

WHICH TASKS OF A JOB ARE SUSCEPTIBLE TO COMPUTERIZATION?

Philipp Brandes and Roger Wattenhofer

ETH Zurich, Switzerland

Abstract

In their paper, the two economists Frey and Osborne quantified the automation of jobs, by assigning each job in the O*NET database a probability to be automated. In this paper, we refine their results in the following way: Every O*NET job consists of a set of tasks, and these tasks can be related. We use a linear program to assign probabilities to tasks, such that related tasks have a similar probability and the tasks can explain the computerization probability of a job. Analyzing jobs on the level of tasks helps comprehending the results, as experts as well as laymen can more easily criticize and refine what parts of a job are susceptible to computerization.

1 Introduction

Computerization is considered to be one of the biggest socio-economic challenges. What is the foundation of the recent worries about many jobs being affected by automation [7, 8, 11, 12]?

Why did the last few years see dramatic technological progress regarding self-driving cars [14], board games [20], automatic language translation [1], or face recognition [16]? One reason is big data. While “intelligent algorithms” in the past were restricted to learning from data sets with a few thousand examples, we now have exabytes of data. Learning becomes even more powerful if you combine big data with a highly parallel hardware, stirred by the success of graphics processing units (GPUs). However, both of these technological advancements needed to be harvested, and they are with the advent of so-called deep learning algorithms, which have blown the competition away, starting with voice recognition [10]. As a consumer, you can already witness some of these advancements on your smartphone, but a lot more will come soon. We believe that these advancements will revolutionize white collar work and (with a little help from sensors and

robotics) also blue collar work. In contrast to previous waves of innovation, this time new emerging jobs might not be able to compensate jobs endangered by the new technology.

Computer scientists have been pondering the consequences of AI for a long time; starting with the possibly most famous paper in computer science “Computing Machinery and Intelligence” by Alan Turing, where he proposed the Turing test [21]. An early pioneer of AI, Marvin Minsky, claimed already 1967 that “within a generation ... the problem of creating ‘artificial intelligence’ will substantially be solved” [19]. More recent work in computer science focuses on the effects on the economy and the society [22–24].

In their paper, Frey and Osborne [12], two economists, quantitatively study job automation, predicting that 47% of US employment is at risk of automation. In order to calculate this number, Frey and Osborne labeled 70 of the 702 jobs from the O*NET OnLine job database¹ manually as either “automatable” or “not automatable”. Then, for the remainder of the jobs in the O*NET database, they computed the automation probability as a function of the distance to the labeled jobs.

But the results of Frey and Osborne are opaque, one either believes their “magic” computerization percentages, or one has doubts. We want anybody to be able to easily understand and argue about our results, by incorporating the unique tasks of each job. This additional depth will help laymen as well as job experts to argue about potential flaws in our methodology.

If we know that a job is 100% automatable, we also know that every task of that job must be completely automatable. But what if a job is 87% automatable? Is every task 87% automatable? Or are 87% of the tasks completely automatable, and 13% not at all? We want to forecast which tasks of a job are safe and which tasks are automatable.

In order to calculate the automation probability for a task, we first need to determine its share of a job (Section 4). Based on this, we are able to assign each task a probability to be automated such that the weighted average of the probabilities is equal to the probability of the corresponding job (Section 6.1). During our evaluation (Section 6.2), we discover a few suspicious results in the probabilities by Frey and Osborne, e.g., a surprisingly high automation probability of 96% for the job *compensation and benefits managers*.

We analyze the correlation between various properties of a job and its probability to be automated (Section 7). E.g., we show that there is a strong negative correlation between the level of education required for a job and its probability to

¹O*NET is an application that was created for the general public to provide broad access to the O*NET database of occupational information. The site is maintained by the National Center for O*NET Development, on behalf of the U.S. Department of Labor, Employment and Training Administration (USDOL/ETA); see <https://www.onetonline.org/>

be automated. In Section 8, we try to determine whether the effects of automation can already be seen in our data.

Our complete results can be found at <http://jobs-study.ethz.ch>.

2 Related Work

The current effects of automation have been studied intensively in economics. Most studies agree that some routine tasks have already fallen victim to automation [2, 4, 13]. A task is routine if “it can be accomplished by machines following explicit programmed rules” [4]. With computers being able to do routine tasks, the demand for human labor performing these tasks has decreased. But on the other hand, the demand for college educated labor has increased over the last decades [6, 25, 26]. The effect is more pronounced in industries that are computer-intensive [3]. As a consequence of this, the employment share of the highest skill quartile has increased. In addition to more people being employed in the highest skill quartile, the real wage for this quartile has increased faster than the average real wage. Service occupations, which are non-routine, but also not well paid, have also seen an increase in employment share and in real hourly wage. Thus, both, employment share and real wage, are U-shaped with respect to the skill level [2]. This employment pattern is a phenomenon that is called polarization. This is not unique to the US, but can also be observed, e.g., in the UK [13]. These papers make important observations about the effects that automation already has. Until now, routine tasks are the ones most affected, but more and more tasks can nowadays be performed by a computer. We focus on the future and try to predict which tasks will be automated next.

John Keynes predicted already in 1933 that there will be widespread technological unemployment “due to the means of economising the use of labour outrunning the pace at which we can find new uses for labour” [15]. Automation might be the technology, where this becomes true [7, 8, 11, 12]. “Automation of knowledge work”, “Advanced robotics”, and “Autonomous and near-autonomous vehicles” are considered to be 3 out of 12 potentially economically disruptive technologies [17]. Computer labor and human labor may no longer be complements, but competitors. Automation might be the cause for the current stagnation [7]. There might be too much technological progress, which causes high unemployment. A trend that could be going on for years, but was hidden by the housing boom [9].

The paper by Frey and Osborne is the first to make quantitative claims about the future of jobs [12]. Together with 70 machine learning experts, Frey and Osborne first manually labeled 70 out of 702 jobs from the O*NET database as either “automatable” or “non automatable”. This labeling was, as the authors admit, a subjective assignment based on “eye balling” the job descriptions from

O*NET. Labels were only assigned to jobs where the whole job was considered to be (non) automatable, and to jobs where the participants of the workshop were most confident. To calculate the probability for non-labeled jobs, Frey and Osborne used a probabilistic classification algorithm. They chose 9 properties from O*NET as features for their classifier, namely “Finger Dexterity”, “Manual Dexterity”, “Cramped Work Space, Awkward Positions”, “Originality”, “Fine Arts”, “Social Perceptiveness”, “Negotiation”, “Persuasion”, and “Assisting and Caring for Others”.

The results from Frey and Osborne for the US job market were adopted to other countries, e.g., Finland, Norway, and Germany [5, 18]. This was done by matching each job from O*NET to the locally used standardized name. Due to differences in the economies, a different percentage of people will be affected by this change, e.g., only one third in Finland and Norway are at risk compared to 47% in the US.

3 Model






Figure 1: This figure shows a small example consisting of three jobs represented by circles. The share of each task of a job is shown by its sector. Two tasks are connected by a line if and only if they are related, i.e., similar according to O*NET. Task t_1^1 from job j_1 and task t_2^1 from job j_2 are related as indicated by the line connecting them. Note that this relationship is not transitive. Thus, tasks t_1^1 and t_3^1 do not need to be related.

We are given a set of jobs $J = \{j_1, \dots, j_n\}$. Each job j_i consists of a set of tasks $T_i = \{t_i^1, \dots, t_i^m\}$, where every task belongs to exactly one job. We call two tasks t_i^k and t_j^k related if and only if these tasks are similar according to O*NET. Two jobs



(a) Electro-Mechanical Technicians



(b) Environmental Engineering Technicians

Figure 2: The number of tasks and the frequencies assigned to them can differ significantly even for related jobs.

with related tasks are also called related. An example with 3 jobs is depicted in Figure 1.

O*NET provides us for each task t_i^k with the information how often it is performed. This information was gathered by asking job incumbents and occupational experts. The options are “yearly or less”, “more than yearly”, “more than monthly”, “more than weekly”, “daily”, “several times daily”, and “hourly or more”. O*NET provides a percentage for each of the 7 options. We denote these frequencies of task t_i^k with $f_1(t_i^k), \dots, f_7(t_i^k)$. Since these values are percentages, for every task t_i^k they sum up to 100%, i.e., $\sum_{\ell=1}^7 f_\ell(t_i^k) = 1$.

Each job j_i has a given probability $p(j_i)$ to be automated. We want to use $p(j_i)$ to calculate a probability to be automated for each task of this job.

4 From Task Frequencies to Task Shares

We use the frequencies with which a task is performed to assign each task t_i^k its share $s(t_i^k)$. For every task t_i^k , the share denotes how much time is spent doing this task, such that $\sum_{t_i^k \in T_i} s(t_i^k) = 1$. The frequency values from O*NET do not fulfill this property; their values are very consistent for one job, but they can vary a lot between different jobs and might even seem to contradict each other. An extreme example can be seen in Figure 2. The seven frequency options provided by O*NET are on the x-axis and on the y-axis is the corresponding value of each option.

To make use of the high consistency within a job, we decided that the share of a task is a weighted average of its frequencies, i.e., $s(t_i^k) := \sum_{\ell=1}^7 x_i^\ell f_\ell(t_i^k)$. We want to calculate the job specific coefficients x_i^ℓ . Let us illustrate these coefficients with a simple example. If $x_i^7 = 0.1$, then a task t_i^k that is done exclusively “hourly or more” (i.e., $f_7(t_i^k) = 1$) makes up 10% of job j_i .

We want these coefficients to satisfy a few assumptions. If O*NET states that

a task is done “hourly or more”, then the share of this task should be higher than the share of a task that is done “several times daily”. This translates to $x_i^\ell \leq x_i^{\ell+1}$ $\forall \ell \in \{1, \dots, 6\}$ and $0 \leq x_i^1$.

These constraints neither use that jobs are related nor do they define the coefficients uniquely. Both issues are solved if we require the coefficients x_i^ℓ and $x_{i'}^\ell$ for two related jobs j_i and $j_{i'}$ to be similar. The intuition behind this is that the frequencies of O*NET for related jobs are not independent of each other either, but rather should be similar as well. Occupational experts who have rated the frequency in which a task is done for one job, are likely to have rated the frequencies of related jobs.

The coefficients cannot be identical without violating the other constraints. Jobs have a different number of tasks and the frequencies are task specific. The example in Figure 2 highlights this. It is therefore easy to see that we cannot have the same coefficients for two related jobs and fulfill the equality $\sum_{t_i^k \in T_i} s(t_i^k) = 1$ for both jobs simply because the number of tasks can differ a lot.

Thus, we allow a bit of slack in the coefficients of related jobs. We use the variable $x_{i,i'}^\ell$ to express the difference between the coefficients x_i^ℓ and $x_{i'}^\ell$ for two related jobs $j_i, j_{i'}$. Formally, we define it as $x_{i,i'}^\ell := \max\{x_i^\ell - x_{i'}^\ell, x_{i'}^\ell - x_i^\ell\}$. This yields the following linear program, which minimizes the overall slack:

$$\begin{aligned}
& \text{minimize} && \sum x_{i,i'}^\ell \\
& \text{s.t.} && \\
& && x_{i,i'}^\ell \geq x_i^\ell - x_{i'}^\ell \quad \forall \ell \\
& && \quad \quad \quad \forall j_i, j_{i'} \in J \text{ that are related} \\
& && x_{i,i'}^\ell \geq x_{i'}^\ell - x_i^\ell \quad \forall \ell \\
& && \quad \quad \quad \forall j_i, j_{i'} \in J \text{ that are related} \\
& && \sum_{t_i^k \in T_i} \sum_{\ell=1}^7 x_i^\ell f_\ell(t_i^k) \leq 1 + \varepsilon \quad \forall j_i \in J \\
& && \sum_{t_i^k \in T_i} \sum_{\ell=1}^7 x_i^\ell f_\ell(t_i^k) \geq 1 - \varepsilon \quad \forall j_i \in J \\
& && x_i^1 \geq 0 \quad \forall j_i \in J \\
& && x_i^\ell \geq x_i^{\ell-1} \quad \forall j_i \in J \quad \ell \in \{2, \dots, 7\}
\end{aligned}$$

We set ε to 0.01. The resulting LP has 169,372 variables in its objective function. Since there are 735 jobs,² this means that a job is related to approximately 32.9

²We consider slightly more jobs than Frey and Osborne, since we use the finest granularity available from O*NET.

other jobs on average. The value of the objective function is 24.6, i.e., for two related jobs the coefficients differ only by 0.000145 on average. For comparison, the average value of a coefficient is 0.060. Our complete results can be found online at <http://jobs-study.ethz.ch>.

5 From Jobs to Tasks

Knowing the shares of the tasks enables us to set up a linear program that calculates for each task the probability to be automated. We want that the weighted average of the automation probabilities $p(t_i^k)$ of the tasks of a job j_i can explain the automation probability $p(j_i)$ of the job, i.e., $\sum_{t_i^k \in T_i} p(t_i^k) \cdot s(t_i^k) \approx p(j_i)$. Furthermore, we want to assign related tasks similar automation probabilities. To do this, we define a variable $t_{i,i'}^{k,k'}$ for each pair of related tasks t_i^k and $t_{i'}^{k'}$. It denotes the probability difference that we assign to the two tasks. Formally, it is defined as $t_{i,i'}^{k,k'} := \max\{p(t_i^k) - p(t_{i'}^{k'}), p(t_{i'}^{k'}) - p(t_i^k)\}$. We want to minimize the sum of these variables, i.e., the sum of the probability difference of all related tasks.

Combining these requirements with necessary conditions to have meaningful probabilities, i.e., $0 \leq p(t_i^k) \leq 1$, yields the following linear program:

$$\begin{aligned}
& \text{minimize} && \sum t_{i,i'}^{k,k'} \\
& \text{s.t.} && \\
& p(t_i^k) - p(t_{i'}^{k'}) \leq && t_{i,i'}^{k,k'} \quad \forall t_i^k, t_{i'}^{k'} \text{ that are related} \\
& p(t_{i'}^{k'}) - p(t_i^k) \leq && t_{i,i'}^{k,k'} \quad \forall t_i^k, t_{i'}^{k'} \text{ that are related} \\
& \sum_k p(t_i^k) \cdot s(t_i^k) \leq && p(j_i) (1 + \varepsilon) \quad \forall j_i \in J \\
& \sum_k p(t_i^k) \cdot s(t_i^k) \geq && p(j_i) (1 - \varepsilon) \quad \forall j_i \in J \\
& p(t_i^k) \geq && 0 \quad \forall j_i \in J, t_i^k \in T_i \\
& p(t_i^k) \leq && 1 \quad \forall j_i \in J, t_i^k \in T_i \\
& t_{i,i'}^{k,k'} \geq && 0 \quad \forall t_i^k, t_{i'}^{k'} \\
& t_{i,i'}^{k,k'} \leq && 1 \quad \forall t_i^k, t_{i'}^{k'}
\end{aligned}$$

We set ε to 0.01.

6 Linear Program Results

We now analyze the results of the linear program as described above. Later on, we will look at a small refinement to automatically detect outliers in our results.

6.1 Task Probabilities

The linear program as described in Section 5 has 105,748 variables in its objective function and it has a minimal value of 9,846. This means that two related tasks differ, on average, with regard to their probability by 9.3%. The complete results can be found online at <http://jobs-study.ethz.ch>. The histogram of the probability difference between two related tasks is shown in Figure 3. A majority of related tasks is assigned a similar probability. A small fraction of related tasks is assigned diametrically opposed probabilities, which seems startling. It can be reconciled by considering that neither the classification by Frey and Osborne nor the classification of tasks being related by O*NET are perfect.




Figure 3: A histogram of the probability with which two related tasks differ.

One example that highlights this are the two jobs *computer programmer* and *software developers, applications*. These two jobs have many related tasks, but the probabilities of these jobs differ a lot (4% for *software developers, applications*, 48% for *computer programmers*). Hence, the diametrically opposed probabilities are necessary to meet the constraints of the linear program.

In the following, we present a few selected jobs to illustrate our results. The first example is *chemists*. This job has an automation probability of 10% according to Frey and Osborne. Only one task has, according to our linear program, a high probability of being automated: “Induce changes in composition of substances by introducing heat, light, energy, or chemical catalysts for quantitative or qualitative

Task Description	p	Share
Write decisions on cases.	1	5.1
Instruct juries on applicable laws, direct juries to deduce the facts from the evidence presented, and hear their verdicts.	1	3.4
Monitor proceedings to ensure that all applicable rules and procedures are followed.	1	8.0
Advise attorneys, juries, litigants, and court personnel regarding conduct, issues, and proceedings.	1	6.2
Interpret and enforce rules of procedure or establish new rules in situations where there are no procedures already established by law.	1	5.4
Conduct preliminary hearings to decide issues such as whether there is reasonable and probable cause to hold defendants in felony cases.	1	3.9
Rule on admissibility of evidence and methods of conducting testimony.	0.94	5.3
Preside over hearings and listen to allegations made by plaintiffs to determine whether the evidence supports the charges.	0.46	5.9
Perform wedding ceremonies.	0.39	2.7
Read documents on pleadings and motions to ascertain facts and issues.	0	10.1
Research legal issues and write opinions on the issues.	0	6.5
Settle disputes between opposing attorneys.	0	4.6
Participate in judicial tribunals to help resolve disputes.	0	6.6
Rule on custody and access disputes, and enforce court orders regarding custody and support of children.	0	6.3
Sentence defendants in criminal cases, on conviction by jury, according to applicable government statutes.	0	4.0
Grant divorces and divide assets between spouses.	0	4.7
Award compensation for damages to litigants in civil cases in relation to findings by juries or by the court.	0	3.8
Supervise other judges, court officers, and the court's administrative staff.	0	8.5

Table 4: The automation probability and the share of each task of *Judges, Magistrate Judges, and Magistrates*". The automation probability of this job is 40%.

analysis." Other simple mechanical tasks have been assigned low automation probabilities. We will revisit this job in Section 6.2.

Next up: *judges*. Their automation probability is 40%. The tasks, their probabilities, and their shares are shown in Table 4. The tasks that can be automated can be grouped in two sets: preliminary hearings which includes making first assessments, and ensuring that the procedures in court are followed. The tasks that involve sentencing (or the preparation thereof) have been assigned low automation probabilities.

Figure 5 shows the histogram of the probabilities from our linear program. The probabilities for most tasks are either very high or very low and only a few tasks have a probability in-between. This desired side effect of our linear program helps us to achieve our goal of allowing job experts (and laymen) to argue about the validity of our results. We invite the reader to have a look at other jobs at <http://jobs-study.ethz.ch>.

6.2 Outlier Detection

To evaluate our approach and check for outliers, we use a variant of cross-validation. For every job j_i , we create a linear program without job j_i . This yields a probability




Figure 5: A histogram of the probabilities that tasks can be automated. Nearly all tasks are assigned either 0 or 1. This simplifies arguing whether our classification is correct.

$p_i(\cdot)$ for every task but the tasks from j_i . Afterward, we calculate the new probability $p'(j_i)$ that job j_i can be automated. We do this by setting the probability $p'(t_i^k)$ of each task t_i^k to the average of all tasks that are related to it. We denote the set of related tasks by $N(t_i^k)$. Formally, we set $p'(t_i^k) := \frac{1}{|N(t_i^k)|} \sum_{t_i^{k'} \in N(t_i^k)} p_i(t_i^{k'})$.




Figure 6: A histogram of $p(t_i^k) - p'(t_i^k)$ for all tasks. This plot shows that most tasks have a similar probability in both approaches. The values are sharply distributed around 0.

We first compare $p'(t_i^k)$ with $p(t_i^k)$. The difference between these two probabilities should be small for the majority of the tasks. This is indeed what can be seen in Figure 6. The histogram of $p'(t_i^k) - p(t_i^k)$ shows that nearly all tasks have similar probabilities in both approaches. The average absolute difference is less than 20%. The distribution is centered around 0. Its mean is less than 0.05%.

By combining the new probability of each task with its share, we can calculate

the new probability of job j_i by using a weighted average. This allows us to compare $p'(j_i)$ with $p(j_i)$.

We have plotted this difference, i.e., $p(j_i) - p'(j_i)$, in Figure 7. We can see that the difference is centered around 0%; with the average absolute difference being less than 20%. For more than half of the jobs our probability differs by less than 20% from Frey and Osborne [12]. Most interesting are the jobs whose probability differs significantly. We now have a look at a few of them.

There are jobs where our probability is more than 80% smaller than the one by Frey and Osborne. One job is *compensation and benefits managers*. We assigned it a probability $p'(j)$ to be automated of 9.1%; compared to $p(j) = 96%$ by Frey and Osborne. We do not claim to know the true value, but we can look at the job and compare it to the probabilities of jobs we consider similar. Notice that this is conceptually similar to what our linear program does and thus might be biased. The tasks of this job are shown in Table 8. If we manually compare them to related tasks, we conclude that they do not seem to be automatable in the next few decades. We do favor our result over the result of Frey and Osborne.

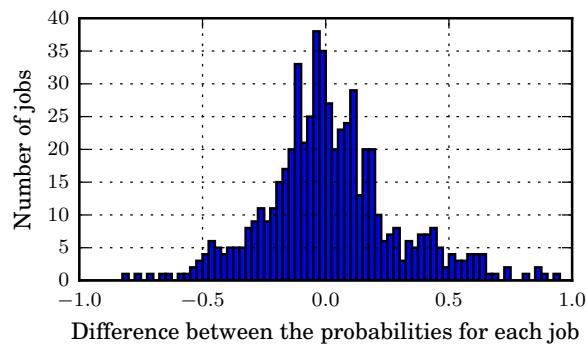


Figure 7: A histogram of the difference between the probability by Frey and Osborne with our probability. The distribution is centered around 0 and a majority of the jobs differs by less than 20%.

There is only one job that we assign a much higher probability than Frey and Osborne. The job *First-Line supervisors of production and operating workers* has been assigned a 83% automation probability by us and only 1.6% by Frey and Osborne. A close inspection of the tasks makes us believe that the true value is between these extremes. Quite a few of the tasks are clearly automatable, e.g., “Keep records of employees’ attendance and hours worked.” and “Observe work and monitor gauges, dials, and other indicators to ensure that operators conform to production or processing standards.” Others, e.g., “Read and analyze charts, work orders, production schedules, and other records and reports to determine

Task Description	<i>p</i>	<i>p'</i>
Advise management on such matters as equal employment opportunity, sexual harassment and discrimination.	1	0.15
Study legislation, arbitration decisions, and collective bargaining contracts to assess industry trends.	1	0
Fulfill all reporting requirements of all relevant government rules and regulations, including the Employee Retirement Income Security Act (ERISA).	1	0.20
Investigate and report on industrial accidents for insurance carriers.	1	0.12
Represent organization at personnel-related hearings and investigations.	1	0
Analyze compensation policies, government regulations, and prevailing wage rates to develop competitive compensation plan.	1	0.5
Mediate between benefits providers and employees, such as by assisting in handling employees' benefits-related questions or taking suggestions.	1	0.42
Prepare detailed job descriptions and classification systems and define job levels and families, in partnership with other managers.	1	0
Prepare personnel forecasts to project employment needs.	1	0
Direct preparation and distribution of written and verbal information to inform employees of benefits, compensation, and personnel policies.	1	0
Manage the design and development of tools to assist employees in benefits selection, and to guide managers through compensation decisions.	1	0
Design, evaluate and modify benefits policies to ensure that programs are current, competitive and in compliance with legal requirements.	1	0
Administer, direct, and review employee benefit programs, including the integration of benefit programs following mergers and acquisitions.	1	0
Prepare budgets for personnel operations.	1	0.03
Maintain records and compile statistical reports concerning personnel-related data such as hires, transfers, performance appraisals, and absenteeism rates.	1	0
Contract with vendors to provide employee services, such as food services, transportation, or relocation service.	1	0.38
Identify and implement benefits to increase the quality of life for employees, by working with brokers and researching benefits issues.	1	0
Plan, direct, supervise, and coordinate work activities of subordinates and staff relating to employment, compensation, labor relations, and employee relations.	1	0
Negotiate bargaining agreements.	1	0.67
Plan and conduct new employee orientations to foster positive attitude toward organizational objectives.	1	0
Conduct exit interviews to identify reasons for employee termination.	1	0
Develop methods to improve employment policies, processes, and practices, and recommend changes to management.	0.51	0
Formulate policies, procedures and programs for recruitment, testing, placement, classification, orientation, benefits and compensation, and labor and industrial relations.	0.23	0.01

Table 8: The tasks and their corresponding probability that they will be automated for *Compensation and Benefits Managers* according to our original linear program and the cross-validation.

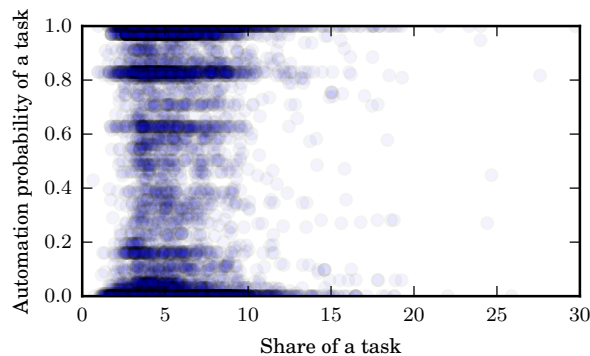


Figure 9: Our results show almost no correlation between the share of a task and its probability that it will be automated.

production requirements and to evaluate current production estimates and outputs.” seem difficult to automate. The complete results for this can job be found at <http://jobs-study.ethz.ch>.

We continue by comparing the previous results with the approach described in this section. To do this, we return to the jobs that we have looked at previously. First off is the job *chemists*. The automation probability of most tasks has increased. Consequently, the automation probability of this job has increased from 10% to 42%. Due to the large difference, this job should be analyzed in-depth by job experts.

The changes in the automation probability of the tasks of *judges* are much smaller. Most tasks have a similar automation probability as before and the overall probability of this job has changed marginally, i.e., increased only from 40% to 50%. Therefore, we are confident that the classification by Frey and Osborne is correct.

We conclude that our approach can also be used to detect outliers in the results of Frey and Osborne. We can then manually inspect the automation probabilities of the tasks of such an outlier to determine the truth. We think our results allow us to fine tune the results from Frey and Osborne, but not replace it, as we need their results to bootstrap our linear program.

7 Further Analysis

In addition to inspecting every task of every job, we consider a broader picture. We do this by looking at general properties of a job that correlate with the probability that it can be automated.

7.1 Tasks




Figure 10: The probability that a job can be automated over the level of the abilities “deductive reasoning” and “originality” used in this job. These levels are defined by O*NET and for each they provide an anchor point. Level 2 of “deductive reasoning” means “knowing that a stalled car can coast downhill” and level 5 “deciding what factors to consider in selecting stocks”. Level 2 of “originality” means “using a credit card to open a locked door” and level 6 “inventing a new type of man-made fiber”. Every point represents one job. A higher level of either of these two abilities correlates, as expected, with a lower probability to be automated.

We first analyze the share of a task. The higher the share, the more often a task is performed. Hence, from a machine learning perspective this means that much more training data is available. This might lead to the conclusion that such a task is easier to automate. To disprove this claim, we plotted the share of a task over the probability that a task can be automated according to our linear program. The resulting graph is shown in Figure 9. Every dot represents one task, with its share on the x -axis and its probability on the y -axis. We see that there is barely any correlation between these two. We conclude that tasks that are done more frequently are not more likely to be automated.

7.2 Jobs

We continue our analysis by looking at the correlation between the properties that a job has, e.g., what kind of degree is necessary to do a job, and the probability that this job can be automated. Correlation does not imply causation, but nevertheless, these results reveal some interesting nuggets.

O*NET provides the level that the ability “deductive reasoning” is used in a job. The level ranges from 1 to 7, where for example level 2 means “knowing that a stalled car can coast downhill” and level 5 “deciding what factors to consider

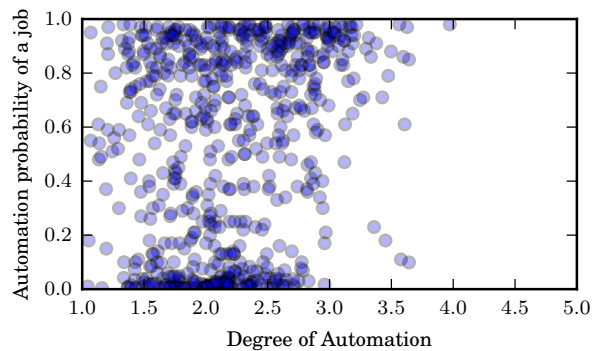


Figure 11: The probability that a job can be automated over the level of the current “Degree of Automation”. This level ranges from 1 (Not at all automated) to 5 (Completely automated). The rather small correlation of 0.23 implies that different jobs will soon be affected.

in selecting stocks”. For every job, we have one value between 1 and 7. The resulting graph can be seen in Figure 10. Every job is represented by one dot; its x -coordinate being its level and the y -coordinate its probability. We can see that jobs that require a high level of deductive reasoning tend to have a lower probability of being automated. A similar result can be seen for “originality” (see Figure 10).




Figure 12: The probability that a job can be automated over the amount of “On the Job Training”, which ranges from 1 (none or short demonstration) to 9 (over 10 years) and the “Required of Level of Education”, which ranges from 1 (less than a high school diploma) to 12 (post-doctoral training).

Level 2 of “originality” means “using a credit card to open a locked door” and level 6 means “inventing a new type of man-made fiber”. This confirms our expectation that these abilities will remain difficult for a computer.

O*NET even has an explicit value for the current level of the “degree of

automation” for each job. This level ranges from 1 (not at all automated) to 5 (completely automated). As depicted in Figure 11, the already existing level of automation barely correlates with the probability that this job will be automated. This indicates that not only jobs that are already affected by automation are in danger, but also a whole new set of jobs. This is aligned with the recent worries about many new jobs soon being affected by computerization.

We conclude this section by looking at the effect that the level of required education for a job has on the probability to be automated. Jobs that require only very little education (level 1, i.e., less than a high school diploma) tend to have a higher probability than jobs that require an associate degree (level 5) which in turn have a higher probability than jobs that require post-doctoral training (level 12). Most jobs that require little education are in danger. It is noteworthy that the effect of training before the job is much stronger than the effect of on the job training. Jobs that require more on the job training only have a marginally smaller probability to be automated. Both plots are shown in Figure 12.

8 Existing Trends

In this section we analyze to what extent automation is already happening. To investigate this, we used historical job data from 2001 until 2015. For every year we know for every job in the O*NET database how many people were employed. Instead of using the absolute values (which tend to increase for most jobs since more and more people live and work in the US), we consider the relative values and look at the percentage of how many people are employed in this job. We consider the fraction a job j_i has in 2001 and focus on the factor c_i by which this fraction has changed in 2015. Note that due to the financial crisis, where a lot of jobs have been lost, the data shows to some extent how recession-proof a job is.

We plot the factor c_i against the automation probability of this job for every job j_i . Every job is represented by a dot; its x -coordinate being the factor c_i and the y -coordinate being its automation probability. We expect that jobs with a high automation probability tend to have a factor smaller than one, i.e., its fraction of the labor market has decreased and vice versa. The result can be seen in Figure 13. Even though there are quite a few outliers, the overall trend follows our expectation. The most extreme outlier is the job *service unit operators, oil, gas, and mining* whose fraction has increased by more than a factor of 4, but is deemed 93% automatable. This can easily be explained due to the rise of fracking. A job that is considered non-automatable, but whose fraction has been reduced to about a third of its initial value, is *advertising and promotions managers*. We suspect that this job is highly susceptible to recessions, i.e., the financial crisis in 2008/2009.

The above analysis shows that the demand for jobs that are highly automatable

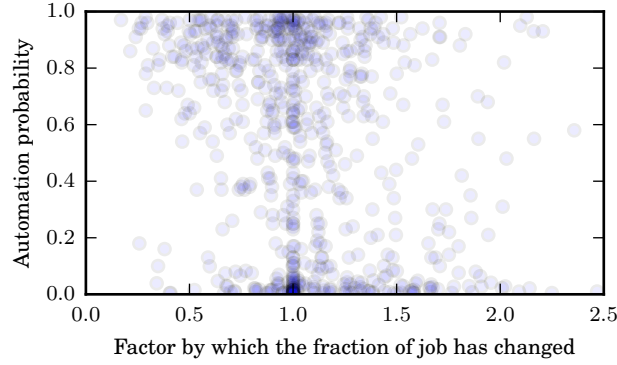


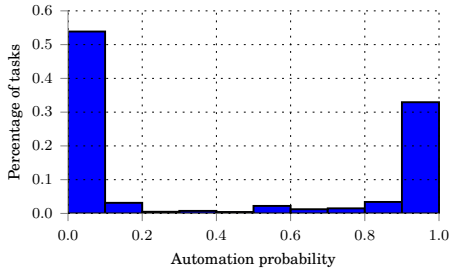
Figure 13: The factor c_i by which the fraction of a job has changed over the automation probability. We can see that jobs that have a factor smaller than one tend to have a high automation probability.

has already started to decrease. We once again look at the problem from a task level perspective. We assume that the demand for a job decreases because the demand for highly automatable tasks decreases since some of these tasks are already being automated.

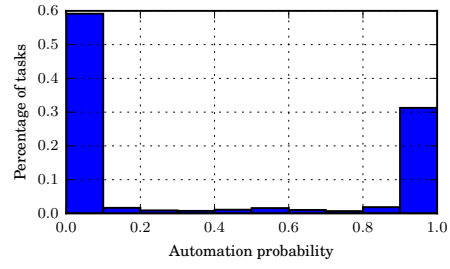
Our analysis works as follows: We initially set the share $s(t_i^k)$ of a task t_i^k as described in Section 4. We allow the share $s(t_i^k)$ of a task t_i^k to change by a factor which we denote by c_i^k . The weighted average of the adjusted share should be equal to c_i , i.e., $\sum_{k \in T_i} c_i^k \cdot s(t_i^k) \approx c_i$ to reflect the change in demand for a job. We assume that if one task can already be automated, i.e., has a factor smaller than one, then related tasks also tend to decrease and vice versa. We want to minimize the overall difference $s_{i,i'}^{k,k'} := |c_i^k - c_{i'}^{k'}|$ between two related tasks $t_i^k, t_{i'}^{k'}$.

This is modeled in the following linear program:

$$\begin{aligned}
& \text{minimize} && \sum s_{i,i'}^{k,k'} \\
& \text{s.t.} && \\
& c_i^k - c_{i'}^{k'} \leq && s_{i,i'}^{k,k'} \quad \forall t_i^k, t_{i'}^{k'} \text{ that are related} \\
& c_{i'}^{k'} - c_i^k \leq && s_{i,i'}^{k,k'} \quad \forall t_i^k, t_{i'}^{k'} \text{ that are related} \\
& \sum_k c_i^k \cdot s(t_i^k) \leq && c_i (1 + \varepsilon) \quad \forall j_i \in J \\
& \sum_k c_i^k \cdot s(t_i^k) \geq && c_i (1 - \varepsilon) \quad \forall j_i \in J \\
& c_i^k \geq && 0 \quad \forall j_i \in J, t_i^k \in T_i
\end{aligned}$$



(a) Tasks whose share decreases by at least a factor of 0.5



(b) Tasks whose share increases by at least a factor of 2

Figure 14: The histogram of the automation probability of all tasks whose share either decreases by at least factor of 0.5 or increases by at least a factor of 2.

To verify our claim, we first look at tasks whose factor c_i^k is less than 0.5, i.e., whose share has more than halved. We expect that these tasks to have a high automation probability. The histogram is depicted in Figure 14. Similarly, tasks whose share has increased by more than a factor of 2 should have a low automation probability. The results are also shown in Figure 14. Despite various other economic factors that come into play like economic cycles or major political events, we can see that tasks whose share has more than doubled tend to have a smaller automation probability.

9 Conclusion

We believe that automation is one of the main challenges for society. In our opinion, the seminal work of Frey and Osborne did an excellent job of getting the discussion going. In this paper we dug a bit deeper, by looking not only at jobs – but at the tasks that make up a job. We hope that opening the Frey/Osborne black box will contribute to the discussion. The professionals that are actually doing a job are the main experts to decide what parts of their job can or cannot be computerized. The Frey/Osborne work only tells these experts that their job is 87% automatable, but what does it actually mean? With our work, job experts can look inside the box, and understand which tasks of their job are at risk. Our hope is that the job experts have a discussion which results are believable and which are not, and why. To facilitate this discussion, we have created a web page (<http://jobs-study.ethz.ch>) that allows users to comment upon our results.

References

- [1] Michael Auli, Michel Galley, Chris Quirk, and Geoffrey Zweig. Joint language and translation modeling with recurrent neural networks. In *EMNLP*, volume 3, page 0, 2013.
- [2] David H. Autor and David Dorn. The growth of low skill service jobs and the polarization of the us labor market. Technical report, National Bureau of Economic Research, 2009.
- [3] David H. Autor, Lawrence F Katz, and Alan B Krueger. Computing inequality: have computers changed the labor market? Technical report, National Bureau of Economic Research, 1997.
- [4] David H Autor, Frank Levy, and Richard J Murnane. The skill content of recent technological change: An empirical exploration. Technical report, National Bureau of Economic Research, 2001.
- [5] Holger Bonin, Terry Gregory, and Ulrich Zierahn. Übertragung der Studie von Frey/Osborne (2013) auf Deutschland. Technical report, ZEW Kurzepertise, 2015.
- [6] John Bound and George E Johnson. Changes in the structure of wages during the 1980's: An evaluation of alternative explanations. Technical report, National Bureau of Economic Research, 1989.
- [7] Erik Brynjolfsson and Andrew McAfee. *Race against the machine: How the digital revolution is accelerating innovation, driving productivity, and irreversibly transforming employment and the economy*. Brynjolfsson and McAfee, 2012.
- [8] Erik Brynjolfsson and Andrew McAfee. *The second machine age: work, progress, and prosperity in a time of brilliant technologies*. WW Norton & Company, 2014.

- [9] Kerwin Kofi Charles, Erik Hurst, and Matthew J Notowidigdo. Manufacturing decline, housing booms, and non-employment. Technical report, National Bureau of Economic Research, 2013.
- [10] L Deng, D Yu, and G Hinton. Deep learning for speech recognition and related applications. In *NIPS Workshop*, 2009.
- [11] Martin Ford. *Rise of the Robots: Technology and the Threat of a Jobless Future*. Basic Books, 2015.
- [12] Carl Benedikt Frey and Michael A. Osborne. The future of employment: How susceptible are jobs to computerisation?, 2013.
- [13] Maarten Goos and Alan Manning. Lousy and lovely jobs: The rising polarization of work in britain. *The review of economics and statistics*, 89(1):118–133, 2007.
- [14] Erico Guizzo. How google’s self-driving car works. *IEEE Spectrum Online*, October, 18, 2011.
- [15] John Maynard Keynes. Economic possibilities for our grandchildren (1930). *Essays in persuasion*, pages 358–73, 1933.
- [16] Quoc V Le. Building high-level features using large scale unsupervised learning. In *ICASSP*, pages 8595–8598. IEEE, 2013.
- [17] James Manyika, Michael Chui, Jacques Bughin, Richard Dobbs, Peter Bisson, and Alex Marrs. *Disruptive technologies: Advances that will transform life, business, and the global economy*, volume 12. McKinsey Global Institute New York, 2013.
- [18] Pajarinen Mika and Anders Ekeland. Computerization and the future of jobs in norway. 2015.
- [19] Marvin L Minsky. *Computation: finite and infinite machines*. Prentice-Hall, Inc., 1967.
- [20] David Silver, Aja Huang, Christopher J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of go with deep neural networks and tree search. *Nature*, 529:484–503, 2016.
- [21] Alan M Turing. Computing machinery and intelligence. *Mind*, 59(236):433–460, 1950.
- [22] Moshe Y. Vardi. The great robotics debate. *Commun. ACM*, 56(7), July 2013.
- [23] Moshe Y Vardi. Is information technology destroying the middle class? *Commun. ACM*, 58(2):5, 2015.
- [24] Moshe Y. Vardi. The moral imperative of artificial intelligence. *Commun. ACM*, 59(5), April 2016.
- [25] Adrian Wood. North-south trade, employment and inequality. *New York*, 1994.

- [26] Adrian Wood. Globalisation and the rise in labour market inequalities. *The economic journal*, 108(450):1463–1482, 1998.