

# Occupational Exposure to Text- and Code-Generating Artificial Intelligence in Finland



---

## Antti Kauhanen

ETLA Economic Research, Finland  
antti.kauhanen@etla.fi

## Mika Pajarinen

ETLA Economic Research, Finland  
mika.pajarinen@etla.fi

## Petri Rouvinen

ETLA Economic Research, Finland  
petri.rouvinen@etla.fi

---

### Suggested citation:

Kauhanen, Antti, Pajarinen, Mika & Rouvinen, Petri (25.10.2023). "Occupational Exposure to Text- and Code-Generating Artificial Intelligence in Finland". ETLA Brief no. 127. <https://pub.etla.fi/ETLA-Muistio-Brief-127.pdf>

---

## Abstract

About 19% of Finnish employment is in occupations with at least 50% of tasks exposed to Generative Artificial Intelligence (GAI) with text- and code-generating abilities, such as ChatGPT. Most jobs need some adjustment due to recent advances in GAI, but relatively few will be heavily disrupted.

Our results do not support the "end-of-work" narrative. GAI's long-term impact on human employment is ambiguous; its effects could certainly be positive, especially if GAI turns out to be a sustained source of productivity growth. Whatever the outcome, our findings suggest that a labor market change induced by GAI is brewing and that individuals, organizations, and society all need to make a conscious decision to adapt.

In our view, the biggest risk of GAI in the Finnish labor market is that we will not explore the opportunities it offers with any enthusiasm. Its impact is best faced head-on, and early adopters stand to benefit the most from it. More broadly, the biggest societal risk – in our view – is that we are less and less capable of separating human- and GAI-generated digital content (including audio, images, and video), with a heightened risk of disinformation and highly targeted cyber-attacks.

This research brief replicates the analysis by Eloundou et al. (2023) in the context of Finland.

## Tiivistelmä

### Ammatillinen altistuminen tekstiä ja ohjelmistokoodia tuottavalle tekoälylle Suomessa

Noin 19 % suomalaisista työskentelee ammateissa, joiden työtehtävistä vähintään 50 % on altistunut Chat-GPT:n kaltaiselle tekstiä ja ohjelmistokoodia tuottavalle generatiiviselle tekoälylle. Melko suuressa osassa ammatteja on vähintään lievää altistumista mutta vain harvoissa ammateissa altistuminen on suurta.

Havaintomme eivät tue tekoälyyn toisinaan liitettyä ”työn loppu” -narratiivia. Generatiivisen tekoälyn vaikutus ihmistyön määrään jää häilyväksi. Vaikutus voi hyvin olla positiivinen – varsinkin, jos odotukset tekoälyn merkittävistä, pidempiaikaisista tuottavuusvaikutuksista realisoituvat. Joka tapauksessa generatiivinen tekoäly aiheuttaa merkittäviä työmarkkinamuutoksia, joihin yksilöiden, organisaatioiden ja yhteiskunnan on syytä alkaa varautua.

Valmistumme edessä olevaan murrokseen parhaiten kokeilemalla ja hyödyntämällä generatiivista tekoälyä mahdollisimman etupainotteisesti. Työmarkkinoita laajempi generatiivisen tekoälyn suurin yhteiskunnallinen uhka on se, että aidon ja synteettisen digitaalisen sisällön erottaminen on hyvää vauhtia käymässä mahdottomaksi, mikä kasvattaa valeinformaation ja kyberhyökkäyksiin liittyviä riskejä.

Tässä muistiossa esitetyt havainnot perustuvat keskeisesti alun perin Yhdysvaltoja koskien tehdyn Eloundou ym. (2023) tutkimuksen toistamiseen suomalaisella aineistolla.

---

D.Sc. (Econ.) **Antti Kauhanen** is a Research Director at ETLA Economic Research and a Professor of Economics at Jyväskylä University School of Business and Economics.

M.Sc. (Econ.) **Mika Pajarinen** is a Researcher at ETLA Economic Research.

Ph.D. (Econ.) **Petri Rouvinen** is a Research Advisor at ETLA Economic Research.

KTT **Antti Kauhanen** on Elinkeinoelämän tutkimuslaitoksen tutkimusjohtaja ja taloustieteen professori Jyväskylän yliopiston kauppakorkeakoulussa.

KTM **Mika Pajarinen** on Elinkeinoelämän tutkimuslaitoksen tutkija.

Ph.D. **Petri Rouvinen** on Elinkeinoelämän tutkimuslaitoksen tutkimusneuvonantaja.

---

**Acknowledgements:** We would like to thank the TT Foundation for supporting this research. Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock have kindly provided the input needed to replicate their work, and we are most grateful for their assistance. We acknowledge helpful comments and suggestions by Arash Hajikhani (VTT) and Otto Kässi (ETLA).

**Kiitokset:** Tämä tutkimus on osa TT-säätiön tukemaa hanketta. Kiitämme Tyna Eloundounia, Sam Manningia, Pamela Mishkiniä ja Daniel Rockia tutkimusaineistoon liittyvästä yhteistyöstä sekä Arash Hajikhania ja Otto Kässiä arvokkaista huomioista.

---

**Key words:** Generative artificial intelligence, Technological change, Employment, Labor market, Occupations

**Avainsanat:** Generatiivinen tekoäly, Teknologinen muutos, Työllisyys, Työmarkkinat, Amatit

**JEL:** E24, J21, O33

---

## Both hyped and real

In this brief, we directly replicate analysis by Eloundou et al. (2023), which concerns the United States, for Finland. This work kicks off a project on the impact of artificial intelligence supported by the TT Foundation and conducted by ETLA Economic Research.

The focus is on occupational exposure to **Generative Artificial Intelligence with text- and code-generating abilities** (in what follows, **GAI** for short; Generative Pre-trained Transformers or Large Language Models are commonly used terms closely related to GAI)<sup>1</sup> in the form of ChatGPT by OpenAI and GitHub Copilot by Microsoft.<sup>2</sup>

The key measure in this brief is occupational **exposure** to GAI, without taking a stance on whether human effort is being enhanced or replaced. Exposure merely indicates technological feasibility. Actual deployment in a person's day-to-day work will depend on the economic, ethical, regulatory, social, and technical realities of relevant organizations and – if used as tools, as opposed to completely replacing a human in their job – on the capabilities, incentives, and motivations of the person involved as well.

The first manifestation of institutional adaptation has already appeared. On 25 September 2023, the scriptwriters' strike in Hollywood led to a new labor contract for the Writers Guild of America – a global first with a provision regarding GAI.<sup>3</sup>

From the outset, we note that focusing on exposure ignores two important aspects of any new technology:

- the change in the overall demand for work (or rather, what it delivers), as it becomes more affordable and accessible (including possible budget effects), and
- the creation of completely new types of work, both of which elevate the demand for human involvement.

Upon drafting this brief, we found ourselves in the odd position of both hyping and downplaying the GAI phenomenon. On the one hand, GAI is a fundamental step in the long evolution of artificial intelligence,<sup>4</sup> and is something that will alter the nature of human work. On the other hand, some recent estimates of GAI's impact seem

overblown, and we do not think GAI in itself “changes everything”, even in the longer term. As for Finland and its labor markets, we conclude this brief on an optimistic note.

This brief in English is accompanied by a Finnish version (Etlä Muistio nro 128), which summarizes the key findings here and further elaborates on national policy implications and the potential economic impact.

## Occupations as bundles of tasks

We start from the premise that an occupation is a collection of tasks carried out by

- workers,
- machines, or
- workers using machines as tools.

Technological innovations, such as advances in artificial intelligence, may replace or augment humans in their tasks. Many occupations face both replacing and augmenting innovations simultaneously (Autor, 2022). Whether GAI will lead to replacement or augmentation of human labor in a given occupation depends on the specifics of the job (Felten et al., 2023).

Many occupations have some tasks that can be automated, but only in very few can all tasks be automated (Arntz et al., 2016; Brynjolfsson & Mitchell, 2017). This means that it's unlikely that occupations will disappear altogether but it is likely that the composition of tasks carried out by humans will change (Milanez, 2023).

GAI differs from other recent technological innovations in that it might possibly replace human labor in non-routine tasks (Autor, 2022). In recent decades, new technologies have replaced workers in tasks that can be described by a fixed set of rules, whether the tasks are physical or cognitive. These changes have affected occupations mostly in the middle of the wage distribution, such as clerical and manufacturing jobs. On the other hand, GAI can handle more complex tasks found in areas like business, engineering and science that are found near the upper ends of skill and wage distributions.

Recent experimental studies, some of which are summarized in Table 1, show that GAI improves productivity at the individual level (Brynjolfsson et al., 2023; Noy & Zhang, 2023) and changes human task content towards more creative tasks (Noy & Zhang, 2023). Results from the online labor market for knowledge workers suggest that AI may supplant human labor in writing-related tasks (Hui et al., 2023).

Following the seminal study of Eloundou et al. (2023), the labor market consequences of GAI have been analyzed in China (Qin et al., 2023) and Germany (Oschinski, 2023).

These studies show that its impact on the labor market is more widespread than the impact of other recent technological advances.

**Table 1** Some experimental studies on the impact of GAI

|                            |   |
|----------------------------|---|
| Brynjolfsson et al. (2023) | GAI-based conversational assistant to 5,179 customer support agents.<br>Average productivity (issues solved per hour) increased 14% with the greatest impact on novice and less-skilled workers.<br>Qualitative improvements: better customer sentiment, less need for managerial intervention, and improved employee retention.  |
| Choi and Schwarcz (2023)   | Law school exams with and without access to GAI (GPT-4).<br>Access to GAI improved performance on multiple-choice but not on essay questions.<br>Students at the bottom of the class saw performance gains, while those at the top of the class saw declines.   |
| Dell'Acqua et al. (2023)   | 758 consultants from the Boston Consulting Group, some randomly assigned to use a GAI system (GPT-4 in two different variants).<br>On average, consultants with GAI completed 12% more tasks, 25% more quickly, and at 40% higher quality. Those with initially fewer skills benefited the most.<br>For tasks considered beyond GAI's current capabilities, consultants using AI were 19% less likely to produce correct solutions. |
| Gaube et al. (2023)        | Physicians diagnosed X-rays with or without a visual annotation from an AI or a human radiologist.<br>Receiving annotated advice from an AI resulted in the highest diagnostic accuracy.<br>Non-task experts benefitted more from AI advice.  |
| Girotra et al. (2023)      | Comparison of a GAI system (ChatGPT-4) and students in elite US universities in idea generation.<br>On average, ChatGPT-4 generated ideas of higher quality and variability.<br>Most of the best ideas overall were generated by ChatGPT.   |
| Noy and Zhang (2023)       | 453 professionals assigned to a writing task, with half randomly exposed to a GAI system (ChatGPT).<br>GAI decreased the average time to complete the task by 40%.<br>Average output quality rose by 18% with GAI.<br>Differences in the quality of writing between workers decreased, as GAI helped less skilled writers.  |
| Peng et al. (2023)         | 95 developers performed a standardized computer programming task, some with the support of a GAI system (GitHub Copilot).<br>Average completion time reduced by 55.8% with a GAI system. The least experienced benefited the most.<br>Ability to complete the task was not influenced by the support of a GAI system.   |

**Sources:** The authors' reading of the mentioned articles.

## What's new?

As clearly demonstrated by platforms such as ChatGPT, GAI can take unstructured, natural language as an input and generate new, unstructured output. In doing so, it enters a domain of **creativity** that was considered exclusively human not so long ago. In early 2023, with easy reach and straightforward interfaces, GAI systems sent imaginations racing on what their future incarnations could do.

Of course, there is nothing in GAI that could reproduce human creativity. Rather, GAI systems learn patterns and structures in training data, in dimensionality that even an Einstein-level human mind cannot possibly understand. Via mathematical wizardry, they employ these learned patterns and structures to produce synthetic output.

The fact that a GAI system plausibly mimics human creativity highlights that most of what we label “being creative” are, variations of the same old thing. Thus, the bulk of human creativity is actually “repetitive and formulaic”, as David Ferrucci, the CEO of Elemental Cognition, noted in a Goldman Sachs interview on 7 August 2023.<sup>5</sup>

Regardless of how we split hairs in defining creativity, by adding a new aspect to the earlier digital toolbox, GAI can be used to address new kinds of problems, which in turn exposes new tasks and occupations to technological advance. This brief is an exercise in quantifying that exposure.

Even though the focus in this brief is on text and code generation, we should note that GAI capabilities are turning out to be quite generic and are equally applicable to such diverse areas as audio, biology (e.g., protein sequences), and video. This general applicability might seem puzzling, but once problems are numerically coded and put into a mathematical form, they are all the same to a computer.

## Why now?

The GAI systems that came online in late 2022 and early 2023 demonstrated huge leaps over their predecessors and were readily accessible to laypeople without specialist computing knowledge.

This advance could be directly experienced by anybody who provided input, and if anyone needed further proof, ChatGPT-4 and other systems matched talented humans across a range of academic and professional exams calling for mastery of complex legal and medical concepts (Webb et al., 2023), among other things.

As we discuss in what follows, the key innovation underlying GAI was made in 2017 – it simply took the intervening time to make good of the promise embodied in the technology. The earlier evolution in digital technologies from advanced microprocessors to cloud computing services were necessary building blocks – to paraphrase Einstein,<sup>6</sup> everything stands on the “shoulders of giants”.

## How do they work?

This brief is about the impact of GAI on human labor, not about GAI as a technology. To contextualize our analysis and discussion, we would nevertheless like to make a few notes on the nuts and bolts of GAI.

In 2017, eight Google employees conceived and published (Vaswani et al., 2017)<sup>7</sup> a new architecture for processing natural language known as the “transformer”. It is a computationally efficient way to convey interactions in pieces of inputs such as words in a sentence, notes in music, pixels in an image, or gene sequences in a protein. Transformers process larger sequences of text (or other input) at once giving them a sense of the broader context and the ability to evaluate words by their importance.

One part of a transformer studies the input (say, a prompt in ChatGPT) and another produces output by predicting what is “the most probable continuation” in a sequence, given the input prompt and the output produced so far (conditional on what was in the training data and how the underlying mathematical model is parameterized). In generating words, simply looping the prediction of the next word for a long enough time produces an essay.

In a GAI system, a word is a vector that consists of numbers in the order of hundreds. Together, the numbers span a “word vector space”, in which the similarities and differences of words have a mathematical presentation and one can perform arithmetical operations between words.

Take a word like “cat”. A GAI system trained on Internet online forums “knows” that, along with words such as “dog”, a cat belongs to a broader group of “pets” and that it is a “feline”, whereas a dog is not. It also “knows” that people typically think a dog believes that its caretaker is a god, whereas they also think that a cat believes itself to be a god and that its caretaker is a servant. Somewhere in the vector space it also knows that in casual slang, a “cat” may refer to a person with certain characteristics and being a “cat lady” is considerably different from being just a “lady”. The grand beauty here is that patterns and structures are not defined by a human – the system itself discovers them from the training data, albeit with support of countless human data labelers and learning enforcers. More data is better, although quality might ultimately matter more than quantity.

Timothy Lee and Sean Trott illustrate the potential (in what follows, the latter part referring to a mathematical operation) and important issues (in what follows, gender bias) in GAIs when they note that “Because these vectors are built from the way humans use words, they end up reflecting many of the biases that are present in human language. For example, in some word vector models, doctor minus man plus woman yields nurse.”<sup>8</sup>

Fundamentally, the output of a GAI system is probabilistic and, while based on patterns and structures, somewhat random. A common misconception is to think of a GAI system as an encyclopedia, a search engine, or a logic operator. With extensions (or by controlling other digital tools), it can certainly serve these functions, but the fundamentally probabilistic and random nature of such a system makes occasional misinformation or “hallucinations” a feature, not a bug.

## Calculating exposure

Eloundou et al. (2023) approach is based on coding the various tasks embodied in 1,016 occupations as described in the O\*NET database concerning the United States.

**Exposure**, at the level of a specific **task**, is defined as follows: a text- and code-generating GAI system, if appropriately deployed, reduces the time for a human to perform the task by at least 50% (either used by the human

as a tool or by completely automating the task). At the level of **occupation**, exposure is the proportion of time spent on exposed tasks.

Exposure is considered across three categories:

- no exposure,
- exposure to current capabilities, and
- exposure to current and anticipated capabilities.

Task-level exposures are coded by

- human experts and
- a machine (namely, ChatGPT).

Eloundou et al. (2023) provide estimates for human- and machine-coded lower (current capabilities) and upper bound exposures (current and anticipated capabilities) and for their midpoints. Thus, they consider six measures:

- lower bound, human-coded;
- lower bound, machine-coded;
- midpoint, human-coded;
- midpoint, machine-coded;
- higher bound, human-coded; and
- higher bound, machine-coded.

For comparison, we provide the above six estimates for Finland in the Appendix 1.

To ease discussion and interpretation, in the main text of this brief we provide the average of human- and machine-coded midpoints, although we would like to emphasize that the range of the six estimates is a good gauge of the considerable uncertainty of our estimates.

## Finnish adaptation of the US approach

As noted, the key input to our analysis is the six measures across about one thousand occupations by Eloundou et al. (2023), although, for simplicity in the main text, we distill the six into one as noted above.

Since Finland and the United States use different occupational classifications, we have had to make some adjustments (see Appendix 2 for details), which leaves 410 Finnish occupations to consider.

We note that there are two dimensions of interest:

- Exposure as a proportion of occupations and.
- Exposure as a proportion of workers in Finland.

These two dimensions are different because occupations vary in headcount (at a point and across time; our analysis covers the year 2021, although – due to data availability – estimates in Figure 3 refer to the year 2020).

Since we implicitly assume that the task composition of occupation is the same in Finland and in the US, differences between the Finnish and US findings are solely driven by occupational composition.

## Modestly high GAI exposure quite common in Finland

For Figure 1, occupations are first sorted from the least to the most exposed. The vertical axis represents the proportion of occupations in this order. Even though all occupations are in the 0–100% range, the line cuts the vertical axis at less than 100%, because some occupations have no exposed tasks. The horizontal axis refers to the intensity of exposure. Due to the initial sorting of occupations, it is indeed the minimum proportion of tasks exposed in the following manner:

- Pick a point of interest on the horizontal axis, e.g. 50%.
- From this point, move up to the curve.
- From the curve, move left to the corresponding point on the vertical axis. In this case, it is 21%.

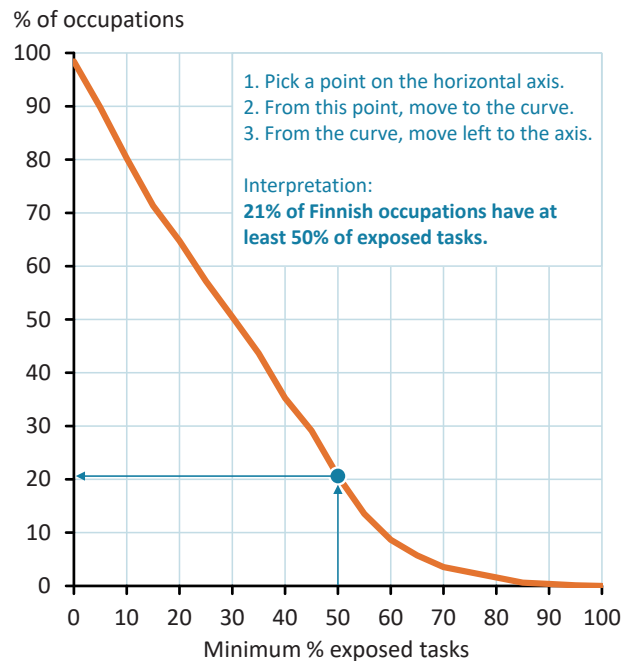
The interpretation: 21% of Finnish occupations have at least 50% of tasks exposed to GAI. Repeating the same exercise reveals that a sizable 65% of occupations have at least 20% exposure and only 1.6% of occupations have at least 80% exposed tasks.

Initial conclusions from Figure 1:

- Few occupations have either very high or very low exposure.
- Modestly high intermediate-level exposure is quite common.

In other words, GAI is about to cause a widespread need to adjust how humans work but generally (outside certain pockets in the labor market), the induced shift does not appear to be highly disruptive.

**Figure 1 Exposure to GAI within Finnish occupations**



**Sources:** Data from Statistics Finland. Calculations by the authors based on Eloundou et al. (2023).


Figure 2 repeats the analysis in Figure 1, but with occupations weighted by their headcounts to examine exposure at the level of the Finnish labor force. As can be seen, **19% of Finnish workers have at least 50% of tasks exposed to GAI**. This is about 2 percentage points less than the same level of exposure in Figure 1, suggesting that more exposed occupations employ somewhat fewer people than the less exposed ones.

Further data points in Figure 2: sizable 68% of employment has at least 20% exposure and just 2% of employment has at least 80% exposure. Immediate conclusions from Figure 2:

- A small proportion of employment has either very high or very low exposure.
- Also, in terms of employment, intermediate-level exposure is quite common.

A comparison to Eloundou et al. (2023) suggests that Finland and the US are quite similar in terms of GAI exposure.

**Figure 2 Exposure to GAI within the Finnish labor force**



**Sources:** Data from Statistics Finland. Calculations by the authors based on Eloundou et al. (2023).

## Can we say more about where exposure is?

Figure 3 starts from the median exposure in Figure 2, according to which 19% of Finnish workers have at least 50% of tasks exposed to GAI.<sup>9</sup> The idea is to evaluate how median exposure cuts across various societal groups.

As can be seen in Figure 3, more educated individuals (Pane A) with higher incomes (Pane B) and socio-economic statuses (Pane C) are more exposed. In terms of business sectors (Pane D), exposure is overwhelmingly in information and communications technology (ICT). In terms of employment geography (Pane E), more exposed individuals tend to live in urban inner-city locales.

With the above, it’s fair to say that exposure to GAI is distinct from previous technological discontinuities, in which groups opposite to the categories most affected in Figure 3 tended to have more exposure.

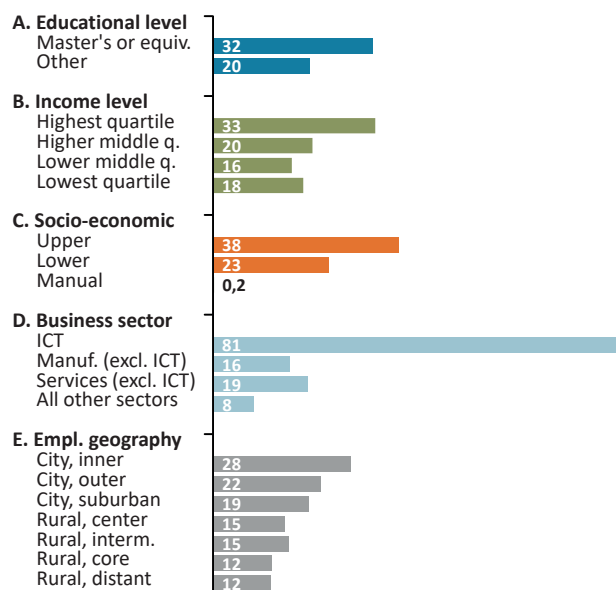
## Discussion

As Figures 4 and 5 in Appendix 1 suggest, the numbers on GAI exposure in this brief come with a considerable margin of error. Besides technological evolution, organizational awareness, and interests, availability of capabilities and resources greatly influence adoption. Furthermore, ethics, privacy, security, and intellectual property are huge issues in both developing and using GAI systems.

As is often the case with digital technologies, optimal timing and mode of deployment is difficult to determine from the perspective of any organization starting to use it. Early adopters pay a price in terms of some wasted effort and resources but may reach understanding, productivity benefits, and market access earlier. Later adopters can study recent history to learn what worked and what didn’t and may also have the luxury of tapping into nicely “canned” solutions rather than developing something from scratch. What is “optimal” for a business is ultimately played out in both input and output markets.

GAI is hugely capable but still fundamentally about reproducing existing material in a probabilistic and some-

**Figure 3 50% minimum employment exposure by...**



**Sources:** Data from Statistics Finland. Calculations by the authors based on Eloundou et al. (2023).



what random manner. Creativity in the sense of truly new-new is still exclusively human and in applications where errors are costly – for example, in legal and medical contexts – humans are likely to remain in the loop for the foreseeable future.

Typically, a company, or rather its leadership, decides on the technologies its employees use in delivering goods and services. With GAI, for example, law firms have been forced to issue bans on grassroots deployment for the time being, before the organization has had sufficient time to determine what tools to use and in what manner. This is an interesting dynamic in adopting GAI. Ultimately, employees might either demand their use or use them in secret to gain an edge.

According to Deloitte’s Digital Consumer Trends August 2023 survey, a quarter of UK consumers have used a GAI system.<sup>10</sup> Around a third of users claim to have used GAI at work. Deloitte (p. 30) notes that “Given the lack of corporate policy and governance, it is fair to assume that a portion of this use was unsanctioned; and without clear, mandated education, employees may have been at risk of sharing confidential information, or failing to recognise hallucination and bias.”

There is a subdomain that is exceptionally suited for GAI deployment – namely, computer programs. It seems to us that GAI could cause an explosion in access to (custom) coding and the (fragments of) software that it produces, which in time should lead to ever-increasing volume of use cases for coding. Without considering “general equilibrium” effects, coding as a human profession seems threatened, but the need for human problem formulation, tailoring, and debugging may well increase more than enough to compensate.

One concern with GAI is that less (training) data rich contexts and less (commercially) interesting application domains might not be able to fully benefit from it. Over time, that would be a growing disadvantage in a competitive environment.

As with any emerging and potentially critical technology, one of the concerns is who commands critical resources and services. The current big tech incumbents seem well positioned, although vibrant startup and open-source activities give hope for a level playing field.

Our study has some notable limitations, which we expand on in Appendix 3. This brief kicks off a sizable research project supported by the TT Foundation. It is our aim to remedy some of the limitations in the coming 14 months or so.

## Conclusion

In an interview in the *Financial Times* (Strauss, 2023), David Autor – arguably the most prominent scholar in the changing nature of work and labor markets – emphasizes the tool and human-enhancing aspects of GAI. In his view, it helps educated knowledge workers to perform at a higher level than they otherwise could. He also emphasizes that, collectively, we determine what the impact of GAI will be: “It’s hard to overstate the importance of designing what it’s there for.” Daron Acemoglu – a renowned expert in the relationship of technology to economic growth – has been a global thought leader in arguing that “we are going in the wrong direction” (Rotman, 2023) when it comes to artificial intelligence. He says that the focus is too much on automating human labor as opposed to augmenting it.

While the “end-of-work” narrative is already somewhat prominent in the GAI context, our findings do **not** support it. Furthermore, due to aging, we need innovations that save human labor and, in doing so, we can employ it elsewhere, where it is harder to replace (e.g., in health and social services).

Nevertheless, our findings suggest that **a change induced by GAI is brewing** and that individuals, organizations, and society need to make a conscious decision to adapt. Understandably, this change has hardly started. In a Boston Consulting Group (June 2023)<sup>11</sup> survey of 13,000 people in 18 countries, 86% of respondents saw a need for upskilling with exposure to GAI but just 14% said that they were receiving training.

The OECD’s (Lorenz et al., 2023) thinking is in line with ours: “The Outlook (OECD, 2023) finds that the net impact of AI in general on employment to be ambiguous. While AI displaces some human labour (displacement effect), the greater productivity it brings (productivity effect) could increase labour demand. AI can also cre-

ate new tasks, resulting in the creation of new jobs for which human labour has a comparative advantage (reinstatement effect), particularly for workers with skills complementary to AI.”

The International Labour Organization (Gmyrek et al., 2023) is in close alignment with the OECD, noting that “... most jobs and industries are only partially exposed to automation and are thus more likely to be complemented rather than substituted by AI... Ultimately, we argue that in the realm of work, generative AI is neither inherently good nor bad, and that its socioeconomic impacts will largely depend on how its diffusion is managed.”

Capital Economics, a UK-based economics consultancy, also follows this line of thinking: “Fears of a big rise in “technological unemployment” are misplaced; if anything the net impact on labour demand will ultimately be positive. But the potential for AI to affect a much wider range of sectors than past technologies means there will inevitably be substantial labour market dislocation.” (CE, 2023).

The history of industry for the last two hundred fifty years has been about exactly the kind of human augmentation and automation we discuss in this brief. Yet the rate of improvement in human well-being has remained stable – and stellar by pre-industrial standards – and mass unemployment has not emerged.

To be sure, GAI exposes a new group of professionals to the potentially adverse effects of automation, although our thinking is in line with a recent McKinsey report (El-

lingrud et al., 2023) noting that “... we see generative AI enhancing the way STEM, creative, and business and legal professionals work rather than eliminating a significant number of jobs outright.”

Agrawal et al. (2023) echo this in stating that GAI systems “...can enhance job prospects and potentially widen the scope for employment of many workers. The neglected mechanism we highlight is the potential for changes in the skill premium where AI automation of tasks exogenously improves the value of the skills of many workers, expands the pool of available workers to perform other tasks, and, in the process, increases labor income...” (from the abstract of the paper).

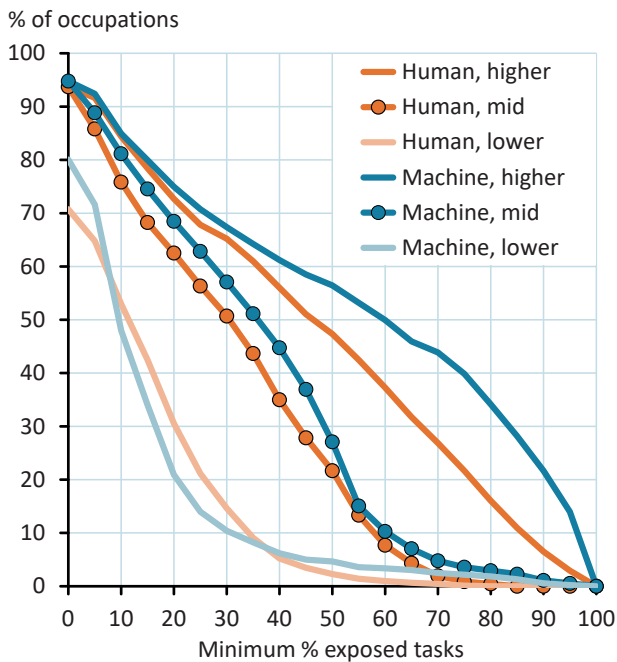
The activities most affected by GAI are those that already take place on a computer screen. Writing computer programs is arguably among the most exposed tasks, as it involves a highly structured language with a well-defined aim and success is easy to measure. While popular interest is in truly generic foundational large language models, short-term economic impact is more likely to come from relatively narrow applications in business-to-business domains.

In our view, **the biggest risk of GAI in the Finnish labor market is that we will not explore the opportunities it offers with any enthusiasm.** Its impact is best faced head-on, and early adopters stand to benefit the most from it. More broadly, the biggest societal risk is that we are less and less capable of separating human and GAI generated digital content with a heightened risk of mis- and disinformation as well as highly targeted cyber-attacks.

## Appendixes


### Appendix 1: A comparison of Figure 3 in Eloundou et al. (2023) with US data and our replication with Finnish data

**Figure 4 Exposure to GAI within US occupations**



**Sources:** Eloundou et al. (2023); re-drawn and re-labeled by the authors.

**Figure 5 Exposure to GAI within Finnish occupations**



**Sources:** Data from Statistics Finland. Calculations by the authors based on Eloundou et al. (2023).

### Appendix 2: Notes on the Finnish data

In moving from the US to the international occupational classification, we were forced to take averages for the occupational groups, which induces a slight “convergence towards the middle” phenomenon.

We converted the data defined for the US Standard Occupational Classification (SOC) to the International Standard Classification of Occupations (ISCO-08 – with virtually one-to-one correspondence to the Finnish Ammattiluokitus 2010).<sup>12</sup> In the US data provided by Eloundou et al. (2023), there are 923 SOC occupations. The number of occupations drops to 410 with mapping to the Finnish classification. Our data nevertheless covers practically all Finnish workers (99.3%) with a valid occupation code at the 4-digit level in Finland in 2021.

### Appendix 3: Limitations

We focused on exposure in the sense of technological feasibility. To have any actual effect, the technology must be deployed. Deployment is held back temporary or permanently by a range of factors that we do not discuss in any detail – including laws and regulations, conventions and standards, attitudes and values, difficulties in implementing complementary organizational changes, and powerful vested interests.

In a sense, our focus on exposure is a trap, as the change in the overall demand for work – by occupation and in the aggregate – gets ignored, as does the creation of altogether new types of work (that fall outside the current, essentially historical, occupational classification).

In building on the US-based analysis, we implicitly assume that the composition of tasks across occupations is the same in both Finland and the United States. Due to factors such as considerable differences in labor market institutions and attitudes towards employee independence and monitoring, the Finland–US mapping is necessarily inexact.

For further discussion on limitations, see Eloundou et al. (2023), section 3.4. Borji (2023) studies ChatGPT’s failures. Risk and challenges of GAI are also discussed in a project report of European Parliamentary Technology Assessment network EPTA.<sup>13</sup>

## Endnotes

- <sup>1</sup> We would like to note that the terminology in this paper departs slightly from that of Eloundou et al. (2023).
- <sup>2</sup> We emphasize that this brief does not concern all forms of artificial intelligence (or digitalization more generally) and not even all forms of GAI.
- <sup>3</sup> <https://www.wgacontract2023.org/the-campaign/summary-of-the-2023-wga-mba>
- <sup>4</sup> Including, e.g., machine learning, computer vision, natural-language processing, and deep learning.
- <sup>5</sup> <https://www.gsam.com/content/gsam/global/en/market-insights/gsam-insights/perspectives/2023/machines-learning-generative-ai.html>
- <sup>6</sup> [https://doi.org/10.1007/978-1-4471-0051-5\\_5](https://doi.org/10.1007/978-1-4471-0051-5_5)
- <sup>7</sup> For an engaging story on the 2017 paper and its aftermath, see <https://www.ft.com/content/37bb01af-ee46-4483-982f-ef3921436a50>. For a good read on the basics of GAI, see <https://www.understandingai.org/p/large-language-models-explained-with>. For the Financial Times visual story on “How generative AI really works”, see <https://ig.ft.com/generative-ai/>.
- <sup>8</sup> <https://www.understandingai.org/p/large-language-models-explained-with>
- <sup>9</sup> Note: Figure 2 uses data from 2021; due to data availability, Figure 3 uses data from 2020.
- <sup>10</sup> <https://www2.deloitte.com/content/dam/Deloitte/uk/Documents/technology-media-telecommunications/deloitte-uk-digital-consumer-trends-2023-deck.pdf>
- <sup>11</sup> <https://web-assets.bcg.com/8c/26/b80dfaa64b1d-92bed7b64d2e19dd/ai-at-work-what-people-are-saying.pdf>
- <sup>12</sup> To do this, we have utilized three crosswalks provided by O\*NET and Bureau of Labor Statistics (BLS). The first crosswalk (<https://www.onetcenter.org/crosswalks.html#soc>) linked O\*NET-SOC 2019 codes to SOC 2018 codes. Because crosswalk from SOC classification to ISCO-08 was available only for SOC 2010 codes, in the second stage we needed to convert SOC 2018 codes to older SOC 2010 codes (<https://www.bls.gov/soc/2018/crosswalks.htm>). Finally, the third crosswalk linked SOC 2010 codes to ISCO-08 occupations (<https://www.bls.gov/soc/soccrosswalks.htm>). The conversion tables unfortunately do not provide a one-to-one match for each occupation. In addition, the original O\*NET SOC codes are at the 8-digit level, whereas the SOC-ISCO crosswalk is defined only at the 4-digit level (as well as occupations in the Finnish data). Due to these aspects, the resulting data from the conversion procedure is more aggregated than the original US data. For multiple matches, we calculated the arithmetic averages of the variables. For instance, for the ISCO occupation *University and higher education teachers* (2310) there are 37 matches in the SOC 2010 (Business Teachers, Postsecondary (25-1011); Computer Science Teachers, Postsecondary (25-1021); Mathematical Science Teachers, Postsecondary (25-1022), etc.). The resulting values for ISCO occupation code 2310 are thus the averages of those corresponding 37 occupations’ values in the US data.
- <sup>13</sup> <https://www.parlament.cat/document/composicio/394503200.pdf>

## References

- Agrawal, A. K., Gans, J. S., & Goldfarb, A.** (2023). The Turing Transformation: Artificial Intelligence, Intelligence Augmentation, and Skill Premiums. *National Bureau of Economic Research Working Paper Series*, No. 31767. <https://doi.org/10.3386/w31767>
- Arntz, M., Gregory, T., & Zierahn, U.** (2016). The Risk of Automation for Jobs in OECD Countries. *OECD Social, Employment and Migration Working Papers*, 189. <https://doi.org/10.1787/5jlz9h56dvq7-en>
- Autor, D.** (2022). The Labor Market Impacts of Technological Change: From Unbridled Enthusiasm to Qualified Optimism to Vast Uncertainty. *National Bureau of Economic Research Working Paper Series*, No. 30074. <https://doi.org/10.3386/w30074>
- Borji, A.** (2023). A Categorical Archive of ChatGPT Failures. *ArXiv*, 2302.03494 ([v8] Mon, 3 Apr 2023 20:02:26 UTC). <https://doi.org/10.48550/arXiv.2302.03494>
- Brynjolfsson, E., Li, D., & Raymond, L. R.** (2023). Generative AI at Work. *National Bureau of Economic Research Working Paper Series*, No. 31161. <https://doi.org/10.3386/w31161>
- Brynjolfsson, E., & Mitchell, T.** (2017). What can machine learning do? Workforce implications. *Science*, 358(6370), 1530–1534. <https://www.science.org/doi/10.1126/science.aap8062>
- CE** (2023). AI, *Economies and Markets – How artificial intelligence will transform the global economy*. <https://www.capitaleconomics.com/ai-economies-and-markets-how-artificial-intelligence-will-transform-global-economy>
- Choi, J. H., & Schwarcz, D. B.** (2023). AI Assistance in Legal Analysis: An Empirical Study. *SSRN Electronic Journal*, 4539836. <http://dx.doi.org/10.2139/ssrn.4539836>
- Dell’Acqua, F., McFowland III, E., Mollick, E., Lifshitz-Assaf, H., Kellogg, K. C., Rajendran, S., Kraymer, L., Candelon, F., & Lakhani, K. R.** (2023). Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality. *Harvard Business School Technology & Operations Mgt. Unit Working Papers*, 24-013. <http://dx.doi.org/10.2139/ssrn.4573321>
- Ellingrud, K., Sanghvi, S., Dandona, G., Madgavkar, A., Chui, M., White, O., & Hasebe, P.** (2023). Generative AI and the future of work in America. *McKinsey Global Institute*, July 26. <https://www.mckinsey.com/mgi/our-research/generative-ai-and-the-future-of-work-in-america>
- Eloundou, T., Manning, S., Mishkin, P., & Rock, D.** (2023). GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models. *ArXiv*, 2303.10130 ([v4] Thu, 23 Mar 2023 21:54:09 UTC). <https://arxiv.org/abs/2303.10130>
- Felten, E., Raj, M., & Seamans, R.** (2023). How will Language Modelers like ChatGPT Affect Occupations and Industries? *ArXiv*, 2303.01157 ([v2] Sat, 18 Mar 2023 19:46:53 UTC). <https://doi.org/10.48550/arXiv.2303.01157>
- Gaube, S., Suresh, H., Raue, M., Lermer, E., Koch, T. K., Hudecek, M. F. C., Ackery, A. D., Grover, S. C., Coughlin, J. F., Frey, D., Kitamura, F. C., Ghassemi, M., & Colak, E.** (2023). Non-task expert physicians benefit from correct explainable AI advice when reviewing X-rays. *Sci Rep*, 13(1), 1383. <https://doi.org/10.1038/s41598-023-28633-w>
- Girotra, K., Meincke, L., Terwiesch, C., & Ulrich, K. T.** (2023). Ideas are Dimes a Dozen: Large Language Models for Idea Generation in Innovation. *SSRN Electronic Journal*. <http://dx.doi.org/10.2139/ssrn.4526071>
- Gmyrek, P., Berg, J., & Bescond, D.** (2023). Generative AI and jobs: A global analysis of potential effects on job quantity and quality. *ILO Working Papers*, 96. <https://doi.org/10.54394/FHEM8239>

**Hui, X., Reshef, O., & Zhou, L.** (2023). The Short-Term Effects of Generative Artificial Intelligence on Employment: Evidence from an Online Labor Market. *SSRN Electronic Journal*. <http://dx.doi.org/10.2139/ssrn.4527336>

**Lorenz, P., Perset, K., & Berryhill, J.** (2023). Initial policy considerations for generative artificial intelligence. *OECD Artificial Intelligence Papers*, 1 (DSTI/CDEP/AIGO/RD(2023)5/FINAL). <https://doi.org/10.1787/dee339a8-en>

**Milanez, A.** (2023). The impact of AI on the workplace: Evidence from OECD case studies of AI implementation. *OECD Social, Employment and Migration Working Papers*, 289. <https://doi.org/10.1787/2247ce58-en>

**Noy, S., & Zhang, W.** (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187192. <https://doi.org/doi:10.1126/science.adh2586>

**OECD** (2023). *OECD Employment Outlook 2023*. <https://doi.org/10.1787/08785bba-en>

**Oschinski, M.** (2023). Assessing the Impact of Artificial Intelligence on Germany's Labor Market: Insights from a ChatGPT Analysis. *Munich Personal RePEc Archive MPRA Papers*, 118300. <https://ideas.repec.org/p/pramprapa/118300.html>

**Peng, S., Kalliamvakou, E., Cihon, P., & Demirer, M.** (2023). The Impact of AI on Developer Productivity: Evidence from GitHub Copilot. *ArXiv*, 2302.06590([v1] Mon, 13 Feb 2023 18:42:46 UTC). <https://doi.org/10.48550/arXiv.2302.06590>

**Qin, C., Jinfeng, G., Huaqing, X., Xingcheng, X., & Yanqing, Y.** (2023). Large Language Models at Work in China's Labor Market. *ArXiv*, 2308.08776. <https://ideas.repec.org/p/arx/papers/2308.08776.html>

**Rotman, D.** (2023, March 25). ChatGPT is about to revolutionize the economy. We need to decide what that looks like. *MIT Technology Review*. <https://www.technologyreview.com/2023/03/25/1070275/chatgpt-revolutionize-economy-decide-what-looks-like/>

**Strauss, D.** (2023, August 10). David Autor: 'We have a real design choice about how we deploy AI'. *The Financial Times*. <https://www.ft.com/content/9c087da3-63d2-4d73-97dc-023025b529aa>

**Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A., Kaiser, L., & Polosukhin, I.** (2017). Attention Is All You Need. *ArXiv*, 1706.03762([v7] Wed, 2 Aug 2023 00:41:18 UTC). <https://arxiv.org/abs/1706.03762v7>

**Webb, T., Holyoak, K. J., & Lu, H.** (2023). Emergent analogical reasoning in large language models. *Nature Human Behaviour*, 7(9), 15261541. <https://doi.org/10.1038/s41562-023-01659-w>

# ETLA



---

## Elinkeinoelämän tutkimuslaitos

**ETLA Economic Research**

ISSN-L 2323-2463  
ISSN 2323-2463

Publisher: Taloustieto Oy

Tel. +358-9-609 900  
[www.etla.fi](http://www.etla.fi)  
[firstname.lastname@etla.fi](mailto:firstname.lastname@etla.fi)

Arkadiankatu 23 B  
FIN-00100 Helsinki

---