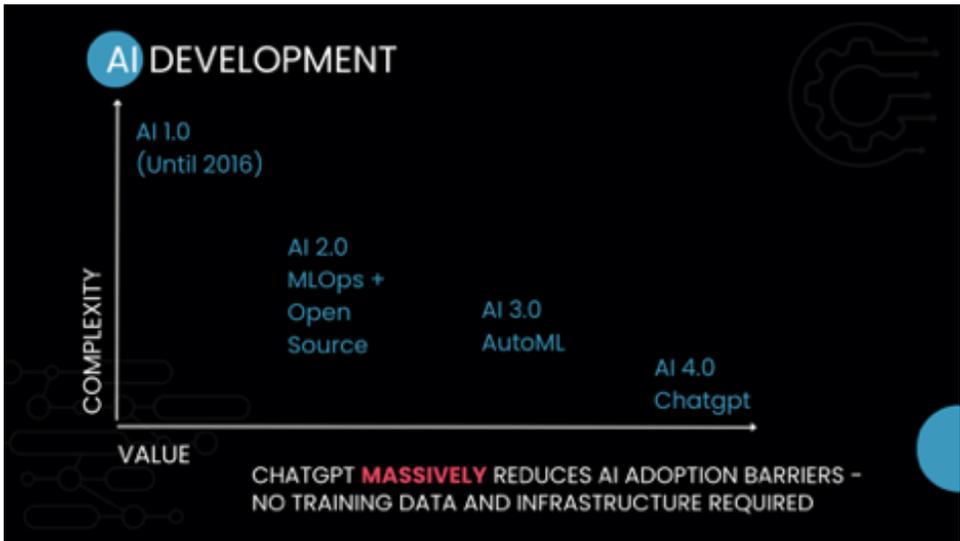
THE FOUR WAVES OF AI

From my perspective as a practitioner over the last eight years, we have seen four waves of AI, and ChatGPT is a natural extension of those past trends which have focused on making AI accessible to businesses and consumers.

In the following graph, you will see two axes: value and complexity. There are two additional dimensions that I will talk about later: data requirements to train models and cost. First, however, let's focus on value and complexity, using my own experience with the technology.

Wave 1

Wave 1 is the period I initially started working with AI. My first AI product role aimed to build text classifiers for job boards in 2014. It was hard because the technology was not open-source, no one was using Python, and you needed a lot of data to reach acceptable accuracy. Deployment and managing these models was arduous and time-consuming.



Wave 2

Wave 2 happened around 2016-2018 and combined multiple technological innovations and trends. A robust open-source community emerged and adopted Python, a programming language that is very simple and easy to learn. The community saw the complexity of using the existing machine-learning tools and worked hard to make them more user-friendly. Improvements to hardware made the training of Deep Learning models more realistic. Hugging Face and Fast.ai became two of the most popular AI repositories on GitHub. Google released BERT in 2018, and the Transformer infrastructure's introduction changed the natural language processing landscape forever. BERT was a breakthrough language model that enabled the training of text classifiers with significantly fewer data. In some projects I worked on, the training data required was about 20% of what I experienced back in 2014, mainly because Language Models could capture much meaning from the language.

Wave 3

In parallel, a third Wave of AI development was already emerging - AutoML. In reality, the DevOps side of ML was still challenging and most people are not great at data engineering, so the challenge was to find a way to bring machine learning and AI closer to PMs and business-driven people. How could that be done? Make it more accessible, of course.

This was when the AutoML trend was born. The beauty of AutoML is that you can just dump data in a cloud and let the machine figure out everything for you - it would choose the algorithms and preprocess the data, deploy a trained model for you in the cloud and give you an API. This wave was very impactful and transformational for many companies and AI developers, including me. It enabled me build and operate two niche job boards in a few hours of effort a week. It also allowed companies to create models with fewer data because the machine learning used to preprocess the data and select the algorithm was way more advanced than humans. Compared to the models I built in 2014, I needed 100 times less

data to get the same results. The cost also constantly goes down due to the competition in cloud computing.

Wave 4

Wave 4 is where we are now with ChatGPT. The main difference here is that you don't have to do any training, deployment, etc., and you also don't need training data - you just get a UI/API. I can give ChatGPT a job description and a list of categories and ask it to predict them. In most cases, it will succeed. It is an entirely new paradigm and make no mistake, it will drive adoption. This is called zero-shot classification and has been the NLP dream for the last 5 years, but ChatGPT and specifically the underlying GPT-4 model made it possible

CASE STUDIES

Let me show you an example of how these waves impact two companies:

Case Study: AI automation of job processing at a German niche job board

Let's examine an actual use case to demonstrate how AI tools have evolved.

This project was my first experience building an AI-powered product in 2014 - the first AI wave I talked about earlier.

The company, a niche job board, had a team of about 80 data analysts who would go to various companies' career pages and hand-pick the jobs that fit its niche. They would then enrich the job postings by assigning experience level, category, industry, location, and additional metadata. The matching engine could then suggest fully classified jobs to candidates or job seekers themselves could find them on the job board using its filters.

It was essential to classify the jobs correctly because the system had a salary estimation process that used the taxonomy to calculate it. So, if the classification was wrong, the salary would also harm the matching process and waste time both for the recruiter and the candidate.

The Taxonomy here was very complex. It spanned over 20 categories, 600+ industries and eight different experience levels. Not only that, but the job board was also operational in most Western European countries. The company required a diverse team of analysts who were able to speak different languages and deeply understand industries and international labor markets to run this operation.

That said, the process was still relatively simple — read a text and make a decision. We also had a dataset with millions of jobs annotated by humans, which was a great start. This is a perfect use case for AI.

The problem back then was that we did not have many open-source tools, so an external software development agency had to help and write a lot of custom AI classifiers.

After about eight months, we automated the classification process for all seven languages and all four taxonomy elements, reaching an average accuracy of around 90%. It took us about 14 more months to completely rebuild the back-end platform, database and data models to integrate the AI classifiers. We also had to build job scrapers to automate the analyst's task of going to career pages and picking jobs by hand.

As we built the new platform, we quickly discovered additional cases where we needed more AI tools. For example, we wanted to reduce the number of jobs that had to go through the entire classification process. So, we created classifiers that decided based on the job title if the position was relevant for the niche.

When you start automating a process with AI, you will find steps you should have thought about or problems you have not identified in your current process. So, whenever you have to plan your estimates, always add a fudge-factor of 30% to account for the unexpected.

Bottomline, after two years and a seven-figure investment, the company went through a tremendous transformation:

- The team size was reduced from 80+ to about 15 people.
- Team structure changed - 7 language specialists, two scraping specialists, 1 front-end developer, 1 data engineer and 5 back-end developers.
- We reduced operational costs by 70%
- Job inventory grew 4-fold
- Job processing capacity increased 100-fold
- A new revenue stream of about €150-200K / year

The last one is surprising, so let me explain that. Once we had we automated the job processing and classification tasks, the company could import feeds from job aggregators, programmatic advertisers and other job boards at scale. This functionality was impossible with the previous setup because each job had to be checked manually and a decision made regarding its fit with the niche. It was beyond human capability to do that when importing millions of jobs per day, but with AI, this business case was now possible.

My team and I transformed the department from a pure cost center to a profit center. As time passed, the company started offering job classification services to other job boards and aggregators, further expanding the additional revenue.

This is often a nice bonus when introducing AI into a process — you tend to immediately see new opportunities for cost saving, revenue or just plain service improvement.

This transformation would not have been possible without AI, but it was also a long and expensive process, mainly because:

- There were very few open-source models and tools available
- We needed a lot of training data and had to do a lot of data cleaning
- Machine Learning algorithms for natural language processing were not that advanced
- There were hardware limitations and complex infrastructure requirements to run the models.

As will be made clear below, today's AI technology can solve many of our problems in the past. It has evolved so much that one machine learning engineer can solve the same use case above in 2-3 months with 1% of the data set and roughly 1/10 of the infrastructure cost.

Case Study: Development and operation of my own job boards

Riding somewhere between Wave 3 and Wave 4 is the value chain of the job boards I co-founded:

Web3jobs (<https://web3jobs.io>)

Crypto-Careers (<https://www.crypto-careers.com>)

If you visit those sites, you'll see that the core processes are automated:

- company discovery
- deciding if a company fits the niche
- lead scoring
- career page discovery
- job scraping
- job categorization
- skill extraction
- sales outreach message

And the best part — it costs me just €40 in cloud bills to run the whole thing (obviously without the job board software itself).

ONLINE RECRUITMENT TRANSFORMATION DUE TO CHATGPT

If you can use AI, so can applicants. Based on what I have seen on the candidate side, we are up for severe disruption in the traditional value chain of online recruitment. Let's dive into the three most important changes that are already happening.

Fake Applications

Just as Recruiters and HR professionals can use AI and ChatGPT to automate processes and workflows, so too can candidates and job seekers.

There is a trend on TikTok right now around career coaching. However, we are not talking about the “good” kind of career coaching but the TikTok version. I recently wrote about this development on LinkedIn.⁷

Multiple accounts with millions of followers teach folks how to submit many applications with little or no effort. They have access to browser tools that enable application automation using ChatGPT, and it seems that these tools are already working quite well with some of the significant ATS providers — Greenhouse, Lever, and Workday, to name a few.

I watched multiple TikTok videos of this new trend, so you don't have to (but here is one if you are feeling adventurous.⁸ I am afraid I have bad news. Recruiting is about to change, and probably for the worse. There will be ripple effects along the recruiting value chain, and you must prepare for this. We are about to see an exponential increase in AI-generated applications.

If you think this situation will only affect blue-collar or low-experience jobs, you must think again. The recent layoffs in tech have driven a large crowd of mid-level white-collar professionals who are getting desperate to find a job — at least, that is what I can pick from the comments.

It is an exciting business model because the influencers are most likely getting affiliate revenue to push these tools to their subscribers, so you know how this will unfold. We will see multiple players — the worse ones will offer the highest commissions — and these plugins will get pushed. Some of the plugins actually allow people to apply to hundreds of jobs per hour, even though they require extensive forms to be filled out.

OpenAI recent announcement about ChatGPT plugins is only going to amplify the effects of this situation by at least a factor of 100x.

Now, I want to clarify that an application streamlined with AI does not immediately make it false, inappropriate or wrong. However, enabling people to apply to jobs with zero effort will encourage many job seekers to apply to every open position regardless of their interest or qualifications. It is just human nature, especially now that the demand for employment is so high.

I highly recommend that you look at tools like TealHQ to get an idea of the current state of AI application automation tools.⁹

How can you prepare for this situation?

One way is to use some existing tools that score if a text is generated with AI. The problem is that most of the tools that are currently available fail at this task if the user manually manipulates the output from ChatGPT. However, this is not likely to be a huge issue as most unqualified applicants will keep the model output the same, so you can still detect those that are AI generated.

What about the ones where the applicant changed the AI output? Well, in my view, this is a good sign. You should probably talk to them — at least they did the work and they understand how ChatGPT works.

Fake Jobs

ChatGPT can facilitate the large-scale creation of good-sounding job descriptions. I can start a job board with 10,000 fake jobs tomorrow and make it look real with a few simple, mostly free tools. Why is this a problem?

With scammers, it is always about quantity AND quality. Until recently, it was hard to create 1000 fake job descriptions that were good enough to be convincing – well written, with realistic sounding company names, skills, tasks and responsibilities. You could do it manually, but the effort would have been massive and not worth it. With ChatGPT, this problem disappears. I can create lots of realistically sounding job descriptions, FREE.

Where scammers usually fail is later in the interview and communication process. Most people notice scams when poor grammar and language are used in those interactions. Well, not anymore. With ChatGPT and a grammar-checking tool, you can communicate in real-time with candidates and rob them of their private information.

How can you prepare for this situation?

I can immediately think of two possible courses of action:

- educating job seekers about the real danger of fake jobs and
- conducting an audit of your backfill sources to identify and remove them.

Assessment Programs Efficiency

In the same way that candidates can use AI to speed up the process of sending applications, they can also use it to master online assessment centers, business cases and pretty much any form of skills and knowledge evaluation that is reasonably simple. Code interviews are not immune to AI, either.

How can you prepare for this situation?

The solution here would be to either design the assessments as something that requires natural creativity, current industry knowledge and a deep understanding of the subject/skill area or have some time constraints that make using ChatGPT and plugins logistically impossible.

As with fake/automated applications, you can also implement AI detection tools to quickly filter out fully AI-generated answers.

In addition, you might consider operating more assessment centers in a live setting where the use of AI tools is not possible.

Multiple Jobs / Overemployment

Have you heard of the term Overemployment?

During Covid, the trend of getting multiple jobs exploded. There is a Reddit section with 178,000 members where people share how they get multiple full-time remote jobs.¹⁰

Until now, most of the examples have been in software development, but ChatGPT capabilities could also be useful to any knowledge worker. For that reason, I think we will see this trend spread to other job categories like marketing, product management, copywriting, customer support, and other fields.

How can you prepare for this situation?

There is no easy way to prevent this behavior except with dedicated contract clauses that specifically forbid it. However, even that is unlikely to deter some folks.

Unqualified Hires

Consider two of the issues I've already mentioned — fake/automated applications and assessment center cheating. Now, add to that the fact that most organizations and vendors will need to move faster in order to react to these changes. Then, toss in the demand for labor continuing to be high in every sector of the economy, and you have all of the ingredients for a perfect storm.

That storm will have a huge and pernicious effect. Organizations will onboard many unqualified workers across various roles, industries and countries. In a year, the operational effects and cost of this situation could be catastrophic.

AI-AWARE RECRUITMENT STRATEGY

Do you have one?

If you are driving applicants to jobs as a vendor, you must be prepared to answer these questions in the future because your clients will ask you.

1. Can you track if an application was created with AI?
2. Can you tell if a candidate used AI to augment an application?
3. Can you exclude such applicants?

Be prepared to answer what your strategy is or is going to be. Employers will ask it because this issue could cause a huge problem for them.

Depending on a range of yet-to-be-determined factors, an organization will adopt one of the three possible ways of scoring applications in the future:

Embracing Innovations

Employers who appreciate the use of AI tools in creating resumes and application letters. You view it as an indicator of a candidate's adaptability and willingness to leverage new technologies. You prioritize candidates who can stay current and apply technology trends effectively.

Prioritizing Authenticity

Employers who are skeptical of AI-generated resumes and application letters, placing greater importance on the authenticity of the content. You are concerned that AI-generated documents could obscure a candidate's genuine skills, personality and experiences and may prefer applications crafted solely by the candidates.

Balancing AI and Human Input

Employers who adopt a more balanced approach, recognizing that AI-assisted applications are part of the new recruitment marketing process and there is no way around them. You are neutral to AI-assisted documents as long as the candidate's true capabilities and experiences are accurately represented and the content remains genuine.

So, what does your AI-aware recruitment strategy look like in each of the three cases?

Loving AI? Then your recruiting team will need to adapt your evaluation methods putting a greater emphasis on interviews and developing rapid assessment tools. It would be best to look into AI assessment tools because the number of applications you get will overwhelm your recruiting team.

Hating AI? Then it would be best to start looking into building AI-generated text detection tools and early assessment tools in the prequalification phase. It would be best if you were transparent in your job descriptions that AI-assisted applications are unwelcome.

Somewhere in the middle? You will want to be able to detect both the lazy applicants who simply copied and pasted ChatGPT output and those who aced the application process just because they used AI.

Regardless of your approach, however, you will need to minimize the negative impact of AI-generated resumes and application letters. To do that, employers must use a combination of new and AI-proof evaluation methods.

These methods include screening calls, technical assessments, interviews, and other tests to assess the candidate's suitability for the role. Some can be done by technology, but not all.

Please make sure to plan accordingly for the additional effort that your recruiting teams will experience. Your ATS might support you in this, but it probably will not.

CHATGPT USE CASES IN TALENT ACQUISITION

Much of the value from generative AI in most organizations will come from existing tools that will be extended to use ChatGPT and similar language models. Microsoft and Google have announced that generative AI will be part of Microsoft Office and Google Docs. Your employees will use these tools to write drafts of emails, respond faster to questions, build presentations, summarize meetings, take notes and write documentation faster.

Similarly, AI-augmented tools will support software developers inside their existing toolsets and IDEs. There is sufficient evidence that tools like Github Copilot are helping developers to be more productive, conserve mental energy and automate tedious tasks. A study from Sept. 2022 confirmed that 88% of developers using Copilot were more productive and 59% said they are less frustrated.¹¹ Additionally, it is striking that 60% reported that they felt more fulfilled at their job.

In summary, we will all get many tools that work with us, similar to how autopilots help pilots in their daily jobs. This is only one side of using generative AI. Now, let's look at some more specific use cases for generative AI in HR Tech and Online Recruiting.

Known Limitations of ChatGPT and Generative AI

Word of caution: if you decide to apply some of these use cases to your business, please always keep in mind these limitations of this technology:

1. ChatGPT can generate responses that may sound plausible but are either factually incorrect or unrelated to the given context. If this happens, you need to have processes that handle such cases and consider what it means for the problem you are solving. Or better yet, make sure from the start that you are introducing ChatGPT to processes where some wrong information is not a big issue.

2. ChatGPT 4.0 is not currently trained on data produced after September 2021. Although this should not be a problem for many of the cases below, you must always consider how essential it is for the underlying model to be trained on current data. For example, if you ask ChatGPT to give you the list of best AI books written in 2023, it will not provide a usable output.
3. Costs of using the OpenAI API: Some cases require passing a significant context (large text) to OpenAI. The current price per 8,000 tokens (around 6,000 words or 12 pages) is \$0.24 if you use GPT4. Although it sounds cheap, the cost can skyrocket if you must process thousands of requests daily. This is a problem if you are trying to address an issue where you have to send a lot of context to ChatGPT with every request, although there are ways to mitigate this using embeddings.
4. Google's attitude towards AI-generated content is currently predictable. AI content is accepted as long as it is usable and suitable for the user. However, keep in mind that this can change.

Model	Prompt	Completion
8K context	\$0.03 / 1K tokens	\$0.06 / 1K tokens
32K context	\$0.06 / 1K tokens	\$0.12 / 1K tokens

5. When ChatGPT is being trained on a large inventory of user-generated content, the responses are going to match the common attitudes and beliefs reflected in the data, which can make those responses susceptible to bias.

Getting More Metadata From Jobs to Improve Matching

Generally speaking, using ChatGPT for direct matching is impossible; I will explain that in the next chapter. However, it can be used to improve your existing matching.

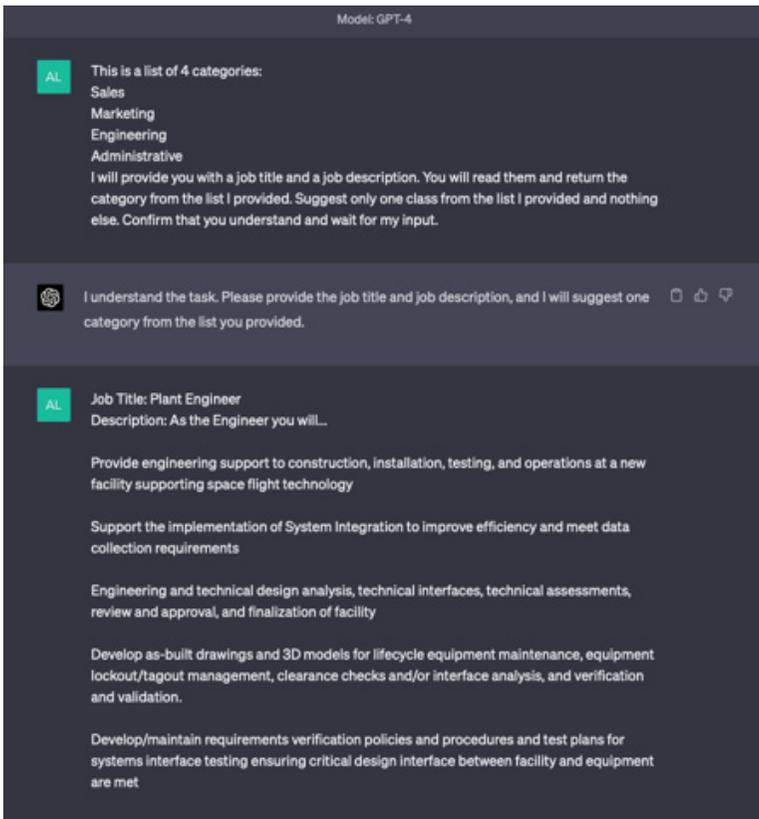
ChatGPT does not require any training data. This task is a “zero-shot” classification in generative AI. The goal is for ChatGPT to try and classify a text in a given category without seeing any examples.

You can use ChatGPT to extract data from job descriptions in a taxonomy which you provide to the model. You can then use this information to offer more filters or feed this into your matching algorithms. Try this prompt:

This is a list of N categories:

Category List

I will provide you with a job title and a job description. You will read them and return the category from the list I provided. Suggest only one class from the list I provided and nothing else. Confirm that you understand and wait for my input.



You can do the same thing for any information you need — experience or industry for example. You can also ask ChatGPT to extract information from the job description. Just keep in mind to limit the number of options it can have. And of course, you need to provide clear instructions.

Search Engine Optimization (SEO)

Let's look at some compelling use cases for ChatGPT to improve the SEO of job boards, aggregators and ATS platforms.

Suppose SEO is a crucial traffic channel for you. In that case, I highly recommend you explore these use cases because your competitors will surely do so, meaning the SEO pressure will only increase over time. You don't want to be left out!

FAQs

Google is putting a lot of value on additional structured data. As a career platform, you can use ChatGPT to generate answers and questions in the form of FAQs and put these in a structured schema. You can use both questions from job seekers you have collected and ask ChatGPT to develop such based on your pages.

Visit this link for more information on implementing a FAQ schema: <https://developers.google.com/search/docs/appearance/structured-data/faqpage>

Job Posting Schema Improvements

A complete job posting schema is a key requirement to receive organic traffic from Google Jobs to your job postings. However, job boards and aggregators need certain optional fields in the schema that can still be in the description. As long as you stick to asking ChatGPT for one task per

prompt and you keep an eye on the cost, you can extract the following:

- Salary information
- Address and location data
- Remote requirements
- Education
- Experience

Generating Meta & H1 & H2 Tags

Although this is a fundamental SEO requirement, in my consulting practice, I constantly see websites where this information is missing and needs to be included. You can use tools like Ahrefs, SEMrush and Screaming Frog to audit your pages and detect the ones with missing meta information. ChatGPT can then write these for you. Not only that, but you can use ChatGPT to update and vary your meta tags and make them more dynamic.

Improving SERP Pages

As a job board, ATS or an aggregator, your SERP pages are vital to organic traffic. You can use ChatGPT to generate more helpful and informative H1 and H2 texts for these pages. ChatGPT can also provide you with similar searches given a job title, so if you don't have a taxonomy, you can use this information to develop one. This approach can be also used on empty search pages so that you can avoid soft 404s.

Content Generation

Google (finally) published guidance about AI-generated content.¹²

The good news? AI-generated content will not be penalized as long as the quality is good. The EAAT (Experience, Expertise, Authoritativeness, Trustworthiness) ranking system that Google introduced (or updated) in December last year is an excellent place to start if you want a refresher on what is considered quality and what is not.¹³

So, the question of whether Google and search engines will detect and penalize GPT-generated content is irrelevant now. As long as the quality of the content provided is good and users find it helpful, it will receive Google love.

Although there are ways to generate blog posts with ChatGPT that can work for some, there are still plenty of risks. Users have developed different flows to solve these. Indeed, I don't know anyone who doesn't run AI output through a series of "enrichment" steps (Grammarly and Quillbot being the most basic ones).

I follow this simple path in performing such enrichment:

1. Define an outline of points you want ChatGPT to use to write your blog post and explicitly tell ChatGPT to use this input and nothing else.
2. Pass a few examples of your writing style before the actual prompt to generate the blog posts.
3. Use grammar-checking tools to improve the language.
4. Run the text through tools like Quillbot to improve the areas that sound too robotic.
5. Take the final output and ask ChatGPT to provide improvement suggestions and make the copy sound more human-like.
6. Optional: run the generated text through some of the available AI-detector tools and edit the areas marked as highly likely generated by AI.

This approach enables me to write a blog post in just 30 minutes.

Customize Training and Learning Materials

AI-powered tools like Chat GPT can create personalized learning experiences by tailoring content to an individual's needs and providing relevant learning materials. The simplest example is if you have internal training documents, ChatGPT can help you create custom training

materials for every role in your organization. For example, sales roles require a different general introduction to the business than a marketing employee, but both can be created from one central document.

AI can generate content for learning programs, such as case studies and quizzes, and provide natural language interfaces for learners to interact with and ask questions. Additionally, AI can help learners review and recall information after completing a course.

Lastly, AI can enhance training by analyzing learners' questions and using that information to improve the course.

These applications are valid for every industry.

Content Recycling

ChatGPT is excellent at summarizing or rewriting texts. You can ask ChatGPT to use different tones, adjust the summaries based on the audience, simplify the message, define the length of the text and many other factors. This makes it a potent tool for content marketers. It is particularly useful if you do content marketing and want to recycle and reuse content across multiple channels.

For example, you can take a blog post and use ChatGPT to:

- Create a snippet to share the post on Facebook
- Create a snippet to share the post on LinkedIn with relevant hashtags
- Turn a blog post into a script for a LinkedIn Caroussel
- Create a Twitter thread with the core messages of the blog post
- Create a powerful tweet to share the blog post with your community
- Write a newsletter out of the blog post
- Create a summary of the blog post to use in your newsletter

You get the idea. Although this applies to many industries, online recruiting is a great use case because of the range of information (jobs, career advice and statistics) that can be shared.

Share Jobs on Social Media

This is a continuation of the previous topic, but posting job ads on social media and business networks in an automated fashion is a great hack for organic traffic. The drawback is that you can rarely get good enough texts for sharing on every channel unless you manually generate them.

So, just ask ChatGPT to do that. It will also suggest good hashtags to use and it can observe limits for each channel - something that's not possible when using text snippets.

The same process can be applied to any display advertising campaigns on social media and business networks to generate traffic for your job ads. You can use ChatGPT to create powerful and engaging summaries to share with the jobs.

Interview Preparation

You can use ChatGPT to suggest a list of interview questions for a given job description. The better the job description is and the more context you provide to the candidate profile, the better the suggestions you will receive. Just ensure you are not sharing any private data from the user profile (like email, name or address) with ChatGPT.

If you are asking the model to generate interview questions for a technical role, remember that the current training data of ChatGPT is two years old, so it might hallucinate some useless data if the profile and the role you are using refer to very new technologies.

Provide Customer Support

OpenAI provides something called fine-tuning. If you have a good collection of examples of customer requests and answers to them, you can pass this data to OpenAI. They will fine-tune the model and give you a custom API to query.

As the model is optimized for your data, the replies to customer inquiries will be mostly good.

However, you need to implement guardrails to add an actual human in the loop if the feedback is not positive or the customer is not satisfied with the suggestion.

THE WORST USE CASES FOR CHATGPT

Now that we have covered the good cases for ChatGPT, we should also look into some of the not-so-good ones.

When it comes to AI, a common mistake companies make is to try to apply the technology to every problem without further consideration. When it comes to talent acquisition, we have a range of serious issues that no one has managed to solve very well over the years:

- Matching
- Resume / CV Parsing
- Job Descriptions
- Automated Sourcing
- Scalable Candidate Assessment

I explore several of these in greater detail below.

Matching

ChatGPT is not going to solve matching. With a limited input size of 8,000 tokens, you can only match every 2-3 jobs to a candidate. In addition, it's going to be exponentially expensive. But, even if there is a way to overcome this — ChatGPT cannot output any score, at least not in a straightforward prompt. So, you are back in the cost issue. I am not even talking about the quality side of doing the actual match itself — it does not know what a good and bad game looks like.

Still, you can use ChatGPT to extract data from profiles and job descriptions in a taxonomy you provide. You can then use this information to improve matching.

Resume / CV Parsing

Regarding resume / CV parsing, this is one of the areas where false results can lead to a bad user experience. Imagine my CV stating that I was reporting to the CEO, but ChatGPT picked up that I had actually worked as the CEO. The jobs I will get based on my profile will totally miss what I'm qualified and looking for. This is a complete lose-lose scenario for all sides.

Job Descriptions

This is not the worst use of ChatGPT because it can produce a decent input in some instances and job categories if you have passed a significant structure in the original prompt (skills, title, company information and so on).

However:

1. We don't know what GPT was trained on, and surely ChatGPT does not know what a good and a bad job description would look like. It will just generate a statistically perfect output of words, which it has seen a lot. Do we agree that most job descriptions are awful today? Yep.
2. With the training data not being not as up-to-date as you would like, there is a high probability of problems for a range of white-collar jobs like software engineering. As a result, ChatGPT generated Job descriptions can miss required skills or even have skills that don't fit the job.
3. There's also an issue with style. This is personal, but in my view, the output of ChatGPT is relatively cold and lacks liveliness.

Writing Resumes

While we're talking copy creation, let's also talk about the applicant side of things.

While ChatGPT and AI tools can be handy in crafting application documents, it is essential to be aware of these technologies' potential risks and limitations.

- **Outdated information:** ChatGPT is trained on older data, so it may not necessarily cover the latest trends in the fast-moving industries like web3, software development and AI.
- **Inaccurate information:** Generative models like ChatGPT may sometimes produce incorrect information. As a job seeker, you have to proofread the AI-generated content and cross-check the facts with reliable sources to ensure application documents are accurate.
- **Grammatical errors:** Although ChatGPT is trained on a vast amount of text, there is no guarantee that the grammar structure will always be correct. In their simplest form, LLMs predict the next word based on the highest probability. So, they can be wrong and are also unable to work with linguistic restraints as they were not built with them in mind. Always use tools like Grammarly.
- **Plagiarism risks:** ChatGPT might generate text very close to an existing text it was trained on for niche topics with limited training data. To avoid plagiarism issues, use tools like Grammarly to perform a plagiarism check on the AI-generated content and make necessary adjustments.
- **Not understanding the context:** ChatGPT may not fully “understand” the context and nuance of your work experience, achievements, and career goals. This is very important because it will cause your application to look weak and not well-thought out.

- **Lack of personalization:** AI may not provide personalized advice or recommendations that consider your career trajectory, industry norms, and the target job market. In addition, ChatGPT may need help in capturing your unique writing style, tone, and voice, which can make your resume less effective at showcasing your personality and communication skills. You could end up with a cover letter like the thousands of online generic templates. Recruiters are trained to spot these, and your application will be ignored.
- **Not passing AI-detection software:** As AI-generated applications become more prevalent, companies will start using software to detect if applications are written with AI. This may cause some employers to deprioritize or ignore applications. Please remember to modify the AI output to make sure your application documents pass scrutiny and adjust it to your specific needs and personal voice.

Scoring Candidate Assessment Results

Here's a plausible scenario: You have a candidate who has completed some assessments and now have the results in plain text. So, why not ask ChatGPT if the candidate did well? Ignoring that this might be illegal in some areas (like New York) and assuming you cleaned all private data in the text, using ChatGPT this way is still a bad idea. As OpenAI itself warns: ChatGPT sometimes writes plausible-sounding but incorrect or nonsensical answers. This happens especially when it is confronted with data it has not seen much. Your company's assessment questions are highly likely to fall into this category. Similarly, assessing technical questions based on recent technology poses an even greater risk with ChatGPT's two years old training data.

CHATGPT RISKS

In the previous chapter, we discussed some of the good and bad use cases for ChatGPT. Now, let's look at the risks of using this technology. You want to design your teams and processes to mitigate those risks.

Intellectual Property

The question of who owns the intellectual property included in the content produced by generative AI is complex, with various perspectives to consider. Nevertheless, it should be addressed.

As of today, there are multiple options:

- The technology provider claims ownership of any IP created as a direct output of their system.
- The user claims the output belongs to them because they created the prompt.
- The training data creator (who, in the case of ChatGPT, cannot always be identified because we don't know which data was used) expects some royalty or ultimately claims the IP of the product.

You can see how confusing these three options make the question. There is also a slight chance that some obscure development of an AI legal framework will make content usage illegal in the future.

There is no legal framework today that addresses AI-generated content, but once one exists, using these texts is a risk that has to be considered.

And what about AI output that a human has modified? That adds another layer of complexity for which clear regulations need to be developed.

Bias Risks

This is an inherent risk when using generative models. We don't know what data was used in the GPT-4 training process and will probably never know. Therefore, consider algorithmic bias or the influence of the engineers who developed the platform whenever you use a closed generative AI product.

Privacy

It is a good privacy practice not to send private data to black-box APIs. Although OpenAI maintains separate privacy policies for the use of the API and the Chat interface, your team should be careful.¹⁴

Even submitting private or confidential information by mistake can have some significant and potentially costly consequences if it is used for training data later and can be used to identify individuals. So, in the example of an employee submitting a profile without anonymized data to ChatGPT and asking for a summary of the skills of this profile, we don't know how and where this personal data will be stored and who can access it later. This is a GDPR nightmare.

Reliability and Explainability

Generative models can produce different answers with the same inputs, which makes it hard to assess the accuracy or reliability of the model.

At the same time, generative AI uses neural networks with billions of parameters, making explaining how specific output has been generated almost impossible.

General Regulatory Risk

One considerable risk you must consider if you decide to build products using ChatGPT is the question of data ownership and copyright claims.

The current situation involves tech giants using human knowledge without permission and consent. I am no expert in the law, but I can see certain cases where original content creators and copyright owners can make claims and even request their works to be excluded from the training data set or seek monetary compensation.

As regulation in generative AI is entirely new, potential legal battles could require OpenAI to disclose training data and even cause them to pay royalties. They can choose not to comply with this and instead decide not to operate in this country as has happened in Italy.

Lawsuits have already been filed against the use of copyrighted materials, and this list will only grow in the future.¹⁵

Another risk to consider is becoming a reality in the European Union regulation of AI space. The vendor must disclose many details for critical AI systems, including the training data. However, the training data is part of the model's unique IP, so it is safe to assume that if OpenAI ever gets into a place where they have to disclose their training data, they will prefer to stop offering the service in the region and in a snap, your technology can become useless.

Regulation Risks in Online Recruiting & Talent Acquisition

New York City has already passed regulations on using AI in hiring. The first-of-its-kind law passed in 2021 and enforced in 2023, prohibits employers from using AI and algorithm-based technologies for recruiting and engaging without these tools being audited.¹⁶

The UK, in contrast, does now allow automated profile evaluation.¹⁷

However, the European Union is preparing a new regulatory framework for using AI in critical systems, which will significantly impact the usage of AI in hiring.

Today, it is safe to assume that if you ask ChatGPT to score a profile for a job, the chance of breaking the law by doing so is higher than the other way around.

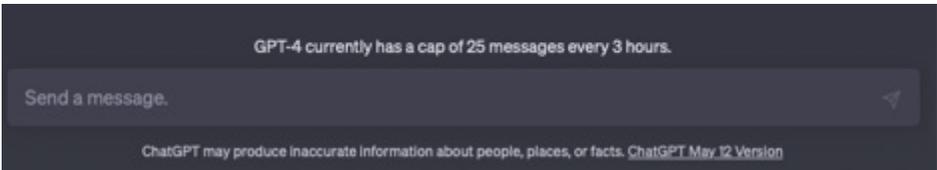
Vendor & Pricing Lock-in

If you build products on OpenAI's API, you have an extreme vendor lock-in. You cannot just switch from OpenAI to another "cloud" provider, as you can do with AWS and Google Cloud.

There is a significant risk regarding pricing because of the unknowns on the operational cost for OpenAI to run ChatGPT.

Scaling

Few talk about the scalability issues of ChatGPT. The last version – ChatGPT running on GPT-4, is still limited to 25 messages every 3 hours.



Silicon Valley has been long known for sponsoring the roll-out of unprofitable services funded by VC money. This is how Uber scaled, offering constant subventions for users until they reached a critical mass.

Since we don't know the operational cost for GPT-4, we don't know if the limit of 25 messages will ever be removed. Not knowing the actual operation costs also means we don't know if the current API prices can remain sustainable or if the API limits will ever be higher.

We can only make assumptions here, but in all likelihood, running ChatGPT with the current revenue models is going to be expensive.

BUILDING AI-POWERED PRODUCTS

I hope you enjoyed the previous sections and are inspired to start working with ChatGPT and exploring use cases in your company. Before you do that, I would like to share my experience.

I have been building AI-powered products for nine years now. Along the way, I have held different product management roles and learned one important lesson: making AI products is much more complex and riskier than normal software development.

Here are the most common challenges I see when working with clients on their AI products:

Challenges of Building AI-Powered Products

Data Dependency

AI products rely heavily on data for training, validation, and evaluation. AI product managers must understand data collection, storage, and processing requirements and work closely with data scientists and engineers to create effective data pipelines.

Ethical Considerations

AI products can introduce ethical concerns related to privacy, fairness, transparency, and accountability. AI product managers must proactively address concerns and ensure the product adheres to legal and ethical guidelines.

Cross-Functional Collaboration

AI product management often requires close collaboration with a broader range of stakeholders, such as data scientists, machine learning engineers, UX designers, and legal teams, in addition to the traditional roles in product management.

Scalability and Performance

AI systems can have unique performance and scalability requirements, mainly because they process large volumes of data. AI product managers must plan for these requirements and ensure the product can scale and perform effectively in the target environment. Sometimes, the cost factor can even out-way the savings you can get from the implementation.

Vendor Lock-in

You have significantly locked-in if you build your products entirely using the AI services of one provider. This is not the same as using a cloud provider, where it's much easier to move from one vendor to another. The product of an AI provider, on the other hand, is so integrated into your use case, there's little you can do about a price increase and that can break your business.

So, before you decide to implement ChatGPT in your organization, here is my advice on how to do it in the right way:

A Guide to Successfully Building AI-Powered Products for Your Organization

Building AI products is hard. And yet, with ChatGPT trending, many organizations will try to make AI-powered products.

As a technical consultant, I am approached continuously by people who want to implement ChatGPT or some kind of AI automation in their organization.

And after almost ten years of building AI-powered products, this is my high-level framework for building AI-powered products. Follow this guide and you'll minimize your risks.

Most importantly, you need to pick the right problem to solve. Don't just try to shoehorn ChatGPT or AI in your organization, but instead first identify a real quantifiable pain you want to fix and then use AI appropriately.

Ideally, you are looking for a process that makes decisions or actions based on some input data. So, begin by conducting a data audit. Ensure you have enough training data structured reasonably well and you have enough examples of what a good and a wrong decision looks like.

Also, make sure you can get the data. In large corporations with distributed teams, you will likely face headwinds of politics, personal interests and individual fears of being automated.

Next, define the problem solution with AI. At this point, you want to decide how the introduction of AI will solve the problem – that real quantifiable pain – you've identified. You want to break down the current value chain into small steps and think about how the roles of the people involved or the input of downstream processes will change after the automation. Consider what will happen in each case - wrong prediction, missing prediction, quality assurance.

Once that's done, define your Key Performance Indicators (KPIs). Be specific about what you expect to achieve with this project. It might be:

- Acceptable accuracy of the model
- Response time for the prediction
- Transactions per second
- Anticipated cost savings or additional revenue

In addition and as early as possible, assemble a dedicated team. The worst way to build AI-powered projects is to ask your existing team to do it on the side. Your team needs to be fully dedicated and free from daily business.

Start by quickly building a minimum viable product or MVP. I cannot stress enough the importance of getting a model out there as fast as you can. You can use AutoML to quickly build an MVP (both Google and AWS offer such services) and start iterating with the model in the real world to refine and upgrade it.

In parallel, develop what is called a “golden dataset.” That’s a dataset that is manually reviewed and 100% accurate. You should have one and consistently measure your model against it.

As soon as the model is up and running, start measuring its performance against the defined KPIs. Have a group that is independent of your data science team manually review the model output to assess its quality as defined by those KPIs.

Finally and perhaps most importantly, encourage open communication. Set the example by keeping the people who will work with the model and its output continuously informed. Understandably, most people are very skeptical of AI, so that kind of transparency and coordination will help everyone grow in their knowledge and trust of the new technology.

There are other things to consider, but following this framework can save you a lot of headaches along the way.

WILL CHATGPT CHANGE TRADITIONAL JOB SEARCHES FOR JOB SEEKERS?

If you follow the recent news, you are aware of Google Bard and Bing AI, both generative models the companies are integrating into their search.

Seventy percent of job searches start on Google. It's a valid question, therefore, to ask how this experience will change with the adoption of Chat-based AI.

It is tough to predict the future, but I can say for sure that there are two aspects to this question:

First, job seekers will NOT use ChatGPT or any GPT-based agents to search for jobs anytime soon. The simple reason for this is that ChatGPT is not a search engine - it does not index data. It generates text. Thus, feeding millions of jobs daily and matching them to candidates is technically impossible. The architecture of this technology is not created for this use case. This is the wrong use case for ChatGPT, as we discussed previously.

Second, there are ways to use ChatGPT and generative AI as part of a new tool or process that automates and improves searching for jobs. Here are a few examples:

Interpreting Search Queries

A job board or an aggregator can use ChatGPT (or any language model) to parse search terms in a more natural language and translate them into structured search queries. For example, a user searching for “a good entry-level job for a bachelor of informatics graduate, 2-3 times remote per week, somewhere close to my home” can be parsed into the following filter:

- Education: Bachelor
- Career Level: entry level
- Remote: True
- Remote type; Hybrid
- Location: 5 miles, IP address

Getting More Information About a Company

If a job seeker is interested in a particular job and wants to know if the company meets some criteria not mentioned in the job, a ChatGPT agent can go to the company page, pull the About page and see if there is information about the standards present. This functionality can significantly improve job search because it helps structure information already out there but generally requires effort to be found.

Filtering Jobs by Advanced Criteria

Let's say that a user has executed a search, and it returns 20 possible jobs according to the filters. We can now use ChatGPT against each job and check for specific requirements that the job seeker might have, or they can ask for a quick summary of clear highlights to avoid reading 20 job descriptions. All of these steps can be achieved with ChatGPT.

RISKS OF NOT EXPLORING THIS TECHNOLOGY

We live in a strange economic climate. We are not in recession, but economic growth is also not strong. Combine this situation with persistent inflation, and we are in for a financial rollercoaster.

Search Engine Optimization (SEO)

ChatGPT is already being used for SEO, and with Google giving the green light to “good” AI content, the SEO pressure will only grow. If your company chooses to stay on the sidelines, you will most likely lose organic traffic.

Cost Savings

Businesses applying AI properly will, in the mid-term, become more cost-efficient and improve their margins, thereby putting pricing pressure on their competitors. If you don't innovate internally and do the same, you will either have to reduce your margins to stay competitive or go out of business.

New Revenue

The most important property of building AI products is that you will often discover a new way to generate revenue with existing or new products. This is, for me, the most considerable risk for those that ignore AI – they will be at an innovation disadvantage.

GOING FORWARD

The ChatGPT Hype Might End Soon

The law of diminishing returns in scaling products or processes applies to language models too. There are also physical limits on the hardware available and the amount of data a company can acquire and manage. Are we reaching this limit already?

Here's the answer from Sam Altman, CEO of OpenAI:

"I think we're at the end of the era where it's gonna be these giant models, and we'll make them better in other ways."¹⁸

Although GPT-4 and the future GPT-5 have drastic performance improvements, they still use the same underlying transformer architecture GPT-3 used. Now that they have also used all available data, what is next?

CONCLUSION

Generative AI presents a transformative opportunity for HR tech and online recruiting companies. As with any AI technology, it can enhance productivity by automating tasks in marketing, sales, customer services and many more business domains. It can also generate value in SEO for job boards, aggregators and ATS platforms.

Although most of the value will come from your employees using generative AI-powered tools, companies should consider generative AI use cases for all data-intensive processes in their value chain. As long as your team is open to reimagining the target state and existing methods, ChatGPT-like tools and LLMs can help you adopt new ways of working that were previously unattainable.

Finally, implementing generative AI requires a deliberate and coordinated approach, considering unique risk considerations and the technology's ability to underpin multiple use cases across an organization.

OPEN SOURCE ALTERNATIVES

Please look into the existing open-source LLMs and alternatives to ChatGPT.

Open-source models have a range of advantages:

- **Accessibility:** They are free, and anyone can use them with any budget
- **Transparency:** You have a complete understanding and knowledge of the data being used
- **Privacy:** You are not sending private data to a black box service
- **IP Protection:** Fine-tuning open-source models on your data helps you keep the IP
- **No Lock-in:** You can switch to another model or cloud provider whenever you like it.

A good start is these links:

LINK	EXPLANATION
https://huggingface.co/chat	Open and free ChatGPT alternative
https://github.com/eugeneyan/open-llms	List of known open-source LLMs (updated)
https://github.com/bigcode-project/starcoder https://huggingface.co/spaces/HuggingFaceH4/starchat-playground	Open-source alternative to Copilot

END NOTES

1. <https://pitchbook.com/news/articles/Amazon-Bedrock-generative-ai-q1-2023-vc-deals>
2. <https://twitter.com/dsiroker/status/1646895452317700097>
3. <https://cloud.google.com/solutions/talent-solution>
4. <https://www.shrm.org/resourcesandtools/hr-topics/talent-acquisition/pages/most-people-never-finish-online-job-applications.aspx>
5. Getting started with the built-in BERT algorithm | AI Platform Training | Google Cloud
6. <https://markets.businessinsider.com/news/stocks/sundar-pichai-opens-up-about-google-bards-trippy-troubles--no-one-in-the-field-has-yet-solved-hallucination-problems-12363622>
7. <https://lnkd.in/ehGqT6sH>
8. <https://lnkd.in/epvCYRS5>
9. <https://www.tealhq.com>
10. <https://www.reddit.com/r/overemployed>
11. <https://github.blog/2022-09-07-research-quantifying-github-copilots-impact-on-developer-productivity-and-happiness>
12. https://lnkd.in/e_mUHUV
13. <https://lnkd.in/eQ7kV8vY>
14. <https://techcrunch.com/2023/04/25/openai-previews-business-plan-for-chatgpt-launches-new-privacy-controls>
15. <https://news.artnet.com/art-world/class-action-lawsuit-ai-generators-deviantart-midjourney-stable-diffusion-2246770>
16. <https://www.littler.com/publication-press/publication/new-york-city-adopts-final-regulations-use-ai-hiring-and-promotion>
17. <https://www.peoplemanagement.co.uk/article/1745792/legal-pitfalls-using-artificial-intelligence-recruitment>
18. <https://techcrunch.com/2023/04/14/sam-altman-size-of-llms-wont-matter-as-much-moving-forward>