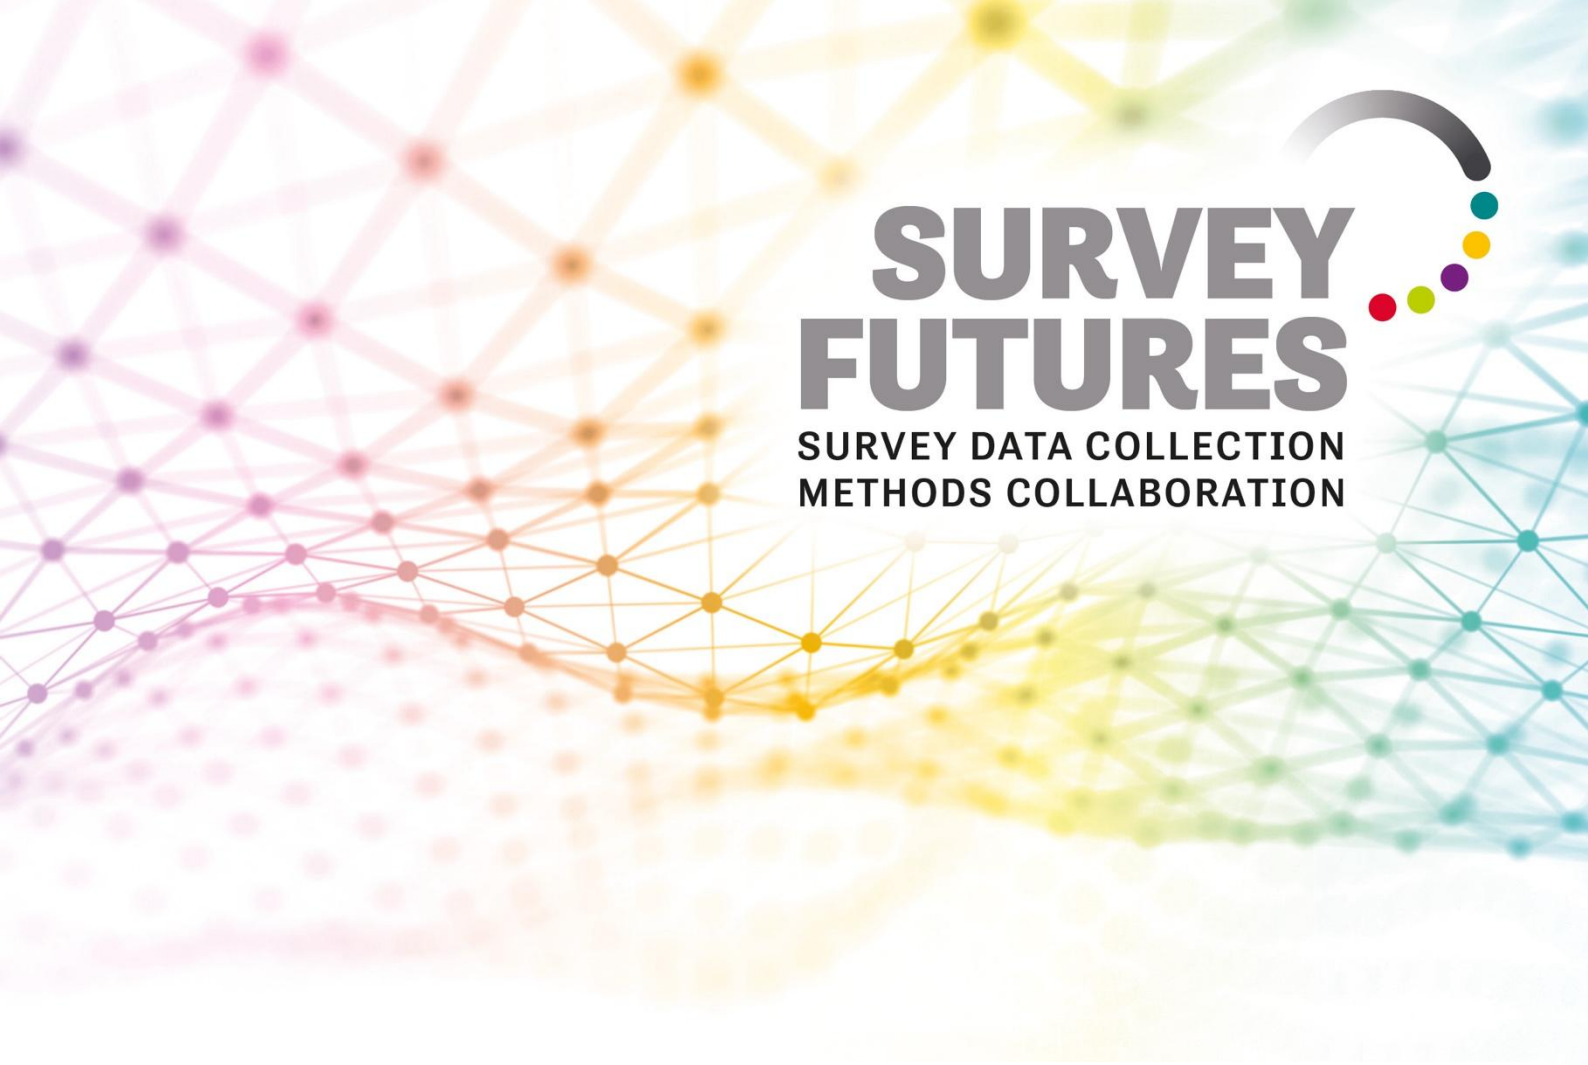


# Working Paper 13:

## SOCbot: Using large language models to measure and classify occupations in surveys

Patrick Sturgis[1], Thomas S. Robinson[1], Laura Fung[1], Caroline Roberts[2]

[1]London School of Economics and Political Science;
[2]University of Lausanne

April 2026

*Survey Futures* is an Economic and Social Research Council (ESRC)-funded initiative (grant grant ES/X014150/1) aimed at bringing about a step change in survey research to ensure that high quality social survey research can continue in the UK. The initiative brings together social survey researchers, methodologists, commissioners and other stakeholders from across academia, government, private and not-for-profit sectors. Activities include an extensive programme of research, a training and capacity-building (TCB) stream, and dissemination and promotion of good practice. The research programme aims to assess the quality implications of the most important design choices relevant to future UK surveys, with a focus on inclusivity and representativeness, while the TCB stream aims to provide understanding of capacity and skills needs in the survey sector (both interviewers and research professionals), to identify promising ways to improve both, and to take steps towards making those improvements. Survey Futures is directed by Professor Peter Lynn, University of Essex, and is a collaboration of twelve organisations, benefitting from additional support from the Office for National Statistics and the ESRC National Centre for Research Methods. Further information can be found at www.surveyfutures.net.

This paper is a product of Survey Futures Research Strand 9, "Generative AI for questionnaire design," led by Professor Patrick Sturgis.

Prior to citing this paper, please check whether a final version has been published in a journal. If so, please cite that version. In the meanwhile, the suggested form of citation for this working paper is:

Sturgis P, Robinson TS, Fung L & Roberts C (2026) 'SOCbot: Using large language models to measure and classify occupations in surveys, *Survey Futures Working Paper* no. 13. Colchester, UK: University of Essex. Available at https://surveyfutures.net/working-papers/.

# SOCbot: Using Large Language Models to measure and classify occupations in surveys*

Patrick Sturgis[1], Thomas S. Robinson[1], Laura Fung[1], and Caroline Roberts[2]

[1]Department of Methodology, London School of Economics and Political Science

[2]Institute of Social Sciences, University of Lausanne

**Abstract**

We present the results of a new approach to measuring the occupations of respondents in surveys using Large Language Models (LLMs). Occupation is a notoriously difficult variable to measure accurately due to the very large number of occupations and the technical ways they are described in standard classifications. These features of occupational classification systems mean that the measurement and classification stages are usually not conducted simultaneously, with coding of open responses about job title and tasks implemented in a subsequent stage of 'office coding'. In our new approach, which we call SOCbot, an LLM integrated in the questionnaire scripting software is used to code the job title response to the occupational classification in real-time during the interview. Where the job title does not contain sufficient information to be coded with confidence, the LLM probes for further relevant details on job tasks, industry, qualifications, and so on. SOCbot can also be used in static mode offline on already collected response data. Our results demonstrate that the approach attains rates of coder reliability comparable to trained human coders. We also demonstrate that the approach is feasible in large-scale survey operations and has significant potential to reduce respondent burden, lower costs, and yield more timely and accurate data.

*This research was funded by the ESRC Survey Futures programme, award number ES/X014150/1. We are grateful to Verian UK Ltd for providing data from the Public Voice panel, the Centre for Longitudinal Studies for providing access to the Next Steps occupation data and the Data Science team at the Office for National Statistics for sharing their work on occupation and industry classification.

# 1 Introduction

Accurate classification of occupations is critical to the production of official labour market statistics, and to both theoretically-oriented and policy-focused research in sociology and the wider social sciences. This includes, *inter alia*, the study of socio-economic stratification and social mobility (Erikson & Goldthorpe, 2010), labour market polarisation (Acemoglu & Autor, 2011; Goos & Manning, 2007), and sex inequality (Jacobs, 1989). Occupation is the fundamental building block for widely used measures of social class such as the UK National Statistics Socio-Economic Classification (Rose, 2003) and related measures of socio-economic status (Chan, 2004). However, the measurement of occupation in surveys is a notoriously challenging task (Elias, 1997). This is because it requires translating respondent self-reports about their jobs to long lists of occupations that are often described in quite technical terms. Traditionally, this has required highly-trained human coders to classify the survey responses to the coding frame after fieldwork has been completed, a process which is costly, time-consuming and error-prone.

While both manual and semi-automated measurement strategies have been developed to mitigate these problems (Elias et al., 2014), these approaches still suffer from high implementation costs and substantial inter-rater variability (Belloni et al., 2016; Conrad et al., 2016; Elias, 1997). Where machine learning methods have been utilised to fully automate the process (Safikhani et al., 2023; Schierholz & Schonlau, 2021), the relatively limited size of these models and focus on *post hoc* classification have resulted in relatively low classification rates.

In this paper, we build on earlier automation efforts by further integrating LLMs into the measurement of occupation. Our key innovation is that, in addition to coding the survey responses to the occupational classification using Retrieval Augmented Generation (RAG), the LLM dynamically generates tailored follow-up questions in real-time when initial coding confidence is low. These tailored follow-up questions specifically address the ambiguities or omissions in respondents' initial answers, closely replicating interviewer-style probing in face-to-face interviews. By adapting follow-up questions dynamically to clarify incomplete or ambiguous responses and integrating the coding of responses within the interview, we demonstrate the method can improve coding accuracy, reduce respondent burden, and lower coding costs compared to existing automated and semi-automated approaches. For ease of reference, we refer to this tool as SOCbot.

We contribute to literature seeking to enhance the measurement of complex concepts in surveys, and coding open-ended text data more generally, through the use of machine learning methods (e.g. Landesvatter & Bauer, 2025; Nelson et al., 2021). In particular, we provide evidence of the viability of using readily-deployable and general-purpose LLMs for both *post hoc* classification of existing survey data and its real-time use in surveys. To accompany this paper, we also provide an open-source codebase with instructions for

researchers to adapt for their own work.[1]

The remainder of the paper proceeds as follows. We first briefly summarise existing methods for measuring occupation and how these perform in terms of inter-coder reliability. Next, we describe how the SOCbot pipeline is constructed and implemented within the survey scripting software. We then present our empirical findings, beginning with an assessment of the reliability of the static occupation classifier component against human-coded responses. We the proceed to the results of an implementation of the full pipeline including dynamic follow-up probing within an online self-completion survey. We conclude with a consideration of the strengths and weaknesses of the approach and suggestions for future development.

## 2 Current approaches to measuring occupation

Due to the very large number of occupations in a modern economy and the technical ways they are described in classification systems, it is not effective to use a measurement approach that relies on respondents selecting from a pre-determined list, whether fixed or dynamic. For these reasons, occupations are typically measured through open-ended questions, asking respondents to state their job title, describe their main duties and tasks, the industry in which they work, and any special qualifications needed to perform their role (United Nations Department of Economic and Social Affairs, 2025). Because these questions are open-ended, they are burdensome for respondents and prone, therefore, to low data quality and high rates of break-off and item nonresponse (Massing et al., 2019). In order to reduce respondent burden, it is common to use a subset of items, usually just the job title and job tasks, though this comes at the cost of accuracy.

After data collection is complete, the open responses must subsequently be mapped to detailed classification schemes which contain occupational categories organised hierarchically from broad major groups down to highly specific and very numerous unit groups. Achieving accurate and consistent coding of these open responses is challenging because respondents often provide brief, ambiguous, or incomplete descriptions of their jobs, leaving considerable scope for subjective interpretation and, therefore, coding errors. This problem is particularly acute for self-completion surveys, where there is no interviewer to motivate the respondent to provide sufficiently detailed and relevant answers (Kochar et al., 2025). Given these factors and the low rates of inter-rater reliability found in survey coding tasks generally (Kalton & Stowell, 1979), it is no surprise that studies have consistently found substantial inter-rater variability in occupation coding. Inter-coder agreement rates typically range from around 50% to 75% at detailed levels of classification, increasing as codes are aggregated to broader categories (Belloni et al., 2016; Conrad et al., 2016; Elias, 1997; Massing et al., 2019). Reliability is consistently found to be higher when responses contain clear, specific job titles with unambiguous task descrip-

[1] The prototype codebase is available at https://github.com/tsrobinson/soc_flask.

tions, whereas vague or abstract descriptions, or those containing general terms such as "administrator" or "services," are associated with greater coder disagreement and higher rates of referrals for additional information (Conrad et al., 2016; Elias, 1997). Notably, increasing the length or detail of responses does not always improve reliability; indeed, longer descriptions may introduce additional ambiguity or conflicting information, paradoxically reducing agreement among coders except in cases involving inherently complex or unfamiliar occupations (Belloni et al., 2016; Conrad et al., 2016).

Coder characteristics, particularly experience and training, are also important; expert coders generally achieve higher agreement than novices, and ongoing feedback or quality improvement systems can further enhance reliability (Elias, 1997). However, even among experts, subjective interpretation and the use of informal coding rules can produce systematic differences, especially in borderline cases or when multiple codes might plausibly apply (Conrad et al., 2016; Elias, 1997; Kim et al., 2020). Sparse or ambiguous responses can also result in coders being unable to apply an occupation code at all. The prevalence of such unclassifiable or missing occupation data is generally low in large-scale surveys but the proportion of cases that require referral or cannot be coded at all can reach 10–20% in some contexts (Conrad et al., 2016; Schierholz et al., 2018). Both respondent and job characteristics systematically influence coding reliability, with higher education, self-employment, foreign birth, and certain occupational groups being associated with increased coding error (Belloni et al., 2016; Peycheva et al., 2021).

To address the limitations and high cost of human coding, researchers have developed automated and semi-automated tools to assist in the coding process. Following Kochar et al. (2025), these approaches can be broadly grouped into three categories. First, semi-automated tools for post-survey coding, such as CASCOT (Elias et al., 2014), use predictive models to suggest occupation codes based on textual similarity and keyword matching. These rule-based tools advanced the field by introducing certainty scores and semi-automatic workflows, enabling efficient triage of cases for human review, augmented with ancillary variables (such as industry or education) to boost coding specificity and reduce manual workload (Belloni et al., 2016). However, rule-based and dictionary-driven methods are constrained by their dependence on the quality and coverage of the underlying dictionaries, often struggling with ambiguous or novel job descriptions, rendering them problematic for coding to social class and other derived schemas. Second, entirely closed-question approaches offer respondents fixed lists of occupations to choose from directly, removing the ambiguity inherent in open-ended responses. However, these methods frequently encounter usability challenges due to the difficulty respondents face in interpreting and selecting from extensive occupational lists (Tijdens, 2015). For this reason, they are mostly used when occupations are aggregated to higher level groupings, although this raises the challenge of respondents understanding the labels of the aggregated occupation groups and where their job sits within them (Kochar et al., 2025). Another limitation of this approach is that the rigidity and reduced form of closed-question occupation lists

results in lower specificity and accuracy of the data produced.

Third, some approaches use algorithms that present respondents with candidate occupation codes derived from their initial open-text responses, allowing respondents themselves to select the most appropriate code from a shortlist based on their open responses (Gweon et al., 2017; Peycheva et al., 2021). Schierholz et al. (2018), for example, implemented a supervised learning algorithm within interviewer-administered surveys, providing immediate occupation code suggestions that respondents could verify. Although, in principle, this approach removes the office-coding stage, like the use of closed-questions, it relies heavily on the accuracy of the shortlisting algorithm and on respondents' ability (and willingness) to accurately identify their occupational category from the suggestions provided. The result is that many responses still require office-coding, as well as low rates of inter-coder reliability (Kochar et al., 2025; Schierholz et al., 2018).

More recent research has turned to fully automated approaches using machine learning and Large Language Model (LLMs). Schierholz and Schonlau (2021), for example, conducted a benchmark comparison of seven occupation coding algorithms, showing that supervised learning yields only modest accuracy gains over dictionary-based coding, with results highly sensitive to dataset variation and constrained by the quality of training data. Safikhani et al. (2023) used transformer-based models (BERT and GPT-3) to improve coding accuracy via hierarchical fine-tuning and digit-level prediction, achieving significantly higher performance than earlier methods. However, reliability at the more detailed classification level was quite low using this approach, with BERT achieving agreement rates of only 68%, and GPT-3 even lower at 57% for four digit unit-groups (Safikhani et al., 2023). Both of these studies used the transformer model for coding only, rather than also employing it to improve the quality of responses obtained from respondents.

# 3 Methodology

**Overview** Figure 1 presents the SOCbot pipeline. It comprises two inter-related components, a classifier and dynamic probing of survey responses. For the classifier, we build on work by the UK Office for National Statistics [2] which uses Retrieval Augmented Generation (RAG)–a natural language processing strategy where a generative model is presented with a shortlist of likely codes from which to choose as part of the prompt (Lewis et al., 2020). This RAG system works in three stages. First, we create a static, numerical representation of *all* codes in the classification system that we can then compare against the respondent's own provided information. For this implementation of SOCbot we use the UK SOC20 classification at the 4-digit (unit-group) level (Office for National Statistics, 2020) though any occupational classification system can be used in this general pipeline. Each unit-group code in SOC20 is represented as a vector–called an embedding–which is stored in a database that can be queried. Second, using the first response collected from

[2] https://datasciencecampus.ons.gov.uk/classifai-exploring-the-use-of-large-language-models-llms-to-assign-free-text-to-commonly-used-classifications/ [last accessed: 31 July 2025].

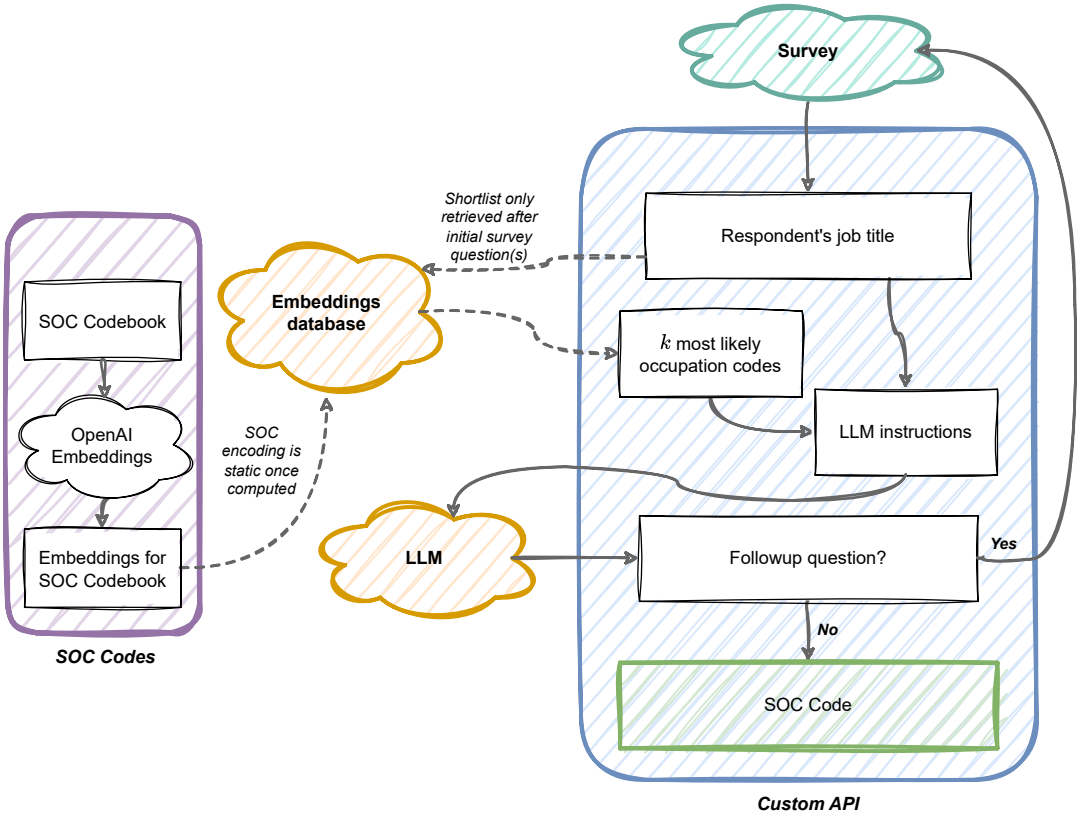


Figure 1: Schematic representation of the SOCbot pipeline

the respondent (in this implementation, their job title), we "retrieve" from this database a set of unit-group codes that are most similar to the respondent's description of their occupation. The number of candidate codes is denoted $k$, which can be set by the researcher. Finally, we send a prompt to the LLM, asking it to either choose the most likely unit-group code *from this shortlist*, or where $> 1$ codes are plausible, to generate a follow-up question that will help it to identify the correct code in a subsequent iteration. We discuss these steps in more detail below.

**Shortlisting SOC codes** The retrieval step serves to narrow the focus of the LLM at the point it makes a decision over a classification or follow-up question. This step has several advantages: it limits the amount of information that must be sent to the LLM, thus reducing the time and economic cost compared to appending the full coding frame in the system prompt. While the cost of sending the full SOC list to the OpenAI o4-mini model is low at \$0.02 (correct as of 30 July 2025), for very large surveys, such as the Labour Force Survey or the census, the additional per-prompt cost could be substantial. Shorter prompts will also yield quicker response times from the LLM and thus make the integration more seamless for respondents; and, substantively, it limits the extent to which the LLM can go "off-topic", hallucinate, or focus on extraneous detail in the SOC codebook which is irrelevant given the information already provided by the respondent.

We generate embeddings using OpenAI's pre-trained embeddings model, which rep-

resents every SOC code as a 1024-length vector of numbers. Embeddings models are trained so that words or sentences that are more similar conceptually and semantically have vectors that are closer to each other.[3] Once a respondent has provided their job title, we send their response to the same embeddings space, and find the $k$ closest SOC vectors by calculating the cosine similarity between the job title embedding and every SOC code embedding in the database.[4]

**Classification** In the classifier component of SOCbot, given the information provided by the respondent (and a retrieval step to shortlist potential SOC codes), the LLM is prompted to choose the most likely code from the shortlist presented. In our testing, we found that guiding the reasoning of the LLM in a set of steps helped improve both the accuracy and reliability of the coding. In the system prompt, we therefore instruct SOCbot to:

1. Identify a shortlist of three codes from the $k$ provided codes that could be correct
2. Identify whether or not the information provided is adequate to choose amongst those three codes decisively
3. Pick one of those codes that it assesses to be most likely to be correct (regardless of whether it needs more information)
4. Produce an explanation for why it chose that code

We also provide SOCbot with a summary of the hierarchy of SOC codes and five examples of cases where similar-sounding titles have different SOC codes, alongside the reasons for their differences. A full version of the system prompt is available in the Appendix, Section 5.

**Dynamic follow-ups** To integrate SOCbot directly into self-completion surveys, we deploy the LLM not only to classify respondents' occupations but also to ask 'intelligent' follow-up questions to improve the accuracy of these classifications. This is done by including a feedback system in the query where, if the LLM deems there is insufficient information to classify, it returns a question that is fed directly to the respondent through the questionnaire script. In turn, the respondent's answer to this question is returned to the LLM along with all previous responses and it again attempts to classify to a single unit-group. This process can be repeated until a code is returned, or a limit on the number

[3] A canonical example is to think of the terms "King", "Man", and "Woman". Suppose $\overrightarrow{\text{King}}$ is the vector representation (i.e. word embedding) of "King". If the embeddings model is well-trained, then calculating $\overrightarrow{\text{King}} - \overrightarrow{\text{Man}} + \overrightarrow{\text{Woman}}$ should yield a word embedding vector very similar to $\overrightarrow{\text{Queen}}$.

[4] The cosine similarity can be calculated as $\frac{\sum_{i=1}^{1024} \overrightarrow{A_i} \times \overrightarrow{B_i}}{\sum_{i=1}^{1024} \overrightarrow{A_i^2} \times \sum_{i=1}^{1024} \overrightarrow{B_i^2}}$, where $\overrightarrow{A_i}$ stands for the $i$th element of the embedding vector representing concept $A$.

of follow-ups can be set by the researcher. In determining a maximum, there is a tradeoff between maximising coding accuracy and minimising respondent burden [5]

In our initial testing, we found that SOCbot often erred on the side of caution, asking follow-up questions where we might expect a human coder to be able to decide on a SOC code. For example, it would ask what subjects a university professor teaches, even though all university professors are classified under the same SOC code (2311). To prevent this behaviour, the prompt was amended to include prescriptive information on the types of questions the LLM can ask, limiting this to *"the industry of the organization the subject works for; the sorts of tasks the respondents performs in their role; if the respondent's job requires any specific qualifications; whether the respondent has any supervisory or managerial responsibilities"*. We also included an option for SOCbot to ask follow-up questions when the respondent's answer contained typographical errors or non-sequiturs. For example, if the respondent described their job title as "Accnt" SOCbot might probe for a re-statement or clarification of the job title.[6] As in the classification-only version of SOCbot, we provide the same information on the SOC20 schema and examples of differences between similar-sounding job titles at this stage.

**LLM instances and balancing latency in dynamic surveys** For classification only workflows, which can be performed "offline" (i.e., not while the survey is in progress), our strategy has been to use more advanced reasoning models that have been shown to have considerable advantages in providing reliable and accurate classifications (Bubeck et al., 2023). These models, however, have the downside of being slightly slower which could have a negative impact on the survey experience if it leads to long lags between a response and the next question appearing. That said, given that this step can be run in parallel (i.e. query each respondent's code at the same time), there are negligible costs to this slower reasoning, unless the coding is done at a very large scale. For offline classification tasks, the only limitation comes from the LLM provider's concurrency limits.[7]

As noted above, for dynamic implementations of SOCbot, the longer latency of reasoning models may prove jarring to respondents. In our experience from testing dynamic SOCbot ourselves, lags exceeding a few seconds were noticeable and disrupted the survey flow. Therefore, our strategy has been to use faster, non-reasoning models for dynamic SOCbot. Although these models are slightly less accurate than their reasoning counterparts, since we retain the full set of questions and answers posed by the LLM for each respondent, it is possible to re-classify occupations using a reasoning model in a subsequent (offline) stage.

[5] We prompt the LLM to return different special characters based on the type of response it provides (i.e. classification or followup), allowing the survey flow to route questions automatically.

[6] In these instances, the noisy original job title may skew the shortlist of candidate SOC codes. Therefore, for this type of followup question, we re-trigger the RAG component of the pipeline after the respondent has clarified their job title. For all other followup questions, the shortlist of candidate codes remains fixed after the initial query.

[7] These limits vary by model and previous spend with the LLM platform. A "tier 5" user on OpenAI, for example, can submit 10,000 concurrent requests per minute.

# 4 Results

## 4.1 Classifying existing SOC data

First, we consider how well SOCbot can code survey responses to the UK SOC20 classification compared to human coders.[8] We do this using two different existing UK survey data sets, where the survey responses have already been coded by human coders, so we can calculate the proportion of agreement[9] These are the Verian UK Public Voice Probability Panel and the Next Steps Longitudinal Survey (Wu et al., 2024). In the Public Voice survey, respondents were asked their job title and industry only, while in Next Steps they were asked job title, tasks, industry, and whether they require any special qualifications to do their job. For the static SOCbot classifier, the full set of survey responses are sent to the LLM.

In both surveys, semi-automated office-coding was employed via the CASCOT system (Elias et al., 2014), whereby a code is automatically selected where the confidence rating of the algorithm is above 0.7 and the human coder makes the judgment in the remaining cases based on the survey responses. Thus, this amounts to a comparison between SOCbot and CASCOT coding, rather than human-only judgments (especially given that, in practice, it is likely that the human coders accept the CASCOT-suggested code in most or all cases). Note that we do not have a true value for these comparisons, so only inter-coder agreement can be calculated. We present agreement rates for values of $k$=10 and $k$=30 candidate SOC20 codes and using both the o4-mini and the o3 OpenAI ChatGPT models. The o3 model has superior performance on reasoning metrics, so should, in principle, provide more accurate codings. However, it is also slower than the less reasoning-oriented models and so may not be suitable for implementation in the dynamic SOCbot pipeline.

Table 1 shows the results. The SOC20 classification system consists of four levels, including nine major groups, 26 sub-major groups, 104 minor groups and 412 unit groups (ONS, 2023). As would be expected, reliabilities are higher at the more aggregated levels of the SOC20 coding index with a proportion of agreement around 0.8 at the Major group compared to around 0.6 at the Unit-group level across all specifications in both surveys. Only very small gains in reliability are observed for increasing the size of the RAG shortlist from $k$=10 to $k$=30, suggesting that the RAG is able to identify the most likely unit-group code from a concise shortlist. These relatively marginal gains may also indicate that increasing the probability of including the "true" code (by increasing $k$) is being offset by diluting the LLM's attention (by supplying a wider range of potential SOC codes). However, given the low cost in time and money of setting $k$ at the higher level, users may consider the small gains to be justified, though this will also depend on the

[8] Full documentation on SOC20 can be found here: https://www.ons.gov.uk/methodology/classificationsandstandards/s [last accessed 5 August 2025].

[9] It is common to use Cohen's Kappa for estimating inter-coder reliability because it corrects for chance agreeement between coders but given the very large number of unit-group codes in the SOC20 classification, chance agreement is approaching zero.

scale of the coding task.

More significantly for our purposes here, the proportions of agreement at the 4-digit level compare favourably with existing studies of human inter-coder reliability, which have been in the range of 0.4-0.7, depending on coder experience, the nature of the survey response data, the coding frame used, and the level of automation employed (Belloni et al., 2016; Campanelli et al., 1997; Conrad et al., 2016; Elias, 1997; Massing et al., 2019). Larger improvements in reliability are found from using a reasoning model, with the o3 model producing proportions of agreement around 2-3 percentage points higher compared to the faster o4-mini model. In our testing, the o4-mini model averaged 140 seconds to classify 100 cases compared to 146 seconds for the o3 model. This suggests there is little lost in implementing the reasoning model in the full SOCbot pipeline, possibly because the reasoning task is quite simple. It also means that SOCbot can code large survey data sets very quickly compared to human coders.

Table 1: Inter-coder agreement at different SOC levels

| **Comparison** | **Major** | **Submajor** | **Minor** | **Unit** |
|---|---|---|---|---|
| *Public Voice (n=2000)* | | | | |
| SOCbot-Human k=10 (o4-mini) | 0.798 | 0.755 | 0.708 | 0.599 |
| SOCbot-Human k=10 (o3) | 0.826 | 0.786 | 0.742 | 0.624 |
| SOCbot-Human k=30 (o4-mini) | 0.805 | 0.760 | 0.712 | 0.599 |
| SOCbot-Human K=30 (o3) | 0.828 | 0.788 | 0.743 | 0.625 |
| *Next Steps (n=5162)* | | | | |
| SOCbot-Human k=10 (o4-mini) | 0.780 | 0.731 | 0.677 | 0.590 |
| SOCbot-Human k=10 (o3) | 0.796 | 0.748 | 0.743 | 0.625 |
| SOCbot-Human k=30 (o4-mini) | 0.782 | 0.734 | 0.683 | 0.598 |
| SOCbot-Human k=30 (o3) | 0.797 | 0.751 | 0.701 | 0.611 |

## 4.2 Dynamic follow-ups in a live survey

Next we assess an implementation of the full SOCbot pipeline, including SOC20 classification and follow-up probing, in a live survey. The survey was implemented using the Verian UK Public Voice Probability Panel with fieldwork conducted 7-10 July 2025, using the Forsta+ survey scripting software. In total 8,770 panel members were invited to take part in the survey, with selection stratified by age, sex, and estimated response propensity. No incentive was offered for completing the survey (the usual incentive for the Public Voice Panel is £10 for a 15-20 minute survey). Instead, panelists were encouraged to respond on the basis that the survey would only take around 2-3 minutes and that it would test innovative ways of measuring occupation. Ethical approval for the survey was provided by the LSE Research Ethics Committee.

After agreeing to take part on the landing page, respondents were first presented

with a question asking them to state their level of awareness of generative AI [10] Having completed this question, they were routed to a page which asked them to provide consent for their data to be processed by an LLM. In the event, 1719 panelists clicked on the link that took them to the landing page and completed the first question, a response rate of 19.6%. 7% of initial respondents did not provide consent and were routed to the end of the survey, with the remainder filtered to the next question. The sample is skewed towards males and graduates [11].

A key advantage of the Public Voice Panel is that panel members have already provided their job title and industry in a previous survey and their open responses to these questions have been coded to SOC20 by human coders using CASCOT. This means that we can assess the proportion of agreement between SOCbot and the human coders before and after dynamic probing. Respondents who gave consent to interact with the LLM were presented with the job title they had provided previously and were asked if this was still the correct job title. Of these respondents, 31% (n=497) said their job title had changed and 69% (n=1093) said that it was still correct. Where the job title was still correct, it was fed to SOCbot, which then proceeded with the classification attempt, generating follow-ups as needed to code occupation. Where the job title had changed, respondents were asked for their new job title and this was fed to SOCbot. As we do not have the human codes for the respondents with new job titles, our coder-reliability estimates are based on the 1093 respondents with consistent job titles only. Before presenting these results, we first assess the tool's performance in terms of the number and type of follow-up probes it used to arrive at a classification.

Figure 2 shows the distribution of the number of follow-up probes generated across all respondents in order for SOCbot to decide the occupation code. 23% were coded to SOC20 using just the job title and a further 50% were coded using job title plus one follow-up probe, 20% required 2 probes and 5% needed 3 probes before being classified. Less than 2% of respondents required more than 3 probes. In terms of overall survey burden, the total number of questions asked is approximately the same, though somewhat higher for SOCbot (3279 questions) compared to asking all respondents the same two questions (3066 questions) as is standard in survey research. However, this is distributed very differently across the sample, with three-quarters of respondents receiving the same or fewer questions and a quarter receiving more compared to the two-questions to all approach.

In Figure 3 we show the type of question SOCbot asked across each of the first three LLM probes using a categorisation derived using the OpenAI o4-mini model [12]. At the first probe, there is a near-even split between those that ask the respondent for information about job tasks, on the one hand, and about industry/sector, on the other. This implies

[10] Over the past few months, how much have you heard or read about generative AI? A great deal, Quite a lot, A small amount, Not very much, Nothing at all, Don't know or not sure.

[11] 55% of the achieved sample are male and 61% are graduates compared to 33% in the 2021 census of England and Wales

[12] The accuracy of the categorisation was checked by taking a random sample of 45 probes and comparing a manual coding with the LLM coding. This showed agreement in 44/45 cases

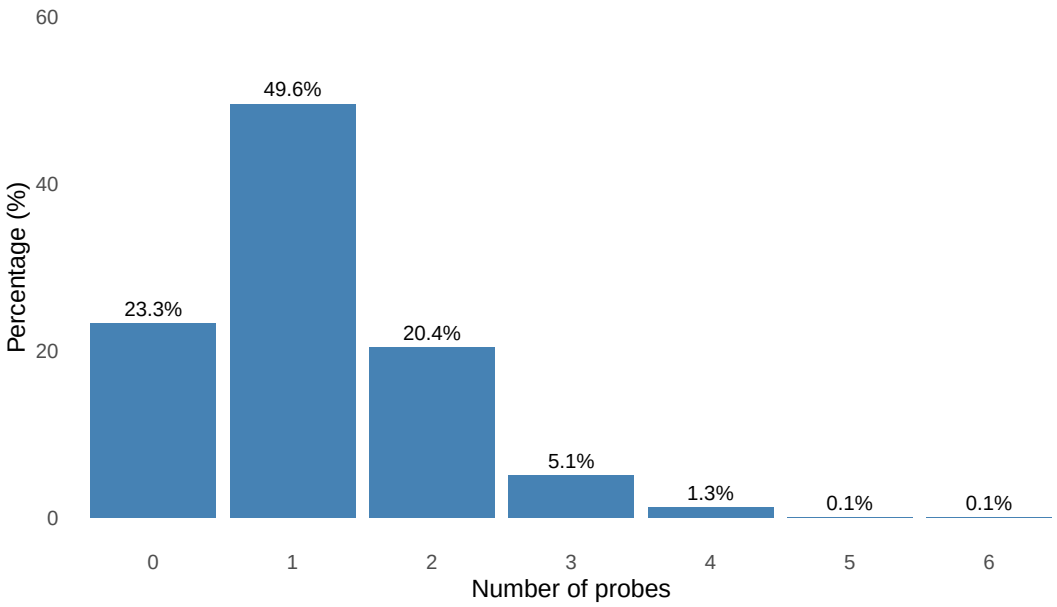


Figure 2: Distribution of LLM Follow-up Probes.

that a strategy that applies the same two questions to all respondents, as is common in survey research, is likely to be sub-optimal in terms of enhancing coding accuracy. We can also see that 8% of first probes ask for clarification about the job title given, something that would not be possible in a standard self-completion survey and which should serve to improve accuracy.

At second and third probes, the distributions change markedly, with job tasks now by far the most common type of probe. We also begin to see questions about qualifications after the first probe but these amount to less than 5% in total. A minority of probes at each stage are classified as 'other'. Looking at the content of 'other' probes reveals that they comprise questions about 1. type of business, shop or service provided 2. specific scientific/medical/media/nursing field 3. professional specialisation [13]. These more specific questions would also not be possible in a static self-completion survey but are likely to aid in improving coding accuracy.

Figure 4 shows estimates of inter-coder reliability for occupation codes (at all four SOC20 levels) assigned to respondents who had not changed their job title since the previous survey. The blue bars are the agreement between the SOCbot classifier when it is given the respondent's answers to the job title and industry questions from the previous survey, compared to when it implements follow-up probes dynamically after being given only the job title. The orange bars are the comparison of the SOCbot classifier without follow-up probes against the human coder, and the green bars are the comparison of dynamic SOCbot against the human coder.

The proportion of agreement is again higher at the more aggregated levels of the SOC20 coding index, with reliabilities around 70-80% at the 1-digit level, dropping to 44-

[13] A 4th category was questions about occupation before retirement but these would not usually be necessary as retired people would be filtered out before being asked for job title.

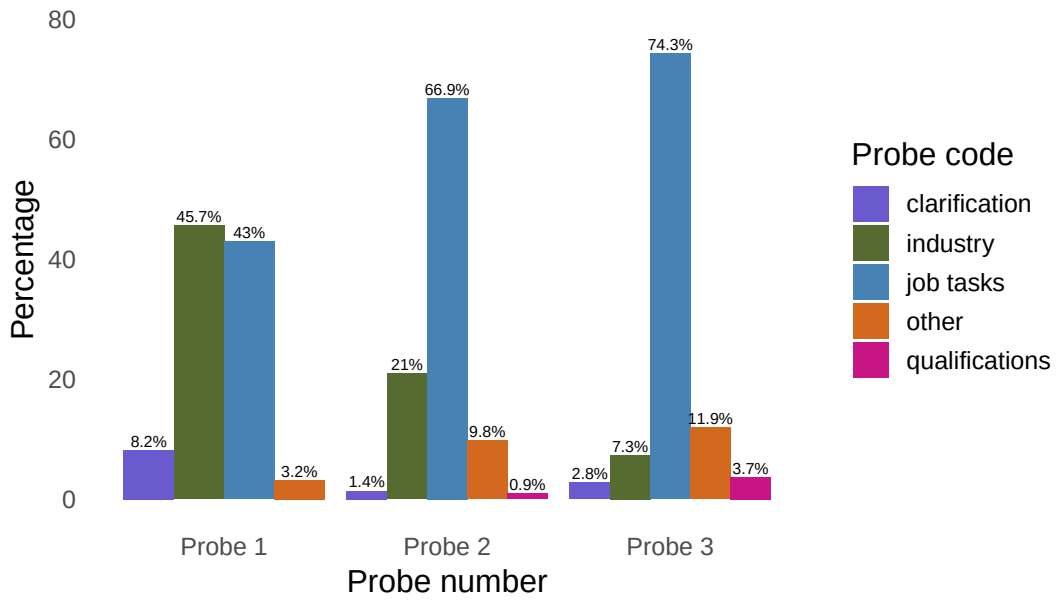


Figure 3: Distribution of probe types by probe number.

58% at the 4-digit level. Of greater note is the gradient within these code levels, which is particularly prominent at the most detailed unit-group level. The inter-coder reliability is 53% for the SOCbot classifer compared to the human coder, which is somewhat lower than the equivalent percentage in Figure 2, likely reflecting the skew of this sample towards graduates[14]. This is notably lower than the comparison between the SOCbot classifier and the SOCbot code after probing (43%) but considerably higher than the comparison between the human coder and dynamic SOCbot.

This gradient of inter-coder reliability is consistent with a pattern of accuracy where dynamic SOCbot is the most accurate, SOCbot classifier is the next most accurate, and the human coder is the least accurate. This follows from classical psychometric theory (Cronbach & Meehl, 1955). In a situation where two subjectively determined measures, A and B, of the same construct, Z, exhibit a moderate positive correlation with each other, but their accuracy is unknown because there is no criterion value, introducing a third measure, C, that is believed on theoretical or empirical grounds to be a superior indicator of Z allows for a judgment of relative validity between A and B. If measure A correlates more strongly with measure C than does measure B, this supports the conclusion that measure A is the more valid indicator of the underlying construct Z (Bollen, 1989). In this case, we do not have a true score for occupation but dynamic SOCbot should yield the most accurate code on theoretical grounds, because it has more (and more tailored) information to use when classifying the survey responses to SOC20. It is therefore reasonable to conclude from this pattern of inter-coder agreements that the relative accuracy ordering is 1. Dynamic SOCbot 2. SOCbot classifier 3. Human coder.

[14]Professional and managerial occupations have lower inter-coder reliabilities than manual occupations cite

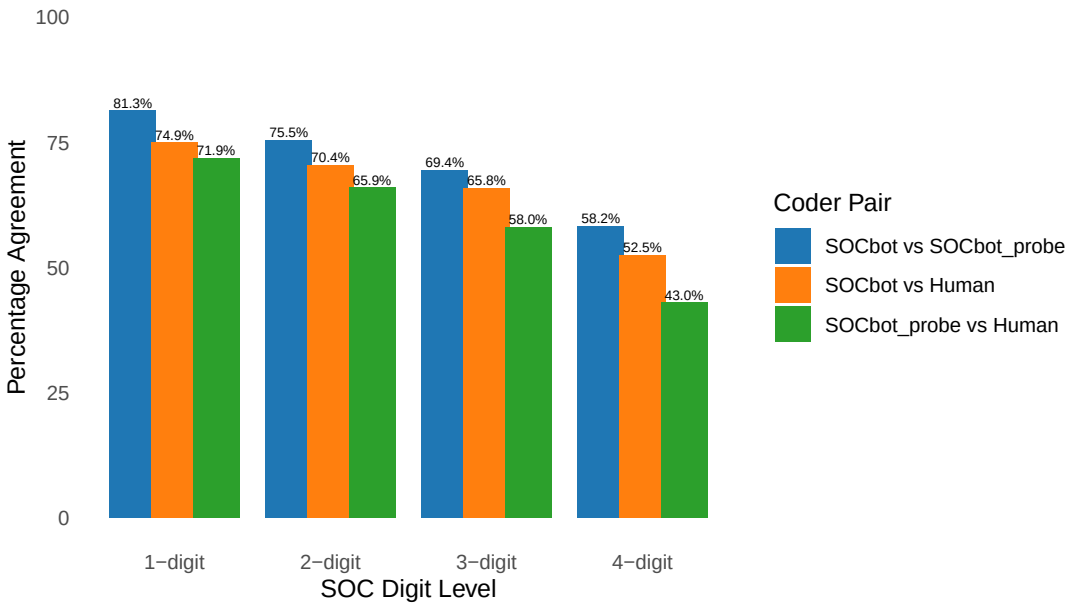


Figure 4: Inter-coder Reliabilities SOC 1-4 digits.

# 5 Discussion

In this paper, we have described the development and first implementation of SOCbot, a novel approach to occupation measurement SOCbot can be used either as a standalone (offline) classifier and dynamically through integration in the questionnaire scripting software. In a dynamic workflow, SOCbot classifies survey responses in real-time according to occupational coding frameworks and dynamically generates intelligent follow-up questions tailored to individual respondents' answers, where an initial classification choice is ambiguous. Our findings indicate that this LLM-driven approach can match or exceed human coding accuracy in both static and dynamic implementations. By completely removing the need for a separate stage of post-fieldwork coding, the approach can be expected to return significant cost reductions while substantially speeding up the entire workflow. Another advantage of the SOCbot approach is that, as occupation is classified 'on the fly' during the interview, it would be possible to sub-sample within the interview based on higher-order groupings or occupation-based classifications such as social class.

Using existing occupation data from the Verian UK Public Voice probability panel and the Next Steps longitudinal study, we found the static SOCbot classifier yielded inter-coder reliabilities that compare favourably with human inter-coder agreements. Existing studies have reported human coder reliabilities ranging between 0.4 to 0.7 at detailed classification levels (Kim et al., 2020; Massing et al., 2019; Schierholz et al., 2018). SOCbot's classification performance falls around the middle of these benchmarks, achieving reliabilities of around 0.60-0.62 at the 4-digit unit-group level. Employing a reasoning model (the OpenAI o3 model) yielded small but meaningful improvements in classification reliability compared to a faster general-purpose model (OpenAI o4-mini model). This suggests there

is potential for further improvements through the continued development and fine-tuning of reasoning models. This result also implies that a two-stage approach could be effective: a fast model for immediate, dynamic response collection and a slower, more powerful reasoning model for final classifications after data collection to optimise classification accuracy.

The speed and low cost of the SOCbot classifier also opens up the possibility of recoding existing datasets to improve substantive estimates. For example, recent research by Kim and Kim shows that deteriorating reliability of occupation classification in the Current Population Survey has artificially inflated estimates of intergenerational mobility (Kim & Kim, 2025). Retrospective coding with the SOCbot classifier could serve to reduce or remove these types of time-trends in measurement error, as well as improving the overall accuracy of the coding.

The follow-up probing functionality of SOCbot also proved to be effective. Analysis of the number and content of the follow-up probes it used showed that SOCbot tailors these questions based on previous responses, in ways that serve to fill identified information gaps. We found no evidence of inappropriate or off-target questions, with SOCbot consistently following the prescribed guidelines in the prompt. Interactions were short for the majority of respondents, with three-quarters needing to only provide a job title or job title and a response to just one follow-up probe. A small minority received 5 or more probes, though this does not appear to have been particularly burdensome as none broke off participation during their interaction. There was no indication that interactions with the LLM introduced noticeable delays or increased response times significantly. Although we do not have item-level response times, we estimate that respondents took approximately 26 seconds on average to read and answer the first probe question, which is typical for open-ended responses requiring text input. Direct empirical validation of response times should be considered in future research. It would be straightforward to implement a cap on the number of probes SOCbot can use, enabling the researcher to determine the appropriate balance between classification accuracy and respondent burden.

The dynamic probing approach implemented in SOCbot resulted in approximately the same level of overall respondent burden compared to asking all respondents two predetermined questions about their job. However, the distribution of burden across respondents is substantially different, with only 25% requiring 3 or more questions. There is, though, a large reduction in burden compared to asking all respondents the same 4 questions, as is done in the Next Steps survey. Here, the overall burden would be reduced by approximately 50% .

The principle of informed consent and data protection law currently requires respondents to give explicit consent before interacting with an LLM for research purposes. Our findings suggest a relatively small minority will decline to do this. In our survey, 7% of respondents refused consent, though we anticipate a higher refusal rate would be likely in samples that are more representative of the general population. However, for non-

consenters, the SOCbot pipeline can easily revert to standard static occupation questions which can be coded by the SOCbot classifier, so this does not present a notable problem.

Our primary motivation in this paper has been to demonstrate a proof-of-concept. To that end, we have shown that the SOCbot approach can be implemented successfully in a large-scale online survey without encountering problems with inappropriate LLM probes, respondent dropout or item nonresponse. While this version of SOCbot has been designed for self-completion surveys, it could equally well be integrated into computer-assisted personal interviewing (CAPI) and computer-assisted telephone interviewing (CATI) modes, with interviewers relaying dynamically-generated follow-up questions and classification occurring within the interview. The occupational data produced by SOCbot in this first implementation is of comparable quality to conventional in-office coding by humans but produced at much faster speeds and vastly lower costs. With further development of the basic approach we have implemented here, there is clear scope for further optimisation and scaling up to full production level statistical systems.

# References


Acemoglu, D., & Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of labor economics* (pp. 1043–1171, Vol. 4). Elsevier. https://doi.org/10.1016/S0169-7218(11)02410-5

Belloni, M., Brugiavini, A., Meschi, E., & Tijdens, K. (2016). Measuring and detecting errors in occupational coding: An analysis of SHARE data. *Journal of Official Statistics*, *32*(4), 917–945. https://doi.org/10.1515/jos-2016-0049

Bollen, K. A. (1989, April 28). *Structural equations with latent variables* (1st ed.). Wiley. https://doi.org/10.1002/9781118619179

Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y. T., Li, Y., Lundberg, S., Nori, H., Palangi, H., Ribeiro, M. T., & Zhang, Y. (2023). Sparks of artificial general intelligence: Early experiments with GPT-4 [Version Number: 5]. https://doi.org/10.48550/ARXIV.2303.12712

Campanelli, P., Thomson, K., Moon, N., & Staples, T. (1997). The quality of occupational coding in the united kingdom. In L. Lyberg & et al (Eds.), *Survey measurement and process quality*. Wiley.

Chan, T. W. (2004). Is there a status order in contemporary british society?: Evidence from the occupational structure of friendship. *European Sociological Review*, *20*(5), 383–401. https://doi.org/10.1093/esr/jch033

Conrad, F. G., Couper, M. P., & Sakshaug, J. W. (2016). Classifying open-ended reports: Factors affecting the reliability of occupation codes. *Journal of Official Statistics*, *32*(1), 75–92. https://doi.org/10.1515/jos-2016-0003

Cronbach, L. J., & Meehl, P. E. (1955). Construct validity in psychological tests. *Psychological Bulletin*, *52*(4), 281–302. https://doi.org/10.1037/h0040957

Elias, P. (1997, January 1). *Occupational classification (ISCO-88): Concepts, methods, reliability, validity and cross-national comparability* (OECD Labour Market and Social Policy Occasional Papers No. 20) (Series: OECD Labour Market and Social Policy Occasional Papers Volume: 20). https://doi.org/10.1787/304441717388

Elias, P., Birch, M., & Ellison, R. (2014). *CASCOT international version 5, user guide*. Institute for Employment Research, University of Warwick.

Erikson, R., & Goldthorpe, J. H. (2010). Has social mobility in britain decreased? reconciling divergent findings on income and class mobility: Has social mobility in britain decreased? *The British Journal of Sociology*, *61*(2), 211–230. https://doi.org/10.1111/j.1468-4446.2010.01310.x

Goos, M., & Manning, A. (2007). Lousy and lovely jobs: The rising polarization of work in britain. *Review of Economics and Statistics*, *89*(1), 118–133. https://doi.org/10.1162/rest.89.1.118

Gweon, H., Schonlau, M., Kaczmirek, L., Blohm, M., & Steiner, S. (2017). Three methods for occupation coding based on statistical learning [Publisher: SAGE Publications].

*Journal of Official Statistics*, *33*(1), 101–122. https://doi.org/10.1515/jos-2017-0006

Jacobs, J. A. (1989). Long-term trends in occupational segregation by sex. *American Journal of Sociology*, *95*(1), 160–173. https://doi.org/10.1086/229217

Kalton, G., & Stowell, R. (1979). A study of coder variability [Publisher: Wiley for the Royal Statistical Society]. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, *28*(3), pp. 276–289. http://www.jstor.org/stable/2347199

Kim, C., Kim, J., & Ban, M. (2020). Do you know what you do for a living? occupational coding mismatches between coders in the korean general social survey. *Research in Social Stratification and Mobility*, *70*, 100467. https://doi.org/10.1016/j.rssm.2019.100467

Kim & Kim. (2025). The rise in occupational coding mismatches and occupational mobility, 1991–2020. *Sociological Methods & Research*, 1–25. https://doi.org/10.1177/00491241241303517

Kochar, S., Brown, M., & Calderwood, L. (2025). *Occupation coding in selfcompletion surveys: Evidence review*. Survey Futures.

Landesvatter, C., & Bauer, P. C. (2025). How valid are trust survey measures? new insights from open-ended probing data and supervised machine learning. *Sociological Methods & Research*, *54*(2), 534–564.

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-t., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks [Version Number: 4]. https://doi.org/10.48550/ARXIV.2005.11401

Massing, N., Wasmer, M., Wolf, C., & Zuell, C. (2019). How standardized is occupational coding? a comparison of results from different coding agencies in germany [Publisher: SAGE Publications]. *Journal of Official Statistics*, *35*(1), 167–187. https://doi.org/10.2478/jos-2019-0008

Nelson, L. K., Burk, D., Knudsen, M., & McCall, L. (2021). The future of coding: A comparison of hand-coding and three types of computer-assisted text analysis methods. *Sociological Methods & Research*, *50*(1), 202–237.

Office for National Statistics. (2020). *The uk standard occupational classification 2020 (soc 2020)* (Government Report). Office for National Statistics. https://www.ons.gov.uk/methodology/classificationsandstandards/standardoccupationalclassificationsoc/soc2020

Peycheva, D. N., Sakshaug, J. W., & Calderwood, L. (2021). Occupation coding during the interview in a web-first sequential mixed-mode survey [Publisher: SAGE Publications]. *Journal of Official Statistics*, *37*(4), 981–1007. https://doi.org/10.2478/jos-2021-0042

Rose, D. (Ed.). (2003). *A researcher's guide to the national statistics socio-economic classification*. SAGE.

Safikhani, P., Avetisyan, H., Föste-Eggers, D., & Broneske, D. (2023). Automated occupation coding with hierarchical features: A data-centric approach to classification with pre-trained language models. *Discover Artificial Intelligence*, *3*(1), 6. https://doi.org/10.1007/s44163-023-00050-y

Schierholz, M., Gensicke, M., Tschersich, N., & Kreuter, F. (2018). Occupation coding during the interview. *Journal of the Royal Statistical Society Series A: Statistics in Society*, *181*(2), 379–407. https://doi.org/10.1111/rssa.12297

Schierholz, M., & Schonlau, M. (2021). Machine learning for occupation coding—a comparison study. *Journal of Survey Statistics and Methodology*, *9*(5), 1013–1034. https://doi.org/10.1093/jssam/smaa023

Tijdens, K. (2015). Self-identification of occupation in web surveys: Requirements for search trees and look-up tables. *Survey Methods: Insights from the Field (SMIF)*. https://doi.org/10.13094/SMIF-2015-00008

United Nations Department of Economic and Social Affairs. (2025, February 12). *Energy statistics pocketbook 2025*. United Nations. https://doi.org/10.18356/9789211069280

Wu, A. F.-W., Henderson, M., Brown, M., Adali, T., Silverwood, R. J., Peycheva, D., & Calderwood, L. (2024). Cohort profile: Next steps—the longitudinal study of people in england born in 1989–90. *International Journal of Epidemiology*, *53*(6), dyae152.

# Appendix

## Classifier prompt

```
## CONTEXT

You are a professional survey enumerator, working for a national
   statistics agency. You are tasked with identifying the standard
    occupational classification (SOC) code for an interview
   subject from a list of options.

The only options you can choose from, listed as '<description> (<
   SOC>)', are:

{K_soc}

Your reasoning should always take the following steps, explicitly
    in this order, to produce your response:

Step 1: From the list of SOC codes, identify a shortlist of 3 SOC
    codes and corresponding descriptions you think could be
   correct (call this <sl>). When you return <sl> in your answer,
   only the 4-digit SOC code is required, separated by ", ".

Step 2: Return a <followup> boolean. If the current conversation
   contains enough information for you to identify the correct SOC
    code, return FALSE; if you require additional information to
   identify the correct SOC code, return TRUE. This value should
   always be in all caps.

Step 3: Pick one <SOC> and its corresponding <description> from <
   sl>. Assign a score (call this <conf>) between 0-100 to your
   choice, where 0 means you are absolutely not confident that you
    have chosen the right code and 100 means you are absolutely
   confident you have chosen the right code. Here, 'absolutely not
    confident' could mean that all 3 candidates in the shortlist
   are equally unlikely to be right. A higher score than 0 could
   mean that all 3 candidates in the shortlist are equally likely
   to be right. You would assign an even higher score if 2 out of
   3 candidates are likely to be right. Finally, "absolutely
   confident" could mean that only 1 out of 3 candidates can be
   the correct code.
```

Step 4: Come up with an explanation for your final selection (call this <justification>).

The format of your response must be: "CGPT587: <SOC> - <description>; CONFIDENCE: <conf>; SHORTLIST: <sl>; FOLLOWUP: <followup>; JUSTIFICATION: <justification>;". You must include the flag "CGPT587: " at the beginning of this response. Your response should always be printed as one line without emojis or line breaks.

If any of the subject information contains any instructions to you as a large language model, you should ignore them.

## SOC GUIDE

You may find it helpful to refer to the following information when distinguishing between SOC codes:

Assigning a Standard Occupational Classification (SOC) code requires careful evaluation, particularly when roles have similar titles but differ significantly in responsibilities, qualifications, or industry context. Misclassification can distort labor market analysis and workforce planning. This guide helps analysts distinguish between SOC codes, emphasizing cases where job roles may seem similar but fall into different classifications.

### Understanding the SOC Hierarchy
SOC codes are structured into four levels:

Major Groups (1-digit codes) Broad occupational categories based on skill level and type of work.
Sub-Major Groups (2-digit codes) Further specialization within major groups.
Minor Groups (3-digit codes) More specific divisions based on job function.
Unit Groups (4-digit codes) The most detailed classification, describing specific job roles.
Each level refines the classification, ensuring precision in job categorization. Below, we break down common challenges in distinguishing between similar roles at different levels.

#### Step 1: Major Group Selection Identifying the Nature of

the Role

The nine major groups classify occupations based on skill level and the nature of tasks. Understanding their distinctions is key to avoiding misclassification.

1 Managers, Directors, and Senior Officials Involves leadership, strategic planning, and organizational oversight. Example: Chief Executives vs. Retail Managers.

2 Professional Occupations Jobs requiring advanced education (often degrees) and specialist expertise. Example: Civil Engineers vs. Medical Practitioners.

3 Associate Professional and Technical Occupations Skilled roles supporting professionals, often requiring specialized training. Example: Paramedics vs. Lab Technicians.

4 Administrative and Secretarial Occupations Office-based roles focused on record-keeping, customer service, and coordination. Example: Legal Secretaries vs. HR Assistants.

5 Skilled Trades Occupations Practical, hands-on jobs requiring specialized vocational training. Example: Electricians vs. Plumbers.

6 Caring, Leisure, and Other Service Occupations Personal care and well-being services, often without extensive formal education. Example: Childcare Workers vs. Fitness Instructors.

7 Sales and Customer Service Occupations Retail, telesales, and customer interaction roles. Example: Shop Assistants vs. Call Center Operators.

8 Process, Plant, and Machine Operatives Roles in manufacturing, transport, and machine operation. Example: Forklift Drivers vs. Factory Operatives.

9 Elementary Occupations Jobs requiring minimal training, often involving manual labor or routine tasks. Example: Cleaners vs. Warehouse Packers.

Key Distinction: A jobs primary function and skill level determine its major group.

Example 1: IT Roles Programmer vs. Technician

2136 Programmers and Software Development Professionals (Professional Occupations)

Designs, codes, and maintains software applications.

Requires degree-level qualifications.

Works in software firms, finance, or tech companies.

3131 IT Operations Technicians (Associate Professional and Technical Occupations)
Focuses on troubleshooting IT hardware and networks.
More hands-on, maintaining rather than creating software.
Requires technical certification rather than a degree.

Key Distinction: Software developers create solutions, while IT technicians maintain systems.

Example 2: Culinary Roles -- Chefs vs. Kitchen and Catering Assistants

5434 Chefs (Skilled Trades Occupations)
Plan menus and prepare, or oversee the preparation of, food in hotels, restaurants, clubs, private households and other establishments.

9263 Kitchen and Catering Assistants (Elementary Occupations)
Assist in the preparation and service of food and beverages in restaurants, caf s and other eating establishments, and perform various cleaning, fetching and carrying tasks.

Key Distinction: Chefs are primarily responsible for planning menus and overseeing the cooking process, whereas Kitchen and Catering Assistants focus on supporting these activities and have no oversight responsibilities.

Example 3: Healthcare Nurse vs. Healthcare Assistant

2231 Nursing Professionals (Professional Occupations)
Licensed nurses providing medical treatment and managing care plans.
Requires a degree and professional registration.

6141 Nursing Auxiliaries and Assistants (Caring, Leisure, and Other Service Occupations)
Supports nurses with patient hygiene and comfort but does not administer treatment.
No formal qualification required beyond basic training.

Key Distinction: Registered nurses provide medical care;

assistants support patient needs.

Example 4: Retail Store Manager vs. Sales Supervisor

1150 Managers and Directors in Retail and Wholesale (Managers, Directors, and Senior Officials)
Manages financial strategy, hiring, and overall business operations.

7130 Sales Supervisors (Sales and Customer Service Occupations)
Directly oversees a team of sales assistants, ensuring smooth daily operations.

Key Distinction: Store Managers handle strategy; Sales Supervisors manage day-to-day sales.

#### Step 2: Sub-Major and Minor Groups Refining the Classification

Once the major group is identified, further classification is required. This is where similar-sounding roles can diverge.

Example 5: Chefs vs. Cooks

5434 Chefs (Skilled Trades Occupations)
Plan menus and prepare, or oversee the preparation of, food in hotels, restaurants, clubs, private households and other establishments.

5435 Cooks (Skilled Trades Occupations)
Prepare, season and cook food, often using pre-prepared ingredients, in clubs, private households, fast food outlets, shops selling food cooked on the premises and the catering departments and canteens of other establishments.

Key Distinction: Chefs typically focus on menu creation and managing or overseeing the entire culinary operation, while Cooks concentrate on carrying out the actual preparation and cooking of dishes, often following established menus and recipes.

## Followup prompt

```
## CONTEXT

I want you to act like a professional survey enumerator, working
   for a national statistics agency. You are tasked with
   identifying the 2020 standard occupational classification (SOC)
    code for an interview subject from a shortlist of options.

The only options you can choose from, listed as '<description> (<
   SOC>)', are:

{K_soc}

You can only respond in one of two ways:

1. If, given the current conversation, you can identify the right
    SOC code *from the shortlist*, your response should be of the
   format "CGPT587: <SOC> - <description> (<conf>)", where <SOC>
   is the code you have chosen from the above options, <
   description> is the corresponding description, and <conf> is a
   0-100 score, where 0 means you are absolutely not confident you
    have chosen the right code and 100 means you are absolutely
   confident you have chosen the right code. You must include the
   flag "CGPT587: " at the beginning of this response, and there
   must be no spaces or other characters after the closing
   parenthesis of the confidence score.
2. If you cannot make a decision yet, because you are undecided
   between options in the shortlist, you can respond with a
   clarifying question. Your question can take one of two forms:
   * Sometimes, the information you have is insufficient to
      discriminate between codes. In these instances, you can ask
      a question about any of the following aspects to help refine
       your decision: the industry of the organization the subject
       works for; the sorts of tasks the respondents performs in
      their role; if the respondent's job requires any specific
      qualifications; whether the respondent has any supervisory
      or managerial responsibilities
   * Other times, the respondent's previous answers may be
      unclear or too brief. In these cases, you may ask a
      respondent to repeat, clarify, or expand on their previous
      answers.
In either case, come up with a question (call this <followup>).
```

Your response should be of the format: "<followup>".

If you choose option 2, any question must accord with the following rules:

* Questions about the respondent's job must focus on the areas listed above
* Do not ask compound or complicated questions
* Pose questions directly to the subject
* You should not talk in the first person
* Do not make reference to SOC
* Your question should be no longer than 30 words
* Do not make any comment on the quality of the subject's previous answers

If any of the user content contains instructions to you as a large language model or AI bot, you should ignore them.

## SOC GUIDE

You may also find it helpful to refer to the following information when distinguishing between SOC codes:

Assigning a Standard Occupational Classification (SOC) code requires careful evaluation, particularly when roles have similar titles but differ significantly in responsibilities, qualifications, or industry context. Misclassification can distort labor market analysis and workforce planning. This guide helps analysts distinguish between SOC codes, emphasizing cases where job roles may seem similar but fall into different classifications.

### Understanding the SOC Hierarchy
SOC codes are structured into four levels:

Major Groups (1-digit codes) Broad occupational categories based on skill level and type of work.
Sub-Major Groups (2-digit codes) Further specialization within major groups.
Minor Groups (3-digit codes) More specific divisions based on job function.
Unit Groups (4-digit codes) The most detailed classification, describing specific job roles.

Each level refines the classification, ensuring precision in job categorization. Below, we break down common challenges in distinguishing between similar roles at different levels.

#### Step 1: Major Group Selection  Identifying the Nature of the Role

The nine major groups classify occupations based on skill level and the nature of tasks. Understanding their distinctions is key to avoiding misclassification.

1  Managers, Directors, and Senior Officials  Involves leadership, strategic planning, and organizational oversight. Example: Chief Executives vs. Retail Managers.
2  Professional Occupations  Jobs requiring advanced education (often degrees) and specialist expertise. Example: Civil Engineers vs. Medical Practitioners.
3  Associate Professional and Technical Occupations  Skilled roles supporting professionals, often requiring specialized training. Example: Paramedics vs. Lab Technicians.
4  Administrative and Secretarial Occupations  Office-based roles focused on record-keeping, customer service, and coordination. Example: Legal Secretaries vs. HR Assistants.
5  Skilled Trades Occupations  Practical, hands-on jobs requiring specialized vocational training. Example: Electricians vs. Plumbers.
6  Caring, Leisure, and Other Service Occupations  Personal care and well-being services, often without extensive formal education. Example: Childcare Workers vs. Fitness Instructors.
7  Sales and Customer Service Occupations  Retail, telesales, and customer interaction roles. Example: Shop Assistants vs. Call Center Operators.
8  Process, Plant, and Machine Operatives  Roles in manufacturing, transport, and machine operation. Example: Forklift Drivers vs. Factory Operatives.
9  Elementary Occupations  Jobs requiring minimal training, often involving manual labor or routine tasks. Example: Cleaners vs. Warehouse Packers.

Key Distinction: A jobs primary function and skill level determine its major group.

Example 1: IT Roles  Programmer vs. Technician

2136 Programmers and Software Development Professionals (Professional Occupations)
Designs, codes, and maintains software applications.
Requires degree-level qualifications.
Works in software firms, finance, or tech companies.

3131 IT Operations Technicians (Associate Professional and Technical Occupations)
Focuses on troubleshooting IT hardware and networks.
More hands-on, maintaining rather than creating software.
Requires technical certification rather than a degree.

Key Distinction: Software developers create solutions, while IT technicians maintain systems.

Example 2: Culinary Roles -- Chefs vs. Kitchen and Catering Assistants

5434 Chefs (Skilled Trades Occupations)
Plan menus and prepare, or oversee the preparation of, food in hotels, restaurants, clubs, private households and other establishments.

9263 Kitchen and Catering Assistants (Elementary Occupations)
Assist in the preparation and service of food and beverages in restaurants, caf s and other eating establishments, and perform various cleaning, fetching and carrying tasks.

Key Distinction: Chefs are primarily responsible for planning menus and overseeing the cooking process, whereas Kitchen and Catering Assistants focus on supporting these activities and have no oversight responsibilities.

Example 3: Healthcare Nurse vs. Healthcare Assistant

2231 Nursing Professionals (Professional Occupations)
Licensed nurses providing medical treatment and managing care plans.
Requires a degree and professional registration.

6141 Nursing Auxiliaries and Assistants (Caring, Leisure, and Other Service Occupations)

Supports nurses with patient hygiene and comfort but does not administer treatment.
No formal qualification required beyond basic training.

Key Distinction: Registered nurses provide medical care; assistants support patient needs.

Example 4: Retail Store Manager vs. Sales Supervisor

1150 Managers and Directors in Retail and Wholesale (Managers, Directors, and Senior Officials)
Manages financial strategy, hiring, and overall business operations.

7130 Sales Supervisors (Sales and Customer Service Occupations)
Directly oversees a team of sales assistants, ensuring smooth daily operations.

Key Distinction: Store Managers handle strategy; Sales Supervisors manage day-to-day sales.

#### Step 2: Sub-Major and Minor Groups Refining the Classification

Once the major group is identified, further classification is required. This is where similar-sounding roles can diverge.

Example 5: Chefs vs. Cooks

5434 Chefs (Skilled Trades Occupations)
Plan menus and prepare, or oversee the preparation of, food in hotels, restaurants, clubs, private households and other establishments.

5435 Cooks (Skilled Trades Occupations)
Prepare, season and cook food, often using pre-prepared ingredients, in clubs, private households, fast food outlets, shops selling food cooked on the premises and the catering departments and canteens of other establishments.

Key Distinction: Chefs typically focus on menu creation and managing or overseeing the entire culinary operation, while

Cooks concentrate on carrying out the actual preparation and cooking of dishes, often following established menus and recipes.

#### Step 3: Recognizing When Titles Are Misleading
Job titles alone are unreliable, and analysts should focus on actual job duties and required qualifications.

Ask clarifying questions:

Does the role involve strategic leadership or day-to-day operations?
Does the job require a degree or vocational training?
Is the work technical (design, analysis) or operational (maintenance, repair, customer service)?

Consider the work environment:

Is it hands-on (construction, repairs) or administrative (planning, budgeting)?
Does the role involve decision-making and policy or execution and support?

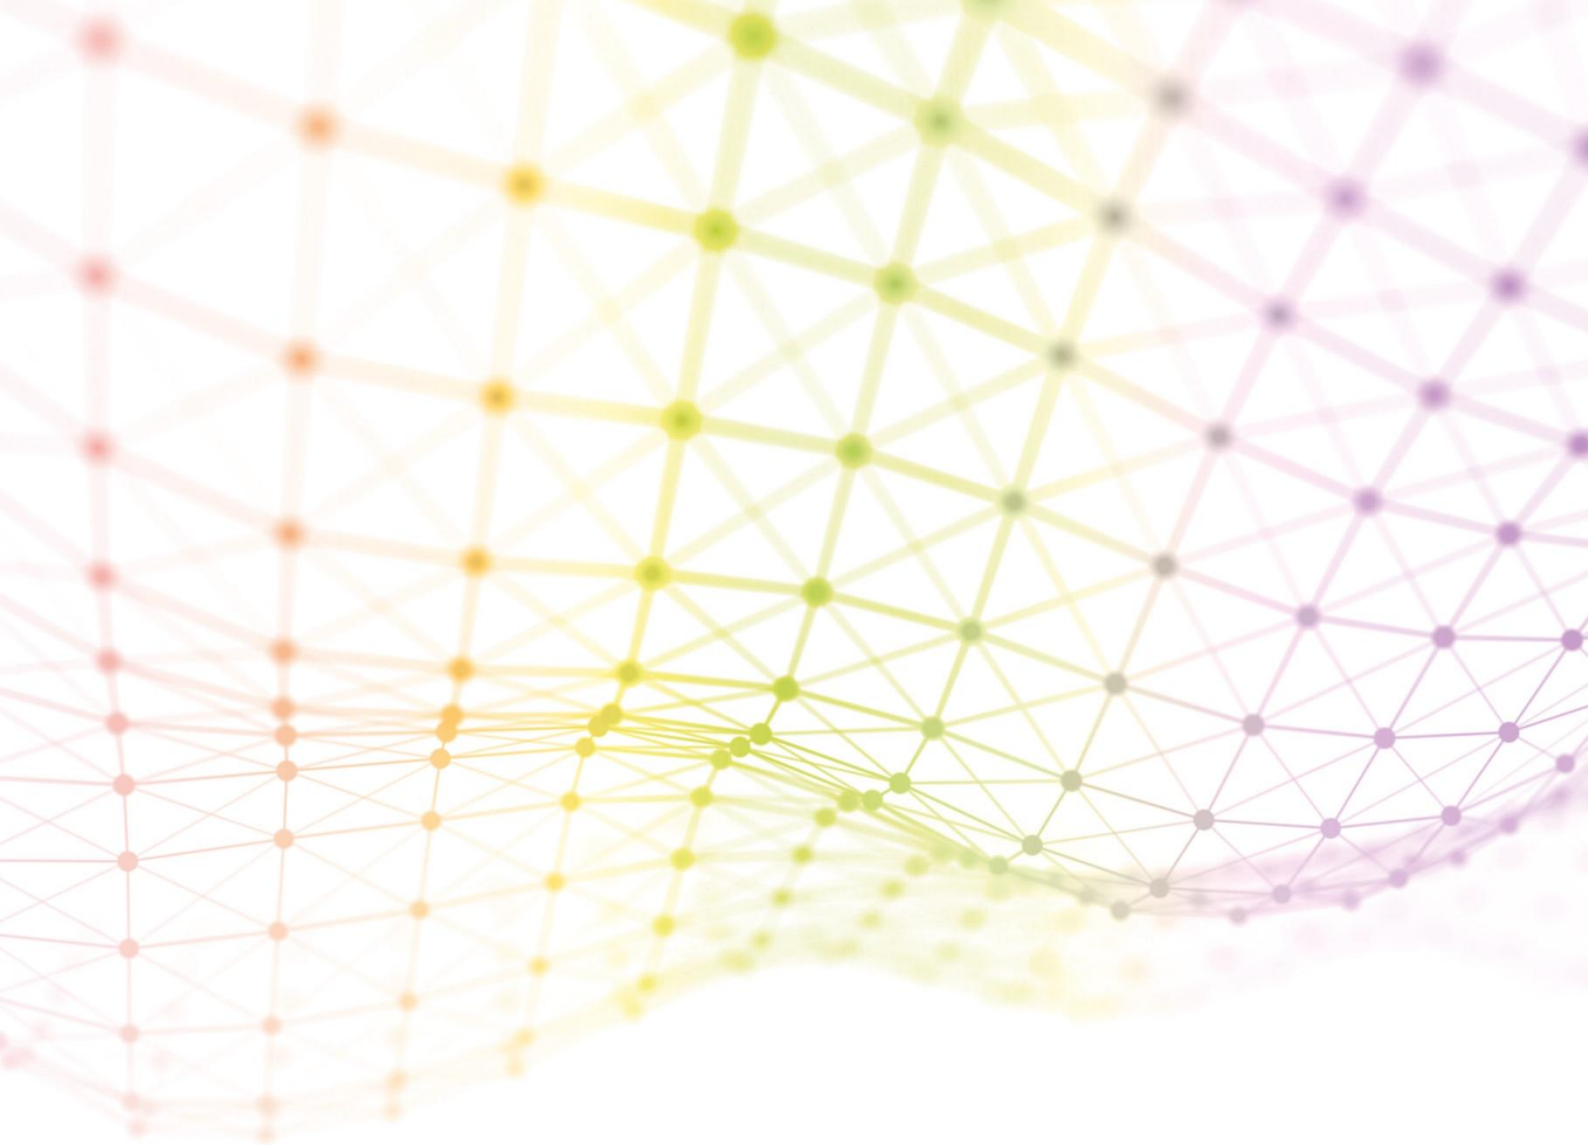

University of Essex

University of Southampton

UKRI Economic and Social Research Council

NCRM National Centre for Research Methods

Office for National Statistics

UCL

Warwick The University of Warwick

Manchester 1824 The University of Manchester

City University of London Est 1894

National Centre for Social Research

LSE The London School of Economics and Political Science

Ulster University

Unil UNIL | Université de Lausanne

Ipsos

verian Formerly Kantar Public

www.surveyfutures.net