

***PUTTING  
IT INTO  
READABLE  
WORDS***

***See your ideas.  
Then write them.***

*after Barbara Minto, The Pyramid Principle, ch. 10*

# ***THE ABSTRACT***

**the 150-word window  
into your entire paper.**

***Readers judge you here.***

# ***CLEAR WRITING HAS TWO STEPS.***


**1.** **DECIDE THE POINT.**

**2.** **PUT IT INTO WORDS.**

***GOOD IDEAS.  
BAD PROSE.***

*the default outcome*

# *Read this slowly.*



*"A primary area of potential improvement is improving cost-effectiveness of field sales-force deployment to reflect the need for redefined selling missions at store and indirect levels..."*

— actual sentence from a corporate memo

# ***THE MYTH***

**Dehydrated prose =  
serious thinking.**

*(it doesn't)*

***GOOD IDEAS DESERVE  
GOOD PROSE.***

*James · Freud · Russell*

rigorous AND readable

***You don't read  
words.***

***You build a picture.***

## ***Read this. Watch your mind.***

*"Near the end of March 1845 I borrowed an axe and went down to the woods by Walden Pond, nearest to where I intended to build my house, and began to cut down some tall, arrowy white pines..."*

— H. D. Thoreau, Walden

# ***MEMORY IMAGE***

***A sketch. Not a photograph.***

*It grows piece by piece as you read.*

***Working memory:***

**7**

**things  
at a time.**

***One image can hold hundreds.***

*Words are slow.*  
*Pictures are instant.*



A is longer than B




B is longer than C



**∴ A is longer than C**

# *Skeletal. Childish. Enough.*



CIRCLE



BOX



ARROW

*"Einstein used shapes like these."*

# ***THREE STEPS.***

- 1. FIND THE NOUNS.**
- 2. DRAW THE RELATIONSHIPS.**
- 3. REWRITE FROM THE PICTURE.**

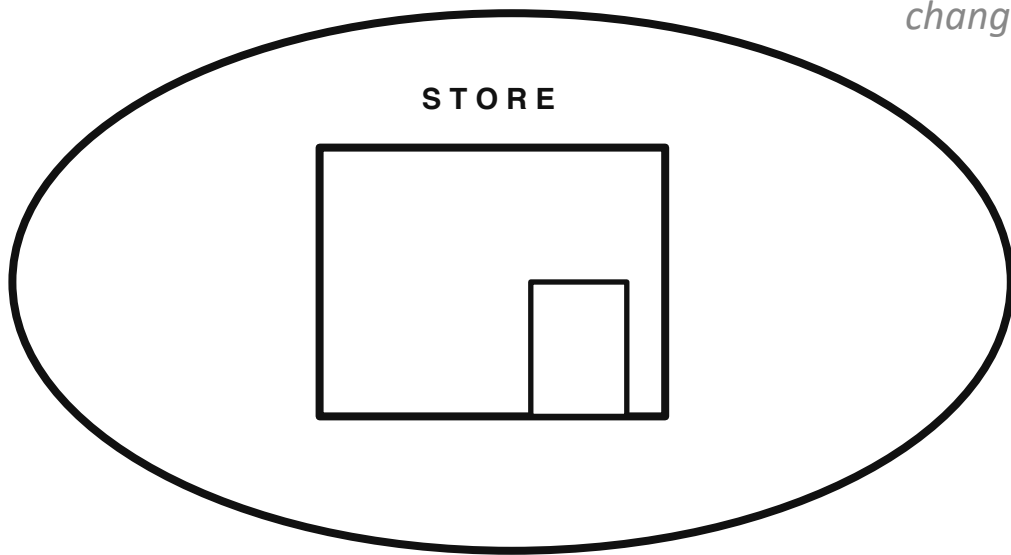
# ***FIND THE NOUNS.***

*"A primary area of potential improvement is improving cost-effectiveness of field **sales-force** deployment to reflect the need for redefined selling missions at **store** and indirect levels dictated by changes in the **trade environment**."*


***3 real things: SALES FORCE · STORE · TRADE ENVIRONMENT***

# ***DRAW IT.***

*new environment*



*"A primary area of potential improvement is improving cost-effectiveness of field **sales-force** deployment to reflect the need for redefined selling missions at **store** and indirect levels dictated by changes in the **trade environment**."*



# ***REWRITE FROM THE PICTURE.***

BEFORE · 34 words

"A primary area of potential improvement is improving cost-effectiveness of field **sales-force** deployment to reflect the need for redefined selling missions at **store** and indirect levels dictated by changes in the **trade environment**."

AFTER · 11 words

***Redeploy the sales force to match the new trading environment.***

# ***YOUR ABSTRACT. YOUR TURN.***

- 1.** Print your abstract.
- 2.** Circle the 3–4 real nouns.
- 3.** Sketch how they connect.
- 4.** Rewrite from the sketch.

*Your reader reconstructs your words on the other side.*

*You owe them **order**.*

To write an abstract, you just need to decide X, M, and Y. That's all.

SITUATION: Problem **X** is important because...

COMPLICATION: In tackling **X**, previous work work did **M** but failed in doing **Y**

PROPOSAL: To tackle Y, we make N contributions:

- 1)
- 2)
- 3)

*Artificial General Intelligence (AGI), which has the capability to perceive, reason, plan, and self-manage tasks within different contexts, is set to completely transform the contemporary workplace. Recognizing its possible impact is not only an urgent need for practice but also an important ethical issue, considering the substantial uncertainties concerning the effects of AGI-induced automation on employment, human agency, and organizations. However, despite much research on workplace risks caused by AI, all of it has been exclusively limited to narrow AI and AI agents without addressing the unique **set of risks** associated with AGI in terms of its ability to autonomously delegate tasks, make autonomous cross-domain decisions, and decide on hiring and firing human employees without any human intervention. To fill this research gap, the present report proposes the following contributions: an AGI definition in the workplace setting based on cognitive autonomy and managerial capabilities of AGI; the multi-level AGI risk framework covering stand-alone failure risk, Human-AGI interaction risks, and systemic risks; a dataset of AGI **risk scenarios** based on O\*NET job tasks as per a specifically designed rubric targeting occupations vulnerable to AGI impact beyond the impact of currently available AI; and finally, a dataset of labeled AGI risk scenarios developed by means of human-LLM hybrid evaluation method in accordance with the EU AI Act risk principles.*

Forecasting **risks** of machines in the workplace is important because it could be an existential threat for the economy. In forecasting those risks, previous work focused on AI but not AGI. To forecast **AGI risk scenarios**, we propose to:

- (1) **Define** AGI based on the literature;
- (2) Generate risk scenarios that are plausible for AGI being used in a **job** based on the previous AGI definition and the job's O\*NET description
- (3) **Evaluate risk scenarios** in terms of whether they characterize AGI as opposed to AI, whether they are plausible in the workplace, ...

*This study investigates how intersecting identity dimensions shape AI harms in workplace settings. A key methodological contribution is our causal filtering system (CQ1/CQ2), which applies two counterfactual questions to each identity marker: asking whether the incident occurred because of the identity, and whether it would have occurred without it. This ensures we isolate true discrimination from generic AI failures. Markers passing this filter are then encoded using a structured template "Because of [identity], the subject was [harmful outcome]" resulting in a highly precise dataset of 488 verified causal markers. Using this refined dataset, our analysis answers four main research questions. First (RQ1), we found that Social Class and Race are the most frequent causal targets of workplace AI harm. Second (RQ2), our intersectional heatmap reveals that harms are significantly amplified when specific identities overlap. Third (RQ3), our media analysis exposes a critical blind spot: the press frequently fails to explicitly name victims' identities (often requiring them to be inferred) and shows disparities in covering privileged versus oppressed groups. Finally (RQ4), we identified an inverse relationship between media virality and incident complexity; the most publicized AI failures often oversimplify the intersectional reality. Taken together, these findings demonstrate that workplace AI harms cannot be understood through single-axis analysis alone. Instead, all four dimensions: frequency, intersectional amplification, media visibility, and representational framing, jointly shape how harm is produced, perceived, and addressed. A comprehensive approach that integrates these perspectives is therefore essential for improving accountability and mitigating bias in AI systems.*

Identifying harms in workplaces is important because harms need to be ideally avoided. In identify harms in workplaces, previous work studied harms caused by a person's individual category and not by the multiple categories of that person.

To identify AI harms caused by a person's multiple categories, we:

- (1) Collect news reporting AI incidents and, from each, extract the categories of the subjects being harmed;
- (2) Develop a metric that quantifies the extent to which the level of risk changes as harmed subject's categories change.
- (3) Analyze the differences between the level of risks for one-category harmed subjects vs. the level of risks for multiple-categories harmed subjects, and found that ... [suprising finding]
- (4) Analyze the differences between the news coverage of one-category harmed subjects vs. the news coverage of multiple-categories harmed subjects, and found that ...

*ChatGPT has become part of everyday work, yet we still lack a clear account of which occupational tasks it supports and what roles it assumes in workplace interactions. Prior research mapped Claude conversations to occupational tasks, but whether comparable patterns hold for ChatGPT—especially in high-EPOCH parts of work, where AI involvement is expected to be more limited—remains unclear. While prior work has examined the roles AI chatbots take in personal and social interactions, a comparable role-based understanding of ChatGPT in workplace settings remains largely missing. We address these gaps with a reproducible pipeline that links publicly available ChatGPT conversations to tasks within the U.S. Department of Labor’s O\*NET database, evaluates task impact in terms of augmentation versus automation, and extracts candidate role descriptions through the lens of high-EPOCH tasks, which represent more human-centered domains of work. These roles form a candidate role taxonomy spanning instructional, analytical, editorial, and relational forms of interaction, suggesting that ChatGPT participates in workplace activity in more varied ways than task-exposure estimates alone can capture. To keep these findings grounded, we complement the automated pipeline with a human-in-the-loop validation protocol designed to separate plausible mappings from spurious ones and to ensure that both task and role labels remain supported by the interaction evidence. Taken together, this framework—and the accompanying reproducible outputs—makes it possible to study not only where ChatGPT shapes work, but also how it participates in it.*

Identifying the occupational tasks AI augments or automates is important because AI could benefit or destroy work. In identifying those work tasks, previous work offered a partial picture: they focused on the tasks exposed to Claude or AI systems but not to ChatGPT, and that's a problem because the tasks exposed to these different systems may be totally different. To map occupational tasks to chatgpt conversations, we:

- (1) Collect a comprehensive dataset of occupational tasks from O\*Net and select those that require EPOCH (e.g., emotional intelligence) skills;
- (2) Collect a large dataset of ChatGTP interactions;
- (3) Design a pipeline that determines whether each EPOCH task is augmented/automated/untouched by ChatGTP, and which role ChatGTP takes (e.g., of an editor, of a friend);
- (4) Evaluate the accuracy of the pipeline with a user study;
- (5) Compare the tasks exposed to ChatGpt with those exposed to Claude and, more generally, AI systems.

*As artificial intelligence increasingly integrates into urban infrastructure, it fundamentally alters the shared environment. However, because these systems are engineered exclusively by humans for human-centric purposes, their underlying algorithms remain structurally blind to non-human life. This anthropocentric blind spot is a critical problem as it unintentionally disrupts how non-human entities move, feed, reproduce, and survive in cities. Since human life depends deeply on these local ecosystems, allowing AI to harm non-human life ultimately threatens our own way of living. Despite the urgency of this ecological threat, **current literature fails to systematically assess AI risks to non-human urban life**. Existing work either focuses **entirely on human subjects or provides theoretical animal ethics frameworks** without applying them for specific urban environments. Related research lacks a methodology to map concrete AI capabilities against the diverse biological and spatial realities of non-human species, making it difficult to quantify and categorize potential harms. To address this gap, we present a systematic framework for evaluating AI impacts on urban ecosystems through four primary contributions. First, we compile a structured **dataset of concrete urban AI deployments** synthesized from literature, practitioner workshops, and incident databases. Second, we develop a novel **taxonomy of 434 non-human urban entities**, categorized across 20 public space types using 10 diverse physical, ecological, and social dimensions. Third, we design a multi-agent large-language-model methodology, grounded in established harm frameworks, to systematically **elicit risk severity and probability** by simulating the perspective of non-human subjects. Finally, we release our resulting risk **dataset** and analytical pipeline as an open-source foundation to enable reproducible research into non-human AI safety.*

Identifying AI risks in cities is important because of the high rates of urbanization and wide deployment of AI in cities. In identifying AI risks, previous work focused on human species but not on non-human species, not least because humans are selfish and design AI for themselves. To identify AI risks on non-human species in cities, we:

- (1) List the variety of AI uses in cities;
- (2) List the variety of non-human species in cities, and categorize them;
- (3) Design a pipeline that determines the list of risks that each non-human specie may suffer when exposed to each AI use;
- (4) Analyze the different risks emerging across all non-human specie categories;
- (5) Evaluate the accuracy of the pipeline.