Skills Builder CAREERS EXPLORER

Skills Builder PARTNERSHIP

GATSBY

Introduction

This paper details our efforts to estimate the essential skills required for all 1,137 jobs in the extended Standard Occupational Classification (SOC) using programmatic methods.

The project began in October 2022 when the Gatsby Foundation kindly agreed to fund a feasibility study into the possibility of generating essential skill scores for jobs using data science. In this feasibility study we explored a number of different approaches, and detailed the advantages and disadvantages of each. By combining different approaches, we concluded it would be possible to produce accurate estimates for all the jobs in the extended SOC framework. With continued support from the Gatsby Foundation, we proceeded to build out the process documented here.

Context

What are essential skills?

Essential skills are those highly transferable skills like problem solving, teamwork and leadership that you need for almost any job. These eight skills are codified and broken down into measurable, teachable components in the Skills Builder Universal Framework.

Essential skills are those highly transferable skills that everyone needs to do almost any job, which make specific knowledge and technical skills fully productive. These are therefore distinct from basic skills (literacy, numeracy and digital skills) and technical skills (specific to a particular sector or role, sometimes drawing off a particular body of knowledge). In the research literature, they are often referred to as “transversal” or “higher order cognitive” skills.

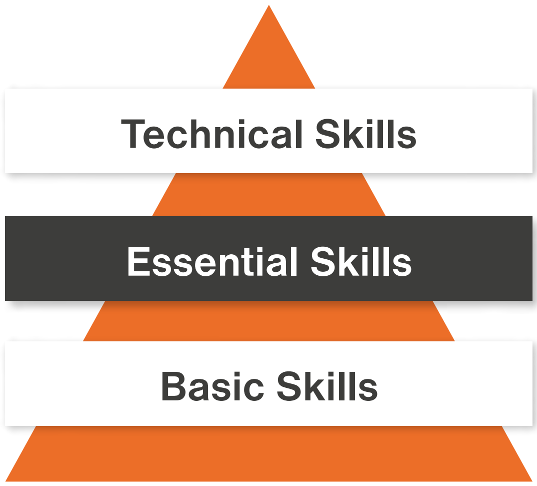

The eight Essential Skills are:

1. LISTENING
receiving, retaining and processing of information or ideas

2. SPEAKING
oral communication of information and ideas

3. PROBLEM SOLVING
the ability to find a solution to a situation or challenge

4. CREATIVITY
use of imagination and generation of new ideas

5. STAYING POSITIVE
ability to use tactics and strategies to overcome setbacks and achieve goals

6. AIMING HIGH
ability to set clear, tangible goals and devise a robust route to achieving them

7. LEADERSHIP
supporting, encouraging and developing others to achieve a shared goal

8. TEAMWORK
working cooperatively with others towards achieving a shared goal

What is SOC?

SOC stands for Standard Occupational Classification. The term is used broadly to describe many different attempts to classify jobs into a taxonomy, usually in several layers of increasing specificity. Government bodies, such as the Bureau of Labour Statistics (BLS) in the US or Office of National Statistics (ONS) in the UK release SOCs for public use, and produce statistics including SOC – for example, wage quantiles per SOC code per region.

As the labour market composition changes over time, so too do SOCs get updated and re-released in new versions. In the UK, the most recent iterations of SOC include SOC2000, SOC2010, and SOC2020. While SOC2020 is intended to best reflect the current state of the labour market, many organisations still use SOC2010. A crosswalk of changes details how to translate between versions.

Most recently, an extended version of SOC2020 has been released, after work led by Professor Peter Elias from the University of Warwick's Institute for Employment Research. This further breaks down the 412 unit groups in SOC2020 into 1,367 jobs (or "sub-unit groups"), providing unprecedented granularity with which to describe the labour market.

	SOC2020				SOC2020 Extended
	Major group	**Sub-major group**	**Minor group**	**Unit group**	**Sub-unit group**
Code	5	51	511	5111	5111/05
Title	Skilled Trades Occupations	Skilled Agricultural and Related Trades	Agricultural and Related Trades	Farmers	Poultry farmers

Figure 1: illustration of the structure of SOC2020 with extension

Essential skills and employment

Over the last four years, the support of the Gatsby Foundation has enabled us to create and scale the Skills Builder Universal Framework of Essential Skills. During that time, we created the Universal Framework, building off an analysis of the existing Skills Builder Framework that had been used in education. We looked at the Framework through four lenses to check its completeness and relevance: other international frameworks and approaches to employability; Burning Glass job advertisement data; universities' graduate attribute statements; and apprenticeship standards. The resulting Framework was launched in May 2020.

There was also huge potential for adoption of the Framework among employers. Over the last three years we ran two 'Trailblazer' cohorts of employers, who we worked with closely to embed the Skills Builder approach into their outreach, recruitment, and staff development work. This approach has helped to build the evidence base for the impact the Universal Framework can have for businesses in supporting their recruitment pipelines and processes and increasing staff engagement and progression. The Skills Builder Employer Programme is already being used by fifty employers who are looking to achieve best practice in essential skills across their recruitment, staff development and outreach.

In parallel, there has been strong progress more widely in the uptake of the Framework. There are now more than 800 organisations who are actively using and championing the Universal Framework as part of the Skills Builder Partnership. These include large national and international institutions like the British Council, the Careers & Enterprise Company and National Citizens Service. Between our partners, we now have a touchpoint with 87% of secondary schools and colleges in the UK. We have also made some progress at a policy level: the Universal Framework is recommended in statutory guidance for secondary schools and colleges; IFATE have made it a recommended model for developing apprenticeship standards; and we have promising ongoing conversations with DWP and DfE.

However, there remains more to do if we are to solidify the Universal Framework's position as the de facto model of essential skills development across the country. One of the routes to supporting widespread uptake of the Universal Framework is to provide tools and insights that can enable its usage. In this project, we aimed to do exactly that by producing a dataset that, for the first time, reveals the essential skills required to do almost any job in today's economy.

Objectives:

- To further enable engagement with, and adoption of, the Universal Framework amongst businesses, other organisations, researchers, and members of the public.
- To enable analysis of Essential Skill requirements in the context of any research incorporating a Standard Occupational Classification – for instance, into apprenticeships, higher education, or the labour market.
- To create opportunities for valuable products and services in the wider ecosystem that incorporate essential skills alongside jobs data.

Process

Identifying a feasible process

A feasibility study assessed the likely success of efforts to estimate skill scores for jobs in the extended SOC framework using different approaches and data sources. We concluded that the best available source data comes from the US jobs and skills database O*NET, which has been compiled from years' of primary research. For us to use this data, however, a map between UK SOC Extended ("SOC Ex") and O*NET SOC is required.

We identified a number of potential approaches, including directly mapping essential skills to O*NET indicator scales and training machine learning models on data collected from experts. We measured success objectively against predetermined criteria. Prototypes of 4 different approaches suggested that each was likely to be successful to varying degrees, and identified the advantages and limitations of each.

The O*NET database comprises a number of scales arranged into categories, such as skills, abilities, knowledge, and work contexts. These are divided into further categories. For example: of the 74 skills, 10 are considered "Basic Skills", and 30 are considered "Cross-Functional Skills". These categorisations, however, do not necessarily align with our equivalent definitions. Each job is given at least one score in each of these measures, based on industry research. For example, in the basic skill 'Active Listening', Psychiatrists score 91 for Importance and 71 for Level, whereas Sewers score 41 and 34.

On the O*NET website measures are presented on a 0-100 scale. In the raw data, however, the scales range from 1 to 7. For guidance, O*NET provides scale level anchor descriptions. These are presented out of 100 on the website, and out of 7 in the raw data. Some measures are scored twice (ie both Importance and Level). Where a measure is scored twice, high collinearity is observed. Scale anchor descriptions are given for Level but not for Importance.

O*NET data has been used in many similar projects. In 2022, Australia released a new national Skills Classification for ANZSCO SOC codes based on O*NET data and existing Australian skills frameworks. To calculate required skill scores in 10 pre-existing core competencies (not dissimilar to essential skills), a small number of closely matching O*NET measures were selected from O*NET Skills, Work Styles, and Work Activities using natural language processing and human judgement. Scale conversion was achieved by placing the 3 O*NET level scale anchors on a 10-point scale derived from existing Australian skills frameworks. Using a map of ANZSCO to O*NET SOC, a 10-point Writing skill can then be calculated for every Australian SOC code. Previous projects have tended towards simplicity by using the fewest number of O*NET scales possible. However, we note that essential skills are less similar to O*NET indicators than the skills targeted in the ANZSCO project.

We used this previous work as a starting point in our efforts to build an analogous process for essential skills, but through testing and iteration we found that more complex methods were necessary to achieve satisfactory results for essential skills.

Overview

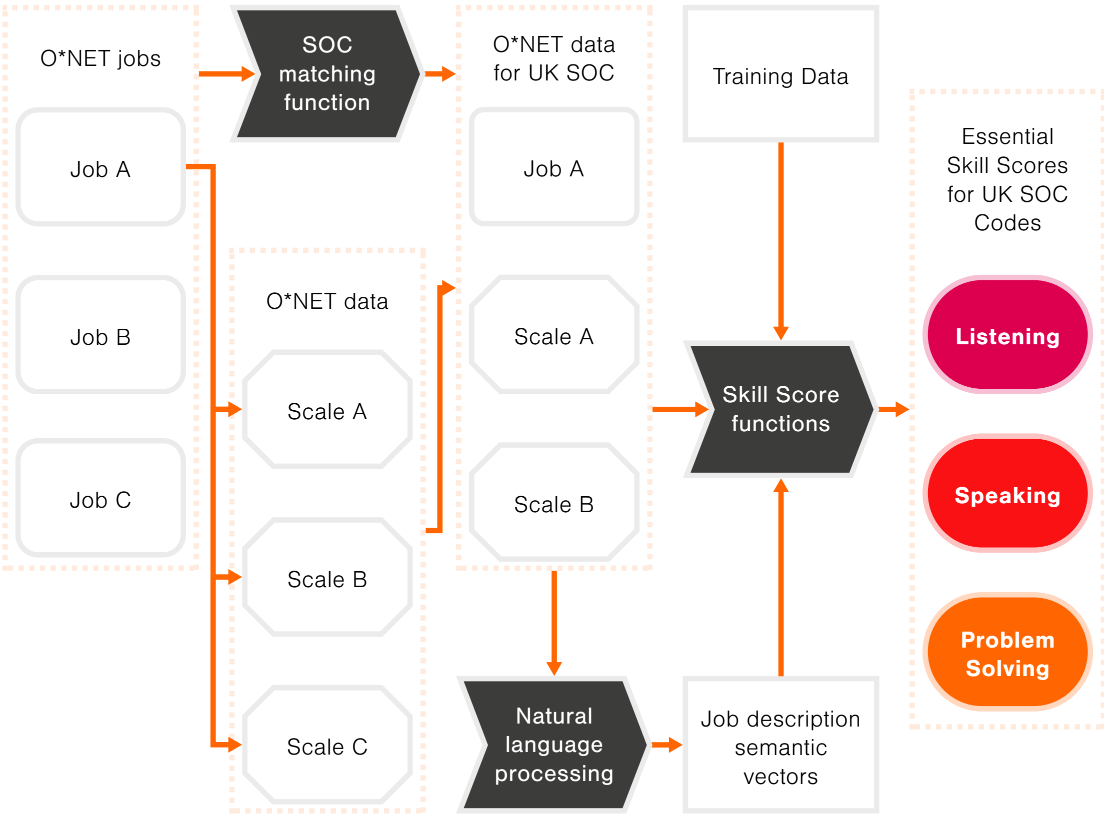

Figure 1: process schematic

1. Matching UK SOC jobs to one or more O*NET SOC jobs

In order to use O*NET data in our modelling, it was necessary to match each job in SOC Ex to at least one job in O*NET SOC. At the time of the project, with SOC Ex newly released, no crosswalk yet exists between these two classifications – we proceeded to build our own.

We used two different natural language models (BERT and GPT) to score the similarity of job descriptions in UK SOC Ex to those in O*NET SOC. For many jobs – for instance "Chief Executive" – this results in a single highly-scored match. For other jobs, however – e.g. "Mystery Shopper" or "Clairvoyant" – we find no close equivalents in the US classification. Matching is complicated by the fact that many jobs in the O*NET database contain incomplete data, and in many cases the best match had to be discarded for this reason.

Because job descriptions describe the tasks and skills that comprise a job, we found that matching jobs to those with descriptions that contained similar tasks and skills usually resulted in good outcomes, even when job titles did not match well. We designed an algorithm that matched each job in UK SOC Ex with one or more O*NET jobs depending on the relative strength of matches. This means that if there is one good match, we use only this match; if there are five roughly equal matches, we average the O*NET scales weighted by the strength of each.

The end result of this step was effectively to port the US data to UK jobs – in other words, each job in SOC Ex now had a value in each of the O*NET scales. There are several hundred of these, split into categories such as Skills, Abilities, Work Contexts, Interests, and so on.

2. Model Development

We asked several experts to score 80 jobs in each of the eight essential skills, on a scale of 0 to 15. This data was then used as training data to train two different models, using the data created above. We found that moderating modelling with semantic data improved stability and performance, and has the benefit of providing input that is not dependent on O*NET. Embedding vectors for each job description were included as independent variables – this allowed for relevant related terms in job descriptions such as “leads a team” to have a bearing on the results.

Elastic Net regression was used in order to ignore irrelevant variables and to mitigate the high collinearity between many input variables. The results were averaged with those from a neural network trained on the same input data. In general, these different models produced very similar results. The predictions from these two models were averaged and rescaled to produce a preliminary skill score.

3. Refinement and adjustment

Inspection of the raw output values produced by this process revealed the existence of some erroneous results, in almost all cases the result of inadequate matching between UK SOC Ex and O*NET SOC. To correct for these, we designed an algorithm that selectively adjusts scores with low certainty towards the sub-unit group average (itself weighted towards more certain results), since sub-unit groups (or 6-digit SOC) contain jobs with very similar skill requirements. The effect of this is analogous to smoothing a signal in order to reduce noise, with the advantage of adjusting only uncertain results towards certain ones, and thus preserving accuracy.

In a final round of adjustments, we also created a small number of custom groups of jobs for which skill scores should be coerced together with a given force. For example, we asserted that certain medical professions, or director roles, should have skill scores that are closer together. In a small number of cases, we also directly pegged one job’s skill scores to another’s when, for instance, the model fails to account for differing seniority between otherwise similar roles.

Limitations, Learnings, and Next Steps

Identifying a feasible process

This process was not without its challenges. In this section we detail some of the things we learned, and reflect on the future of this work.

O*NET's value as a data source for professional skills analysis is reiterated by the number of times it is used as a data source, as in the creation of the Australian National Skills Taxonomy. This also attests to its uniqueness – if such a data source existed in the UK, a major challenge and source of error for such enterprises would no longer exist. Because SOC2020 Ex is new, no crosswalk yet exists between it and O*NET. Our efforts to create one reveal the challenges inherent to this process.

This project leant heavily on natural language processing. With the capabilities and adoption of machine learning language models likely to grow significantly in the near future, the consistency and machine readability of semantic data should be a priority for those managing data sets. Job descriptions are rich sources of semantic information, but approaches that attempt to utilise this will be sensitive to inconsistencies. For example, we note that, before adjustment, surgeons received much lower skill scores than other medical specialists. Why? Likely because the job description says "they take on a range of administrative duties", when this is not mentioned for other similar jobs.

This process will be improved when additional data sources denominated in SOC2020 Ex become available. For this reason we hope to be contributing to a field that will grow rapidly. We hope that adoption of SOC2020 Ex will result in the production of high quality labour market data, job posting analysis, and skills data that can be used to improve this work in a subsequent iteration.

Finally, we note the conceptual limitations of the approach we have taken. There is some inevitable tension between a universal framework that is easily used to build skills in practice and a granular dataset optimised for computation. At Skills Builder, we view "skill score" (in a single skill, or an average of all eight skills) as the product of an assessment of 16 discrete steps per skill. We aim to build skills by building skill steps, the aggregate side effect of which is skill score. The linear relativity of skill score interpreted in this way cannot always be assumed: a score of 10 is "better" than a score of 9, but that score of 9 might comprise much higher steps in which mastery has been achieved.

Because we define skills with steps that are broadly sequential, and cumulative – to master step 1 it is likely you have mastered step 0 – we feel that coercing the rich set of abilities outlined in the Universal Framework into eight linear scales is a justifiable compromise. The end result of this, however, is that those interpreting this information can tell only roughly how much of a skill they need for a job, rather than which specific steps they should focus on building.

With the richer datasets imagined above, we hope that soon we would be able to replicate this work at the skill step level, and provide users with direction as to exactly which steps need to be built for each job. Instead of 8 outputs on a scale of 0-15, this work would produce 128 outputs each with 5 possible values.

With special thanks to:

GATSBY

Skills Builder
PARTNERSHIP

244-254 Cambridge Heath Rd
London E2 9DA
skillsbuilder.org