UK Government

Impact of Growth of Data Centres on Energy Consumption – Follow Up Study

Report prepared by Europe Economics

Department for Energy Security and Net Zero (2026)

Contents

1 Executive Summary

This report, commissioned by the Department for Energy Security and Net Zero (DESNZ), examines how the growth of digital services, and the data centres that support them, affects electricity consumption in the UK, as a follow-up to Europe Economics' earlier work.[1] While digitalisation increases demand for electricity through data centre and network use, it can also displace electricity-intensive activities. Understanding this balance is critical for assessing the implications of digital growth for the UK's net-zero objectives.

Methodological approach

The study applies and extends the methodological framework developed in Europe Economics' earlier work to estimate the electricity consumption of digital services relative to counterfactuals. As before, the analysis is conducted using representative scenarios that deliver broadly the same level of service to end-users in both the digital and counterfactual cases. This approach abstracts from any increase in activity resulting from digitalisation, such as higher demand driven by lower costs or increased output due to productivity gains, and instead focuses on a hypothetical setting in which digital substitution does not alter the overall scale of economic activity. This allows policymakers to isolate the direct electricity implications of digitalisation, separate from the electricity effects of economic growth.

The analysis focuses explicitly on electricity consumption across the full delivery chain, rather than carbon emissions or individual system components in isolation. For digital services, this includes electricity used by data centres (covering both IT equipment and supporting infrastructure), transmission networks, end-user devices and offices. For the counterfactuals, the analysis captures electricity use associated with a wider variety of activities. The analysis considers operational electricity use only and excludes embodied energy associated with the manufacture of capital equipment, including digital infrastructure.

For **digital services**, electricity consumption is disaggregated into four areas, each modelled using activity-based assumptions tailored to the specific use case:

- **Data centre operations**, covering IT electricity used for generic processes (e.g. storage, downloading/streaming), together with supporting electricity, such as for cooling and lighting, captured through a Power Usage Effectiveness (PUE) uplift;
- **Internet transmission**, reflecting the electricity required to transfer data between data centres and end-users;
- **End-user devices**, estimated based on device type and active usage time; and
- **Office electricity consumption**, based on the electricity consumed in office environments by workers performing AI-assisted tasks.

The methodology distinguishes between shared processes, such as uploading and long-term storage, where electricity use may be allocated across multiple users, and user-specific

[1] See our previous study: DESNZ (2025) *"Impact of growth of data centres on energy consumption"* [online]

processes, such as streaming or downloading, where electricity use is attributed entirely to the individual user.

For the **counterfactuals**, the methodology is tailored to the characteristics of each use case rather than applying a single counterfactual structure. In particular, the modelling approach differs between:

- **Office-based services** (AI-assisted tasks): electricity use in the physical counterfactual is estimated based on task duration and average office electricity consumption per worker, incorporating lighting, heating, cooling and office equipment. Our high figure assumes an all-electric office, whereas our low figure assumes that the office also uses another fuel (e.g. for heating).
- **Bespoke services and devices**: UK-attributable electricity consumption is estimated across the relevant activities replaced by the digital service. These activities vary widely, from DTT (in the case of video streaming), to teaching in an educational building (in the case of online education), to manufacture, distribution and use of a gaming CD (in the case of cloud gaming), to local server back-ups (in the case of cloud back-ups).

We carry out sensitivity analyses using low, medium and high electricity assumptions to explore uncertainty in our estimates (e.g. caused by variations in the energy-efficiency of infrastructure).

Findings

We applied the methodology to ten use cases:

- Video streaming versus DTT
- Cloud gaming versus use of a gaming console or gaming laptop
- Online education versus in-person education
- Cloud back-up versus local server back-up
- AI-assisted email drafting versus manual email drafting
- AI-assisted code development versus manual code development
- AI-assisted copy-editing versus manual copy-editing
- AI-assisted data extraction versus manual data extraction
- AI-assisted transcription versus manual transcription
- AI-assisted forecasting model development versus manual model development

The results show a clear split across types of services. For all AI-assisted office tasks, the digital option consistently consumes less electricity than the manual alternative across low, medium and high scenarios. In these cases, the electricity required for AI training and inference, supporting data centre activity and internet transmission is outweighed by reductions in office electricity consumption per task resulting from the increased productivity of workers.

By contrast, the results are mixed for the non-AI use cases. Video streaming and cloud gaming consume more electricity than their counterfactuals across the scenarios considered, driven primarily by data transmission requirements in the digital scenario. Cloud back-up and online education, however, show the opposite pattern: the digital approach consumes less electricity than their counterfactuals. For cloud back-up, the difference is driven by inefficiencies in long-term storage in local server environments, while for online education it is driven primarily by building electricity use and travel in the in-person counterfactual.

We conducted further analysis on the potential UK electricity impact of shifting all current users of broadcast television (DTT) to video streaming, and found that this could increase UK electricity demand by around 1.8 TWh.

The table below summarises total UK electricity consumption for each use case under low, medium, and high assumptions.

Table 1.1: UK electricity consumption by use case and scenario (KWh)

	Scenario	Low	Medium	High
Video streaming versus TV streaming	Digital	0.47	0.77	1.05
	Counterfactual	0.19	0.41	0.62
Cloud gaming versus gaming on computer	Digital	5.23	8.49	20.77
	Counterfactual (UK manufactured)	3.53	7.22	10.94
	Counterfactual (Overseas manufactured)	3.16	6.75	10.37
	Counterfactual (UK, decarbonised)	3.60	7.32	11.07
	Counterfactual (Overseas decarbonised)	3.18	6.78	10.41
Online education versus in-person education	Digital	3.17	3.97	4.81
	Counterfactual	3.99	6.10	9.24
	Counterfactual (decarbonised)	6.17	9.48	14.40
Cloud back-up versus local server back-up	Digital	0.42	0.95	1.89
	Counterfactual	1.25	4.88	17.25
AI-assisted email drafting versus manual email drafting	Digital	1.76	7.54	27.61
	Counterfactual	4.41	12.56	32.29
AI-assisted code development versus manual code development	Digital	4.76	18.84	101.89
	Counterfactual	10.58	31.39	145.31
AI-assisted copy-editing versus manual copy-editing	Digital	0.98	3.09	6.18
	Counterfactual	1.63	3.62	6.71
AI-assisted data extraction versus manual data extraction	Digital	0.28	0.88	2.61
	Counterfactual	0.49	1.31	2.97
AI-assisted transcription versus manual transcription	Digital	0.27	1.03	3.13
	Counterfactual	1.06	2.51	4.84

	Scenario	Low	Medium	High
AI-assisted demand forecasting model development	Digital	11.96	38.91	120.39
	Counterfactual	39.68	94.18	181.64

Source: Europe Economics analysis.

The figure below summarises results across all use cases by showing the ratio of electricity consumption in the digital approach relative to electricity consumption in the counterfactual, based on medium scenario. Values above 100 per cent indicate higher electricity consumption in the digital scenario, while values below 100 per cent indicate lower consumption than under the counterfactual.

Figure 1.1: Ratio of digital to counterfactual electricity consumption (medium scenario)

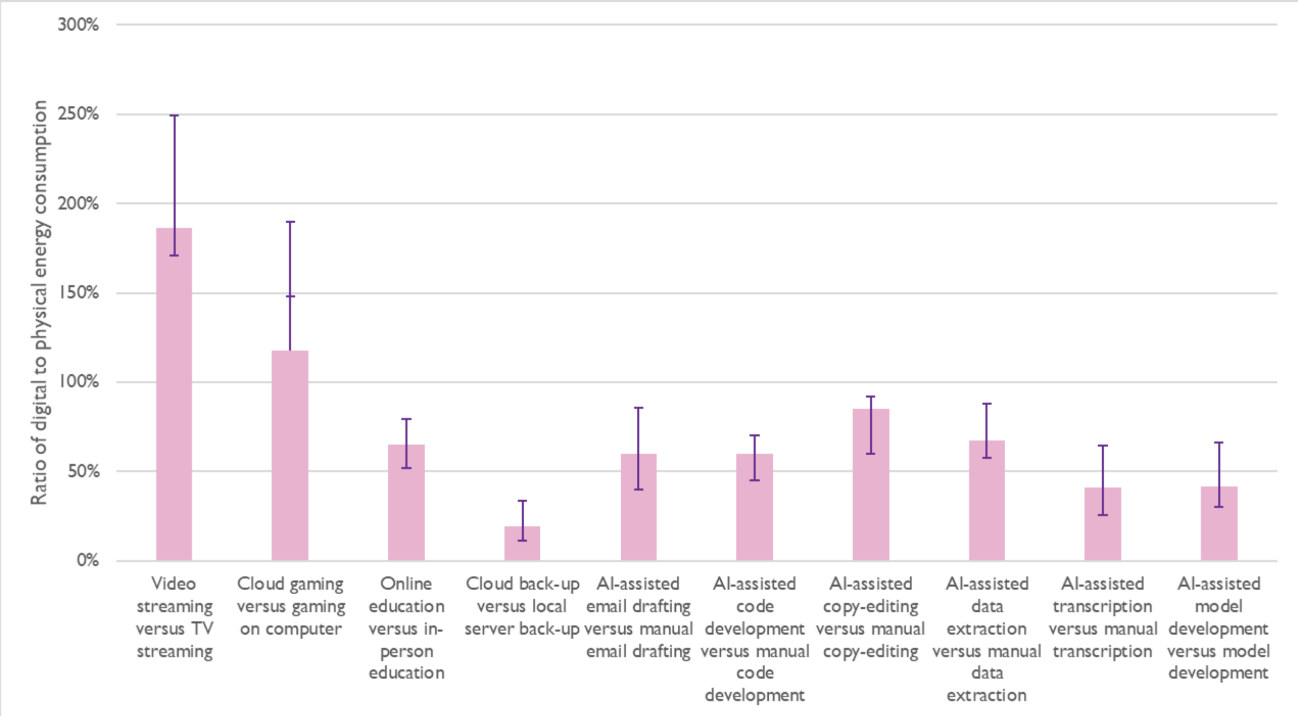

Note: The pink bars show ratios calculated using the medium scenario for both digital and counterfactual pathways. Error bars reflect the range obtained when comparing low digital with low counterfactual estimates and high digital with high counterfactual estimates. In some cases, the top end of the error bar may come from the comparison of the low estimates, while the low end of the error bar may come from comparison of the high estimates. In one case, cloud gaming versus gaming on computer, the ratio obtained using medium estimates is the lowest value, and so the uncertainty range for this use case extends only above the central estimate.
Source: Europe Economics analysis.

Overall, the results suggest that when workers use AI to increase their productivity it typically reduces electricity consumption per task, whereas for non-AI use cases digitalisation can either increase or reduce consumption depending on the nature of the service.

2 Introduction

2.1 Objective of research

Europe Economics has been commissioned by DESNZ to carry out a follow-up study on the electricity implications of data centre growth.[2] The central question is how far rising demand for digital services increases electricity use, once the avoided electricity consumption of the alternative is taken into account. This study extends Europe Economics' earlier work by examining ten additional digital use cases and comparing the electricity required for each digital service with that of its counterfactual. The aim is to broaden the evidence base and provide DESNZ with a clearer understanding of the net electricity impact of digitalisation.

In this study, we have used the methodology which we developed in the last study for quantifying the net electricity implications of digitalisation, by comparing electricity use under a digital scenario with that of a credible counterfactual. The comparison is based on a representative scenario for the service delivered to the end-user which is the same for both digital and counterfactual approaches to delivery. In this way, the methodology abstracts away from the increased activity that may result from digitalisation (e.g. due to increased demand as the cost of the activity falls, or due to the reallocation of labour to other productive activities). This allows for a clearer assessment of the direct electricity implications of digitalisation, separate from its wider impact on economic growth. Whereas previous analyses have typically focused on carbon emissions or examined individual components in isolation, our methodology provides a comprehensive means of comparing energy use across the full delivery chain — spanning data centres, transmission networks, and end-user devices in the digital case. For the counterfactuals, the analysis captures electricity use associated with a wider variety of activities, ranging (for example) from Digital Terrestrial Television (in the case of video streaming) to office electricity consumption (in the case of AI-assisted use cases).

We have applied our methodology to ten use cases (discussed below).

2.2 Our ten selected use cases

We selected ten use cases that are associated with clear and well-defined physical alternatives, allowing for like-for-like electricity comparisons. The criteria for selecting these use cases were:

[2] See our previous study: DESNZ (2025) *"Impact of growth of data centres on energy consumption"* [online]

Europe Economics makes use of artificial intelligence to enhance the value of the work that it does for clients. Europe Economics may have generated some parts of the text in this document with OpenAI's language-generation model. Upon generating draft language, Europe Economics reviews, edits, and revises the text.

- **Driver of future data centre growth:** The use cases are representative of digital activities that are expected to drive future growth in data centre demand, rather than transitions that are already largely complete.
- **Clear Counterfactual:** Each has a well-defined alternative, allowing for direct, like-for-like electricity comparisons.
- **Data Availability:** There is sufficient data available for both the digital and counterfactual scenarios.
- **Analytical Tractability:** The selected use cases are feasible to analyse within the project’s scope, avoiding unnecessary complexity.
- **Good coverage of AI use cases:** At least half of the selected use cases involve AI-enabled services, reflecting the growing importance of AI as a driver of data centre electricity demand.

To investigate the electricity implications of data centre usage, this study focuses on ten use-cases:

- Video streaming versus TV broadcasting;
- Cloud gaming versus using a gaming console or gaming laptop;
- An online education platform versus a short in-person course;
- Cloud company back-up versus in-office server back-up;
- AI-assisted email drafting versus manual email drafting;
- AI-assisted code development versus manual code development;
- AI-assisted transcription versus manual transcription;
- AI-assisted copy-editing versus manual copy-editing;
- AI-assisted data extraction versus manual data extraction; and
- AI-assisted demand forecasting model development versus manual model development.

We briefly outline these use cases below.

Video streaming versus TV broadcasting

Video streaming was selected as a use case because it is a major and growing driver of data centre demand and represents an important category of digital content delivery. It has a clear and well-defined counterfactual in the form of traditional television broadcasting, allowing a like-for-like comparison for the same underlying activity of watching professionally produced video in the home. In addition, there is a substantial body of existing evidence on data transmission, data centre processing and end-user device electricity use for both pathways, which supports robust modelling within the project’s scope.

Cloud gaming versus using a gaming console or gaming laptop

The rapid growth of cloud gaming platforms has shifted game processing from user-owned hardware to centralised data centre infrastructure. This use case compares cloud-based gameplay with the physical alternative of gaming on a user's own console or gaming laptop, holding the underlying activity constant. It was selected because it represents a highly compute-intensive workload that is expected to contribute to future data centre electricity demand, has a clear and well-defined counterfactual, and can be analysed using available evidence on data centre processing, network transmission and end-user device electricity use within a tractable modelling framework.

An online education platform versus a short in-person course

Online education platforms were selected as a use case because short courses and professional training are increasingly delivered in both online and in-person formats, making it possible to define a clear, like-for-like comparison for the same educational activity. Electricity use in the physical counterfactual is driven by a small and well-defined set of components—principally classroom electricity use and participant travel—while the digital pathway is dominated by data transmission, data-centre processing and end-user device use. This makes the use case well suited to isolating how substituting online delivery for in-person attendance affects overall electricity consumption.

Cloud company back-up versus in-office server back-up

Cloud back-up services were selected as a use case because they represent a growing component of data centre demand associated with storage and data resilience. They have a clear counterfactual in the form of locally managed server back-up, allowing a direct comparison of alternative ways of performing the same function. In addition, electricity use in both the digital and counterfactual pathways is driven by a limited set of well-defined processes, principally storage, data transfer and server operation, supported by available evidence, making the use case feasible to analyse within the project's scope.

AI-assisted email drafting versus manual email drafting

The increasing use of generative AI tools to support everyday office work has made AI-assisted email drafting an emerging application of data centre-based services. This use case compares AI-assisted and manual email drafting for the same underlying task, providing a clear and well-defined counterfactual for like-for-like electricity comparison. It was selected because it represents an AI-enabled service with potential for widespread adoption and future growth in data centre demand, and because sufficient evidence exists on inference workloads, data transmission and office-based electricity use to support analysis within a tractable modelling framework..

AI-assisted code development versus manual code development

The growing integration of generative AI tools into software development workflows has changed how coding tasks are carried out in practice. This use case compares AI-assisted and manual approaches to code development for the same underlying task, allowing a like-for-like comparison of electricity use. It was included because AI-assisted coding is an emerging, compute-intensive application with clear implications for future data centre demand, and because the underlying processes on both the digital and counterfactual sides can be characterised using available evidence in a way that is feasible to analyse within the project's scope.

AI-assisted transcription versus manual transcription

The increasing integration of AI-based transcription tools into everyday workplace software has made transcription a scalable AI-enabled activity. This use case compares AI-assisted and manual transcription of the same recorded speech, providing a clear and well-defined basis for like-for-like electricity comparison. It was included because transcription is an AI application with broad and growing uptake, clear implications for future data centre electricity demand, and sufficient evidence on inference workloads, processing time and human transcription effort.

AI-assisted copy-editing versus manual copy-editing

The growing use of AI-based writing assistance tools in professional settings has changed how proofreading and copy-editing tasks are carried out. This use case compares AI-assisted and manual copy-editing of the same written text, providing a clear basis for like-for-like electricity comparison. It was selected because AI-assisted copy-editing represents a scalable AI-enabled service with potential for widespread adoption, a clearly defined human counterfactual, and sufficient evidence on inference activity and office-based work patterns to support analysis within a tractable modelling framework**.**

AI-assisted data extraction versus manual data extraction

The adoption of AI-based tools for processing large volumes of documents has accelerated across both the public and private sectors, particularly for tasks that involve extracting structured information from unstructured sources. This use case compares AI-assisted and manual data extraction for the same underlying task, enabling a like-for-like comparison of electricity use. It was included because AI-assisted data extraction is an emerging, scalable application with clear relevance for future data centre demand, a well-defined human counterfactual, and sufficient evidence on processing requirements and manual effort to allow analysis within the project's scope.

AI-assisted demand forecasting model development versus manual model development

The use of automated machine learning (AutoML) tools to support the development of statistical forecasting models is becoming increasingly common in data science workflows. This use case compares an AI-assisted approach to developing a short-term electricity demand forecasting model with a traditional manual approach in which analysts design, test and tune models themselves. It is an instructive use case to consider because AutoML represents a growing application of cloud-based and data centre computing in analytical work, particularly for tasks involving model selection, hyperparameter tuning and validation across many candidate models. The counterfactual scenario is a manual modelling workflow in which analysts iteratively develop and evaluate statistical models using conventional analytical tools.

2.3 Representative scenario approach

The study adopts a representative scenario approach as the foundation for estimating and comparing electricity consumption across digital services and their counterfactuals. This method involves defining a typical, well-specified instance of the activity for each use case, based on plausible patterns of use. The purpose of this approach is to ensure comparability across digital and counterfactual approaches and thus to provide a meaningful basis for quantifying the impact of digitalisation on electricity use.

Rather than relying on aggregated or sector-level data, each scenario reflects the actions of a single user engaging in a specific activity under standardised conditions. Parameters such as file sizes, duration of use, number of interactions, and choice of device are explicitly defined and held consistent between digital and counterfactual approaches. This ensures that both delivery models are assessed in relation to the same underlying activity and that electricity use is not overstated or understated due to differences in modelling scope or assumptions.

Our study does not assess the effectiveness or quality of digital approaches (such as AI-enabled services) relative to their counterfactuals. It also assumes no wider behavioural or economic effects from digitalisation and AI. In particular, the analysis does not account for potential economic growth, such as increased activity enabled by AI or the reallocation of time to other tasks. The results should therefore be interpreted as a like-for-like comparison of electricity use for a fixed level of activity.

2.4 Structure of report

The report is structured as follows:

- Section 3 outlines our methodology for estimating electricity consumption in digital use cases, with consideration given to both AI and non-AI cases. It includes a high-level modelling framework, generic input assumptions applicable across use cases, and specific assumptions for the ten use cases.
- Section 4 sets out our methodology for estimating electricity consumption in the counterfactuals. It distinguishes between counterfactuals involving office-based services (AI-related cases) and other bespoke counterfactuals, and presents the relevant input assumptions for each.

- Section 5 presents our results, comparing electricity consumption across digital and counterfactual approaches for each of the ten use cases. Results are presented for each use case, followed by a comparative assessment.
- Section 6 provides our conclusions.

3 Methodology for Modelling Electricity Consumption in Digital Use Cases

This chapter outlines the methodology used to model electricity consumption for the digital use cases. It begins by presenting the overall structure of the analytical approach, followed by a discussion of the generic input assumptions underpinning the key stages of electricity use. The section then sets out the specific assumptions applied to each of the ten digital use cases.

3.1 Overview of the generic modelling framework

3.1.1 Use cases not involving AI

This section outlines the approach used to estimate electricity consumption for use cases in which a service is delivered through digital means and which does not use AI. The framework captures three key components of electricity demand:

- the electricity consumed within data centres, covering both IT and supporting non-IT infrastructure;
- the electricity required to transmit data over internet networks; and
- the electricity consumed by the end-user’s device when accessing the service.

Data Centre Electricity Consumption

Electricity consumption within the data centre is estimated by **identifying the relevant IT processes involved in delivering the service**, such as uploading data, storage, CPU processing and GPU processing, and multiplying the volume of each process (e.g. gigabytes, CPU-hours) by assumed electricity intensities. These intensities represent the electricity required to complete one unit of each process.

Where applicable, the framework distinguishes between **shared and user-specific processes**. Shared processes (e.g. AI training) are used by multiple users, and only a proportion of the total electricity is attributed to the representative user. This proportion is calculated based on the assumed user base. In contrast, user-specific processes (e.g. downloading a file) are allocated fully to the representative user.

Once the IT electric energy is calculated, it is uplifted using standard **Power Usage Effectiveness (PUE)** factors to reflect the additional non-IT electric energy required to operate the data centre. This includes electricity for cooling, power distribution, lighting, and other building services. The result is a combined estimate of IT and non-IT data centre energy attributable to the representative user.

Internet Transmission

Electricity used in internet transmission is calculated by applying **fixed electricity intensity assumptions (kWh per gigabyte transmitted)** to the volume of data transferred in the representative scenario. As with data centre electricity, transmission processes are categorised as either shared or specific. For instance, data uploads may be shared across users in some use cases, while downloads are typically treated as user-specific. Where transmission is shared, only a proportional share of electricity is attributed to the representative user.

End-User Device

The final component captures the electricity used by the end-user device when interacting with the digital service. The device type (e.g. desktop computer, television, or gaming console) and usage time are defined in the representative scenario. Electricity use is then estimated by multiplying the device's power draw by the number of hours (or minutes) of use.

Total Electricity Consumption

The total electricity consumption for the digital use case is calculated by summing the three components: data centre electricity consumption (IT and non-IT), internet transmission electricity consumption , and end-user device electricity consumption. This approach enables a consistent and transparent comparison with electricity demands under the counterfactual (discussed in the next section of this report).

The figure below provides an overview of the modelling framework, illustrating how generic inputs and case-specific parameters flow through to produce electricity estimates for each digital use case.

Figure 3.1: Modelling framework for estimating electricity consumption in non-AI digital use cases

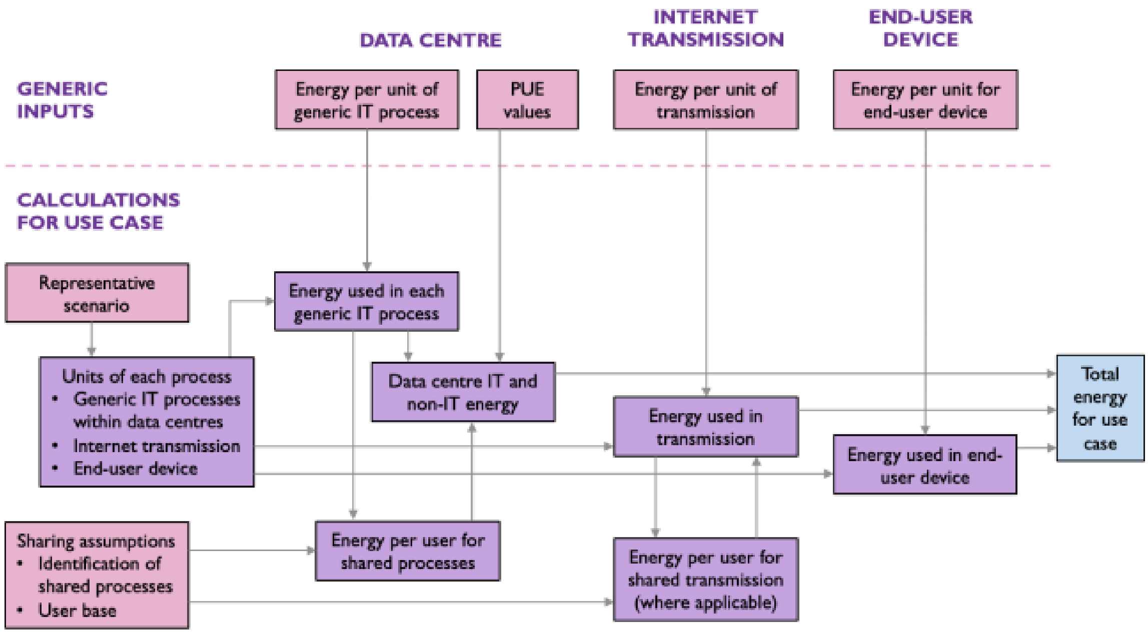

In section 3.2, we set out the input assumptions that we used in the modelling. In many cases, we identified a source for input assumptions (e.g. from the literature).

In cases where we did not find a source, we used AI-assisted imputation to generate plausible values for the model. The AI model was prompted to produce a low, central and high estimate together with the reasoning and assumptions underlying the range.

Before applying AI imputation, we first validated it using a small number of metrics for which reliable benchmark values exist. AI estimates were generated using different prompting styles and model configurations and compared with published evidence. This exercise helped identify which combinations of AI model configuration and prompting style produced estimates most consistent with known benchmarks.

Following this validation, the selected model and prompting approaches were then used to impute parameters for which no suitable published estimates could be identified. For each such parameter, estimates were generated using multiple prompting styles in separate chat sessions, and the resulting values were compared before selecting a representative estimate for use in the model.

3.1.2 AI-assisted use cases

This section sets out the modelling approach for digital use cases involving artificial intelligence. The framework follows the same structure as that used for non-AI digital use cases (Section 3.1.1), with the addition of AI-specific processing within data centres and consideration of the office electricity use of the worker who uses AI.

In these use cases, we focus on AI-assisted workflows in which a human remains involved in carrying out, reviewing or finalising the task, rather than fully automated AI-only delivery. This is intended to reflect how such tools are commonly used in practice. The methodology we have developed in this project for modelling AI-assisted workflows could be applied more widely to other examples of AI-assisted work.

As in the non-AI cases, electricity consumption is estimated across three main components: data centre operations, internet transmission, and user-side activity.

Data Centre Electricity Consumption

Data centre electricity consumption is estimated using the same approach as in non-AI cases, based on the use of IT processes such as data upload and storage. However, for AI-enabled services, this is extended to include the electricity associated with AI model training and AI inference.

AI training is treated as a shared process, as the trained model is typically used by many people. Total training electricity is therefore allocated across the assumed lifetime use of the model, with only a proportional share attributed to the representative user. AI inference is treated as user-specific, as it is directly associated with individual interactions with the AI system.

As above, the electricity required for each IT process is calculated by applying assumed electricity intensities to the relevant activity volumes. IT electricity consumption is then uplifted using PUE factors to account for non-IT electricity use, including cooling and other supporting infrastructure.

Internet Transmission

Internet transmission electricity is calculated in the same way as for non-AI digital use cases, by applying fixed electricity intensity assumptions to the volume of data transferred in the representative scenario. Where transmission is shared across users, electricity is allocated proportionally; otherwise, it is attributed fully to the representative user.

Office-Based Electricity Consumption

For AI use cases, the final component captures electricity consumption associated with the office environment (rather than just the end-user device). This reflects the fact that the AI-assisted activity is assumed to involve a human worker carrying out the task in an office setting.

Office electricity consumption is estimated using a bottom-up approach:

- We apply a **standard electricity energy use intensity (EUI)** for commercial office buildings, measured in kilowatt-hours per square metre per year (kWh/m²/year).
- This is multiplied by an assumption for **average office space per worker** (in m²), yielding an estimate of **annual electricity use per office worker**.

The EUI figures used are all-inclusive: they represent the total annual metered electricity demand of the premises and therefore encompass all forms of electricity use within the building. This includes electricity for IT and computing equipment, lighting, heating (for electrically-heated offices), cooling, ventilation, and other building services. Consequently, they already incorporate the electricity used by the office worker's device and supporting infrastructure. Our high figure assumes an all-electric office, whereas our low figure assumes that the office also uses another fuel (e.g. for heating).

We then scale this annual figure based on the **time required to perform the specific task**. This provides an estimate of the **office electricity consumed during the performance of the activity**.

Total Electricity Consumption

Total electricity consumption for AI digital use cases is calculated by summing data centre electricity energy (including AI training and inference), internet transmission electricity energy, and office-based electricity consumption.

The figure below illustrates the modelling framework for AI digital use cases, highlighting the points at which AI-specific processes enter the calculation.

Figure 3.2: Modelling framework for estimating electricity consumption in AI-enabled digital use cases

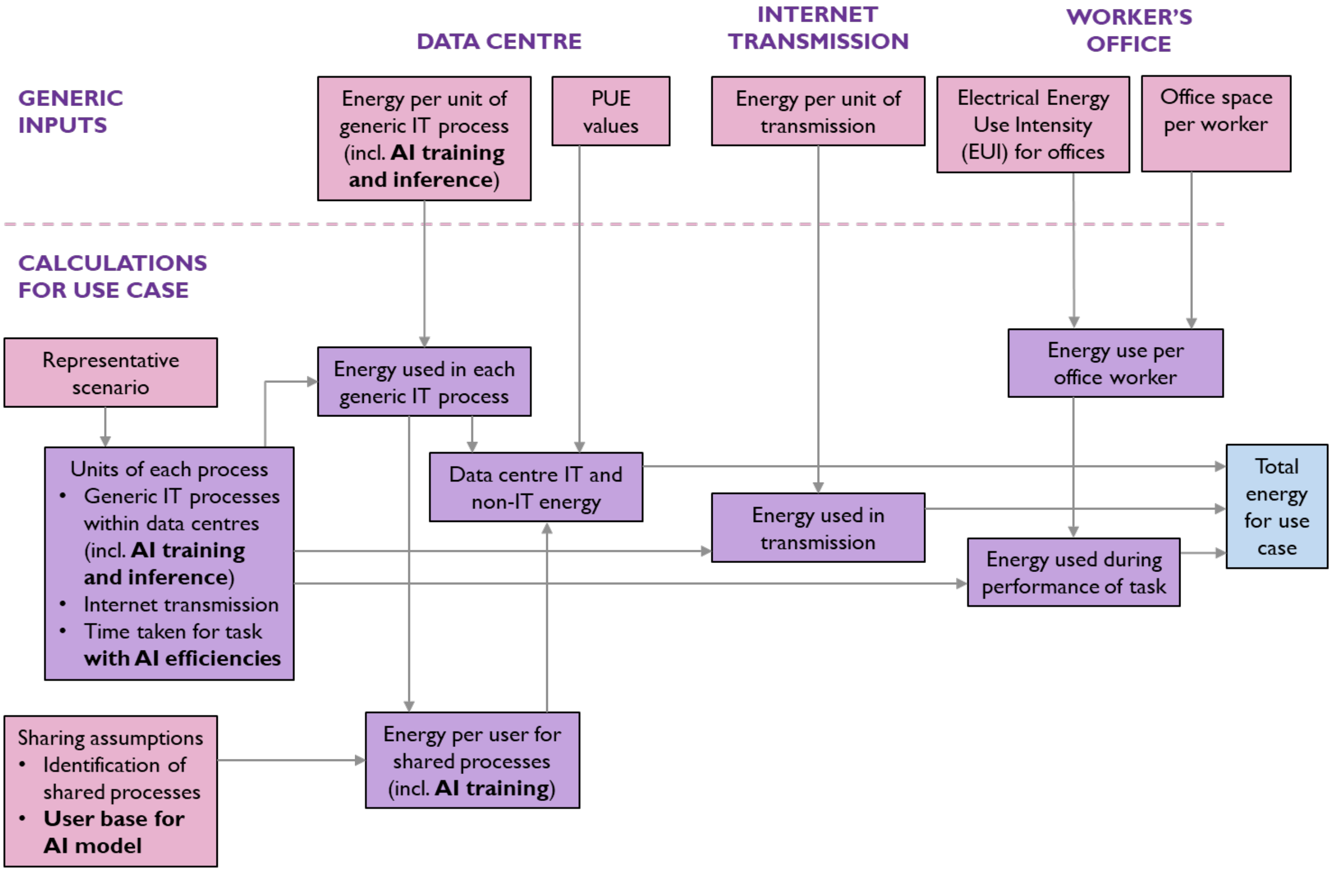

As for the non-AI cases, the following section sets out the input assumptions used in the modelling. Where published sources were available, these were used directly. Where data gaps remained, artificial intelligence was used to generate plausible values, which were subsequently reviewed and sense-checked.

3.2 Generic input assumptions

In this section, we set out the input assumptions underpinning the four main areas of electricity consumption in our digital use case modelling**: data centre operations, internet transmission, end-user devices, and office electricity consumption.**

3.2.1 Data centre

This subsection sets out the electricity intensity assumptions for the main data centre processes used in the modelling. The processes, units and values have been updated from the previous study and are applied consistently across all relevant digital use cases.

Uploading to data centres

Uploading refers to the stage at which data enters the data centre environment before it is stored or used for further processing. The process includes the internal networking activity required to receive and route incoming data within the data centre, such as network interfacing, routing and switching between servers and internal network equipment. It therefore captures the movement and handling of data within the data centre immediately after arrival. Electricity used for external transmission over the internet is excluded, as this is accounted for separately, and long-term storage electricity is also excluded.

We assume a **medium electricity intensity of 0.000050 kWh per GB uploaded**, with **low and high values of 0.000005 kWh per GB** and **0.000500 kWh per GB**, respectively. As no robust published estimates were identified for this parameter, these values were generated using our AI imputation process. The range reflects uncertainty in data centre network architecture, equipment efficiency and utilisation.

Video Processing

Video processing covers the transformation of video content into formats suitable for storage, delivery and playback, including tasks such as encoding, compression and format conversion. These processes are computationally intensive and typically rely on specialised hardware, such as GPUs or dedicated accelerators. As in the previous report, video processing is included only for use cases where video content is a material component of the digital service.

We assume a medium electricity intensity of **0.4 kWh per GB processed**, with low and high values of **0.1 kWh per GB** and **1 kWh per GB**, respectively. These assumptions are consistent with those used in our previous report.

Long-term storage

Long-term storage refers to the enduring storage of digital content that may be accessed repeatedly over time or retained for archival purposes. As in the previous report, this process captures the electricity required to keep storage systems powered, rather than active processing. It is less compute-intensive than other data centre activities but can contribute materially to overall electricity use due to the scale of stored data.

In the earlier study, long-term storage electricity intensities, electricity intensities were based on the specifications reported in the 2024 U.S. Data Centre Energy Usage Report[3] , yielding an annual electricity usage of 0.0025 kWh per GB for HDD-based storage (low), 0.008 kWh per GB for SSD-based storage (high), and 0.005 kWh per GB as a medium estimate.

In practice, however, enterprise storage systems are not operated at full capacity. To maintain performance headroom and resilience, efficiently run data centres typically operate storage below maximum capacity, meaning that the effective usable storage is lower than installed

[3] Berkely Lab (2024) "2024 U.S Data Centre Energy Usage Report" [online]

capacity. To reflect this, we adjust the previously used storage electricity intensities to account for typical capacity utilisation.[4] We assume a utilisation rate of 80 per cent in the medium scenario. As higher utilisation will lead to lower per unit electricity consumption, we assume utilisation of 85 per cent in the low scenario and 75 per cent in the high scenario.

Lower effective capacity increases the electricity intensity per effective GB stored, as the fixed electricity consumption of storage hardware is spread over fewer usable gigabytes. Applying these utilisation adjustments to the storage coefficients used in the earlier study yields updated annual electricity intensities of **0.0029 kWh per GB (low), 0.0063 kWh per GB (central), and 0.0107 kWh per GB (high)**, which are used in the analysis that follows.

RAM storage

RAM storage supports short-term, high-speed data retention during active computation, such as buffering, caching and in-memory processing. As in the previous report, this process captures the electricity associated with keeping server memory powered during periods of active use, rather than long-term data storage.[5]

Electricity intensity assumptions are based on manufacturer-reported power draw for modern DRAM technologies commonly used by hyperscale providers. Taking the average of reported values, we assume an electricity intensity of **0.000392 kWh per GB-hour**, with low and high estimates of **0.000375 kWh per GB-hour** and **0.000408 kWh per GB-hour**, respectively.

General processing

General processing covers routine computational tasks that are not specific to video processing or AI workloads, such as metadata handling, indexing, system management and background services. These activities are typically carried out by general-purpose CPUs.

We assume a medium electricity intensity of **0.025 kWh per CPU-hour**, with low and high values of **0.00375 kWh per CPU-hour** and **0.225 kWh per CPU-hour**, respectively, as no robust published estimates were identified for this parameter, these figures were generated using our AI imputation process. The range reflects variation in processor types and workload intensity.

GPU processing

GPU processing captures computational workloads executed on graphical processing units or similar accelerators within data centres. These workloads typically arise where tasks are highly parallelisable or computationally intensive, and are common in applications such as graphics rendering, simulation, and other high-throughput processing tasks. Electricity use is parameterised in terms of GPU-hours, reflecting the duration for which accelerator hardware is

[4] Paessler (2025) “Data center capacity planning - A guide for IT infrastructure teams” [online]
[5] Carbon Cloud Footprint “Methodology” [online]

actively used. To convert GPU-hours into electricity consumption, we apply assumed electricity intensities based on reported power draw for data centre GPUs.

We use existing industry reports[6] to extract the electricity estimates for our generic inputs, and arrive at a medium estimate of **0.125 kWh per GPU-hour**, with low and high values of **0.1 kWh per GPU-hour** and **0.15 kWh per GPU-hour**, respectively.

AI training

AI training refers to the development of machine-learning models through repeated computation over large datasets. This process is highly compute-intensive and typically relies on specialised accelerators, such as GPUs or TPUs, operating over sustained periods. In this study, AI training is included in addition to the standard data centre processes set out above, and is treated as a shared activity across users of the AI model.

We note that large-scale AI training may not typically be undertaken in the UK due to relatively high electricity costs and is more commonly carried out abroad. Including training electricity within the scope of this analysis may therefore overstate domestic electricity consumption, but ensures that the full electricity requirements of AI systems are captured. We present sensitivity results in Appendix B which show the effect of excluding AI training electricity.

Electricity requirements for AI training are expressed in terms of electricity per floating-point operation (FLOP), following Shankar and Reuther (2022).[7] Based on their reported ranges, we adopt electricity intensity assumptions of **1×10^{-6} μJ per FLOP** (low), **5.5×10^{-6} μJ per FLOP** (medium), and **1×10^{-5} μJ per FLOP** (high).

AI inference

AI inference refers to the use of a trained AI model to generate outputs from new inputs, such as producing text, extracting information, or generating predictions. While less electricity-intensive than training, inference can occur at scale and therefore contribute materially to overall data centre electricity use. AI inference is treated as a user-specific process.

For consistency with AI training, electricity use for inference is expressed in terms of electricity per floating-point operation. We apply the same electricity intensity assumptions as for training, based on Shankar and Reuther (2024): **1×10^{-6} μJ per FLOP** (low), **5.5×10^{-6} μJ per FLOP** (medium), and **1×10^{-5} μJ per FLOP** (high).[8]

[6] Nvidia (2021) [link]
[7] Shankar, S., & Reuther, A. (2022, September). Trends in energy estimates for computing in ai/machine learning accelerators, supercomputers, and compute-intensive applications. In 2022 IEEE High Performance Extreme Computing Conference (HPEC) (pp. 1-8). IEEE [online]
[8] These figures are supported by our calculations from a separate study on AI model inference, which produced a comparable upper-bound estimate. The analysis considered a typical ChatGPT (GPT-4o) response of 500 tokens. approximately 375 words, which was estimated to require 100 trillion FLOPs, based on 100 billion active parameters and an assumed 2 FLOPs per parameter per token. The energy consumption to generate this response was estimated at 0.3 watt-hours, equivalent to 1,080 joules. Dividing energy by FLOPs yields an energy

Downloading and streaming

Downloading and streaming refer to the retrieval of stored content from data centre systems for delivery to users. This includes the internal networking activity required to access and route data within the data centre, such as network interfacing, routing and switching between servers and internal network equipment. As with uploading, this parameter captures only intra-data-centre activity and excludes electricity used in external transmission networks, which is accounted for separately, as well as long-term storage electricity.

We assume the same electricity intensity values as for uploading, with a **medium estimate of 0.000050 kWh per GB downloaded or streamed**, and **low and high values of 0.000005 kWh per GB** and **0.000500 kWh per GB**, respectively.

Table 3.1: Electricity per unit of generic IT process[9]

Data process	Units	Low	Medium	High
Uploading to data centres	kWh per GB uploaded	0.000005	0.00005	0.0005
Video processing	kWh per GB processed	0.1	0.4	1
Long-term storage	kWh per GB stored per year	0.0029	0.0063	0.0107
RAM storage	kWh per GB stored in RAM per hour	0.000375	0.000392	0.000408
General processing	kWh per CPU-hour	0.05	0.1	0.15
GPU processing	kWh per GPU-hour	0.1	0.125	0.15
AI training	µJ per FLOP	0.000001	0.0000055	0.00001
AI inference	µJ per FLOP	0.000001	0.0000055	0.00001
Downloading/ Streaming	kWh per GB downloaded/streamed	0.000005	0.00005	0.0005

Source: Europe Economics analysis.

Power Usage Effectiveness (PUE)

intensity of 1.08×10^{-11} joules per FLOP, or 1.08×10^{-5} µJ/FLOP. This aligns closely with the upper range of our estimates. It is important to note that this calculation is based on inference using Nvidia H100 GPUs with ~10 per cent utilisation and ~70 per cent of peak power draw, which may be higher than typical inference scenarios. However, as ChatGPT is a high-end deployment of AI inference, we consider it appropriate to benchmark the high end of our range using this approach. [online]

[9] Electricity intensity assumptions have been updated relative to our previous report for uploading, long-term storage and downloading/streaming. For uploading and downloading/streaming, the update reflects clearer system boundaries in our AI imputation process, isolating internal processes within data centres and excluding external transmission to avoid double-counting. For long-term storage, the update reflects the inclusion of storage utilisation assumptions. If our updated electricity intensity values for these data centre processes were applied to the use cases in our previous report, the only use case for which the results shown in Table 7.1 of our previous report would change at the reported number of decimal places would be video streaming. In the case of video streaming, these updated electricity intensity values (keeping all other assumptions constant) would imply low, medium and high estimates of energy consumption under the digital approach of 0.50, 0.77 and 1.26 kWh.

Power Usage Effectiveness (PUE) is used to account for non-IT electricity consumption within data centres, such as for cooling, power distribution and other supporting infrastructure. As in the previous Europe Economics study, PUE values are applied to uplift estimated IT electricity consumption to derive total data centre electricity use. The selected PUE values were informed by stakeholder interviews carried out in our previous project with data centre operators and industry experts.

For each use case and scenario, IT electricity energy is multiplied by the relevant PUE value. This ensures that electricity consumption reflects both computational activity and the supporting services required to operate the data centre. Consistent with the earlier report, we apply the following PUE values across the low, medium and high scenarios:

- **Low value of 1.1**: Reflecting best-in-class performance typically achieved by newly built hyperscale data centres. These centres are often owned and operated by major cloud providers (e.g. AWS, Google, Microsoft) and benefit from extensive in-house optimisation and highly efficient cooling systems.
- **Medium value of 1.3**: Represents the industry's current convergence point, consistent with recent trends and targets. This value reflects the typical performance of new or recently upgraded data centres that comply with voluntary efficiency standards such as the European Code of Conduct for Data Centre Energy Efficiency or the Carbon Neutral Data Centre Pact**.**
- **High value of 1.5**: Corresponds to less efficient legacy facilities, particularly smaller co-location centres that may not yet have undergone major infrastructure upgrades. These centres often serve a variety of clients with diverse operational needs, making holistic optimisation more difficult.

Table 3.2: Ranges of PUE for data centres

	Low	Medium	High
PUE for data centres	1.1	1.3	1.5

Source: Europe Economics

3.2.2 Internet transmission

The next data process is internet transmission, and hence we consider the electricity used in moving data from the end-user to the data centre (and vice versa) over the internet. In our modelling, we adopt transmission electricity intensity estimates from the Carbon Trust's report on 'The Carbon Impact of Video Streaming', which provides a well-documented approach across fixed and mobile networks.[10] These figures are applied in combination with use-case-specific estimates of the data volumes transferred.

[10] Carbon Trust (2021) "Carbon impact of video streaming" [online]

The Carbon Trust estimates comprise the total electricity use associated with both the core network (which includes national and international backbone infrastructure) and the access network (the local infrastructure connecting users to the internet, such as fibre lines or mobile base stations). These two components are combined into a single aggregate figure for each network type.

We apply an electricity intensity of **0.0012 kWh per GB for fixed networks**. This figure is ultimately derived from Aslan et al. (2018), which reviewed a broad range of academic studies and developed a regression model to estimate trends in electricity intensity between 2000 and 2015.[11] The Carbon Trust extrapolated these trends forward to 2020 to arrive at the figure used in their analysis, resulting in an electricity intensity estimate of 0.0065 kWh per GB. For the purposes of this study, we have further extrapolated these trends to 2025, applying the same rate of annual improvement.

While our medium values for electricity intensity per GB are based on Carbon Trust estimates, for sensitivity analysis we assume a **± 20 per cent range** to generate low and high values. Our figures are summarised in the table below.

Table 3.3: Electricity consumption of transmission over fixed networks (kWh/GB)

Component	Low	Medium	High
Core + Access	0.00096	0.0012	0.00144
Home router	0.02	0.025	0.03
Total – Fixed Network	**0.02096**	**0.0262**	**0.03144**

Source: Carbon Trust (2021), Europe Economics analysis

3.2.3 End-user device

Another component of electricity consumption in digital use cases relates to **end-user devices**, which are used for rendering, displaying, or interacting with digital content depending on the nature of the service. These include both primary screen devices, such as TVs, e-readers, laptops, and smartphones, and supporting peripherals, such as set-top boxes. We estimate the electricity use of these devices based on average hourly power draw, using values that represent typical performance for commonly used models.

We used the Carbon Trust report as a starting point for estimating the electricity consumption of end-user devices. The figures in the Carbon Trust report were derived through a review of publicly available technical specifications and industry benchmarks, combining manufacturer data with secondary analysis. Where appropriate, we have reviewed these figures against more recent data for UK consumer devices to ensure they remain representative. Where

[11] Aslan, J., Mayers, K., Koomey, J. G., & France, C. (2018). Electricity intensity of internet data transmission: Untangling the estimates. *Journal of industrial ecology*, *22*(4), 785-798. [online]

updated figures indicate a material difference, we apply revised estimates. Our final assumptions are as follows:

- **Desktop computers** typically consume more power than laptops. Our low estimate comes from a lower-end desktop which draws 200W, while the high estimate comes from a higher-end desktop which draws 300W. The medium estimate comes from the midpoint which is **250W.**[12]
- **Gaming PCs** typically consume more power than desktop computers during active gameplay. We assume a mid-range gaming PC, with 16-32GB RAM, which tends to draw between 300W and 500W during active use. We adopt a medium estimate of **400W.**[13]
- **Smart TVs** often have higher power draw than standard TVs of a similar size. Our low estimate comes from a 32" Smart TV which draws 50W while our high estimate comes from a 65" Smart TV which could draw up to 150W.[14] The midpoint of these estimates, **100W**, is adopted as our medium estimate.
- **Gaming consoles** consume different levels of power depending on the intensity of gameplay and the associated processing load. Based on manufacturer disclosures for current-generation consoles, power draw during active use typically ranges from around 110W to 260W.[15] We distinguish between two operating modes: *normal-intensity* use and *high-intensity* use.
    - For **normal-intensity gameplay**, we adopt a medium estimate of **0.13 kWh per hour**, with low and high values of 0.11 kWh per hour and 0.16 kWh per hour, respectively.
    - For **high-intensity gameplay**, we adopt a medium estimate of **0.22 kWh per hour**, with low and high values of 0.17 kWh per hour and 0.26 kWh per hour, respectively.

To reflect uncertainty and variability across device models and user behaviour, we include low and high around the medium estimates. These power draw values are then combined with representative usage durations (e.g. time spent streaming) to calculate total device-level electricity consumption for each use case.

Table 3.4: Electricity consumption for core devices (kWh/hour)

Device type	Low	Medium	High
Desktop computer	0.2	0.25	0.3
Laptop	0.05	0.075	0.1
TV	0.03	0.105	0.18

[12] EcoFlow (2025) "Laptop vs. Desktop: How Many Watts Does a Computer Use?" [online]
[13] SolarTechOnline (2025) "How much electricity does gaming PC use?" [online]
[14] Blue TTI Power (2024) "How Much Electricity does TV use" [online]
[15] Playstation (n.d.) "Ecodesign" [online]

Device type	Low	Medium	High
Smart TV	0.05	0.10	0.15
Set-top box	0.010	0.015	0.022
Gaming PC	0.30	0.40	0.50
Gaming console (high intensity)	0.17	0.22	0.26
Gaming console (normal intensity)	0.11	0.13	0.16

Source: Carbon Trust (2021), Europe Economics analysis

3.2.4 Office electricity consumption

For AI-enabled use cases where tasks are carried out by office-based workers, we estimate office electricity consumption by combining assumptions on the electrical energy use intensity (EUI) of office buildings (kWh per m² per year) with estimates of office floor space per worker. EUI values are taken from UK Government evidence on non-domestic building energy consumption.[16] The high figure for electricity use is based on an all-electric office, the low figure is based on an office which also uses another fuel, and the medium figure is based on the mid-point. Office space per worker is specified using representative figures, capturing variation between more space-efficient and more space-intensive office environments.

Annual electricity consumption per worker is calculated by multiplying electrical EUI by assumed floor space per worker. To express this on a per-task basis, annual values are converted into electricity consumption per working day, assuming electricity is only consumed on working days and that there are, on average, 253.25 working days per year after accounting for weekends and Bank Holidays.

The resulting values are used consistently across all AI-assisted use cases and their counterfactuals to ensure comparability.

Table 3.5: Input assumptions for office electricity consumed by typical worker

Component	Unit	Low	Medium	High
Electrical Energy Use Intensity (EUI) for office buildings	kWh/m^2/year	67	79.5	92
Office space per office worker	m^2	10	15	20
Electricity per office worker per year	kWh/year	670	1192.5	1840
Electricity per office worker per working day	kWh/working day	2.65	4.71	7.27

Source: Europe Economics analysis.

[16] See Annex 2 at DESNZ (2024) "Non-Domestic Building Stock in England and Wales" [online]

3.3 Specific input assumptions for our ten use cases

3.3.1 TV streaming

This digital use case involves a user streaming video content.

Network Transmission and End-User Configuration

We assume that streaming takes place over a **fixed-line network**. The viewing device is a **smart TV**, selected to reflect typical behaviour for watching longer-form content such as films or series. No additional hardware (e.g. set-top box, gaming console) is included in the scenario.

Process Attribution – Shared versus Specific

Most **upstream processes** in the data centre are assumed to be **shared across the user base**, including content uploading, video processing (e.g. encoding, compression), and long-term storage. These processes typically occur once per item of content and benefit from economies of scale as they support multiple users. By contrast, **RAM storage for caching and downloading/streaming are treated as user-specific**, reflecting the per-user nature of these operations.

A similar distinction is made for transmission: upload transmission is considered a shared process (as content is uploaded once), while download transmission is specific to the end-user.

User Base Assumptions for Shared Allocation

To allocate electricity use for shared processes, we define the user base in terms of total hours of video content viewed during the period of the representative scenario. As in our use case for video streaming in the previous report, for the medium case, we assume that the content is viewed 1,000,000 times. A larger user base results in a smaller share of upstream electricity being attributed to each viewing instance, and hence we assume a larger user base of 5,000,000 views in our low scenario. Conversely, we assume a smaller user base of 500,000 views in our high scenario.

RAM Caching During Playback

We assume that **10 per cent of the video file is cached in RAM during playback**, representing short-term memory buffering typically used in streaming platforms to ensure smooth delivery.

The tables below summarise our input for the video streaming digital use-case.

Table 3.6: Configuration assumptions for the TV streaming digital use case

	Assumption
Network type	Fixed-line broadband
End-user device	Smart Television
Use of additional devices?	No
Shared processes	Uploading to data centres, video processing, long-term storage, upload transmission
User-specific process	RAM storage, downloading/streaming, download transmission

Table 3.7: User base and RAM caching assumptions

	Unit	Low	Medium	High
User base for shared allocation (hours of video viewed)	Hours	5,000,000	1,000,000	500,000
RAM caching	Percentage	10	10	10

3.3.2 Cloud gaming

This use case involves a user playing a game on an online cloud platform.

Network Transmission and End-User Configuration

We assume that the cloud gaming platform is accessed using a smart TV connected via a fixed-line broadband connection and operated with a standard wireless gaming controller. Given the controller's very low power draw, its electricity consumption is excluded from the analysis.[17] No additional devices are used in this scenario.

Process Attribution – Shared versus Specific

On the data centre side, uploading a gaming file to the cloud platform is treated as a shared process. Similarly, long-term storage of the game at the data centre is treated as a shared process as well. By contrast, RAM storage, general processing, GPU processing and downloading/streaming are user-specific, reflecting individual demands for each user accessing the game.

We assume that upload transmission consists of two elements: upload transmission when the game is first uploaded to the data centre, and upload transmission (e.g. of game commands) when the game is played by the user. The first process is assumed to be shared, while the latter is user specific. By contrast, there is no similar split for download transmission, which is a user-specific process.

[17] Controllers typically draw around 2-4W during active use [online]

Ratio of upload to download

Since during a gaming session data is transferred both from the end-user to the data centre and from the data centre to the end-user, we assume that the consumed transmission electricity energy is divided between these two elements. In cloud gaming, the data flow is dominated by video streaming from the data centre to the user, while upstream traffic mainly consists of relatively small control inputs.[18] Reflecting this asymmetry, we assume that 90 per cent of transmission energy is associated with download traffic and the remaining 10 per cent with upload traffic.

User Base Assumptions for Shared Allocation

To allocate electricity use for shared processes, we define the user base in terms of total hours of cloud gaming during the period of the representative scenario. For the medium case, we assume a representative game is played for 130,000,000 hours.[19] A larger user base results in a smaller share of upstream electricity being attributed to each instance of gaming, and hence we assume 104,000,000 and 156,000,000 hours of cloud gaming in our high and low scenarios respectively.

RAM Caching During Playback

We assume that 10 per cent of the gaming file is cached in RAM during game time, representing short-term memory buffering used in cloud gaming video streaming to ensure smooth delivery.

The tables below summarise our input for the video streaming digital use-case.

Table 3.8: Configuration assumptions for the cloud gaming digital use case

	Assumption
Network type	Fixed-line broadband
End-user device	Smart TV
Use of additional devices?	No
Shared processes	Uploading to data centres, video processing, long-term storage, upload transmission (from uploading the game to data centre)
User-specific process	RAM storage, general processing, GPU processing, downloading/streaming, upload transmission (from playing the game), download transmission

Table 3.9: User base and RAM caching assumptions

[18] Zetexa (2025) “Data Consumption in eSport: How Much Do You Need” [online]
[19] USwitch (2025) “Online gaming statistics” [online] Ericsson (2024) “Cloud gaming report: Tracing the consumer journey” [online]

	Unit	Low	Medium	High
User base for shared allocation (hours of cloud gaming)	Hours	156,000,000	130,000,000	104,000,000
RAM caching	Percentage	10	10	10

	Unit	Share
User-specific share of download	Percentage	90
User-specific share of upload	Percentage	10

3.3.3 Online education platform

This use case reflects the completion of a short online course.

Network Transmission and End-User Configuration

We assume online education platform is accessed on a desktop computer using a fixed-line network broadband. No additional devices feature in this scenario.

Process Attribution – Shared versus Specific

On the data centre side, this digital use case features similar aspects to the video streaming example, as we assume an online course consists of a series of pre-recorded videos teaching the content of the course, uploaded to an education platform. Uploading a content file to the platform is treated as a shared process. Similarly, video processing and long-term storage of the game are shared processes as well. By contrast, RAM storage and downloading/streaming are user-specific, reflecting individual access by each user of the course.

A similar divide applies for internet transmission: upload transmission is considered a shared process (as content is uploaded once), while download transmission is specific to the end-user.

User Base Assumptions for Shared Allocation

To allocate electricity use for shared processes, we define the user base in terms of hours spent watching the course content during the period of our representative scenario. In our medium scenario, we assume the course content is viewed in total for 36,000 hours. To reflect lower electricity consumption per user in our low use-case scenario, we assume 180,000 hours of content viewing for low scenario. By contrast, we assume 28,800 hours of content viewing in high scenario.[20]

RAM Caching During Playback

[20] ClassCentral (2021) “A Decade of MOOCs: A Review of MOOC Stats and Trends in 2021” [online]

We assume that 10 per cent of the content video file is cached in RAM during playback, representing short-term memory buffering to ensure smooth delivery of the pre-recorded content.

The tables below summarise our input for the video streaming digital use-case.

Table 3.10: Configuration assumptions for the online education platform digital use case

	Assumption
Network type	Fixed-line broadband
End-user device	Desktop computer
Use of additional devices?	No
Shared processes	Uploading to data centres, video processing, long-term storage, upload transmission
User-specific process	RAM storage, downloading/streaming, download transmission

Table 3.11: User base and RAM caching assumptions

	Unit	Low	Medium	High
User base for shared allocation (hours of course viewing)	Hours	180,000	36,000	28,800
RAM caching	Percentage	10	10	10

3.3.4 Cloud-based back-up

This scenario reflects the creation of a back-up copy of company data stored in the cloud. The representative scenario assumes that a fixed volume of data is backed up by copying data from a primary cloud data centre to a secondary data centre located at a different site. The back-up copy is stored for one year.

This back-up activity is separate from, and additional to, any continual saving or storage of files in the cloud during normal business operations. It does not require any interaction with, or activity by, end-user devices.

Electricity consumption is estimated only for the processes required to transfer the data from the primary server to the secondary server and store the resulting back-up copy. These processes comprise retrieving the data from the primary data centre, transmitting it over the internet to a secondary data centre, and long-term storage of the back-up copy.

Network Transmission and System Configuration

We assume that data is transferred directly between two cloud data centres over fixed broadband. The transfer is initiated and managed at the data centre level and does not involve any end-user devices or home networking equipment such as routers. The representative scenario therefore reflects typical automated back-up practices for a small or medium-sized enterprise rather than the behaviour of an individual user.

For internet transmission, only transmission between the data centres is modelled (using the “Download transmission” process in our model). In this context, “downloading” refers to retrieving data from the primary server in order to transfer it to the secondary server, rather than downloading to an end-user device.

Process Attribution – Shared versus Specific

All processes are treated as **specific** to the representative back-up activity, reflecting the fact that the scenario relates to a single company undertaking a discrete back-up operation, with none of the associated electricity consumption shared with other users.

Back-Up Volume Assumption

The representative scenario is parameterised using a single input: the volume of data backed up. We assume that **100 GB** of data is backed up, which is intended to be illustrative of a small enterprise.

The table below summarises our inputs for the cloud back-up digital use case.

Table 3.12: Configuration assumptions for the cloud back-up digital use case

	Assumption
Network type	Fixed-line broadband (excluding home routers)
Data backed up per period	100 GB
End-user device	N/a
Shared processes	None
Specific process	Uploading to data centre, long-term storage, downloading/streaming, download transmission

3.3.5 Assumptions used across AI-assisted use cases

The AI-assisted use cases rely on several common assumptions regarding the user environment, network configuration and the treatment of processes within the data centre. These assumptions are applied consistently across the relevant scenarios to ensure comparability of results.

Assumptions common to all AI-assisted use cases

The following assumptions apply across all AI-assisted scenarios regardless of the specific AI model used. They define the assumed user environment and the treatment of shared and user-specific processes within the digital pathway.

Network and Office Configuration

We assume that the user accesses the AI model from a standard office environment over a fixed-line broadband network connection. Electricity consumed by the user's computer, lighting, heating or cooling, and other background office equipment is captured within office electricity consumption, rather than being modelled as a separate end-user device.

This office-based configuration is applied consistently across the relevant use cases to ensure comparability. Differences between scenarios therefore arise from the specific digital processes involved and the associated data centre activity.

Table 3.13: Network and office configuration assumptions AI-assisted use cases

Assumption	Description
Network type	Fixed-line broadband
Office environment	Standard office setting; all user-device electricity included in office electricity use
Shared processes	Long-term storage, RAM storage, AI model training
User-specific processes	Uploading prompts, AI inference, downloading AI-generated text
Internet transmission	Upload and download transmission treated as user-specific

Source: Europe Economics analysis.

Process Attribution - Shared versus Specific

Within the data centre, AI model training, long-term storage of the model, and RAM storage are treated as shared processes because they support a large user base rather than individual user interactions. Electricity associated with these processes is therefore allocated across users based on an assumed total inference workload over the lifetime of the AI model.

AI inference, uploading of prompt text, and downloading of AI-generated drafts are treated as user-specific processes because they occur separately for each interaction with the system. Internet transmission is also treated as user-specific for both uploads and downloads.

Assumptions common to AI use cases that use the same kind of large language model

The following assumptions apply to AI-assisted use cases that rely on the same large language model, specifically AI-assisted email drafting, data extraction, and copy-editing.

These assumptions are used to allocate electricity from shared processes across the total usage of the model.

User-base assumptions for shared allocation
For the purpose of allocating electricity from shared processes, the user base is defined in terms of the total number of tokens processed by the model over its assumed lifetime.

The assumed model lifetime reflects the fact that large language models are periodically retrained and replaced by updated versions. Earlier generations of models were released at multi-year intervals (for example GPT-3 in 2020 and GPT-4 in 2023), but the pace of development has accelerated, with new model versions and updates now appearing more frequently.[21] Evidence also suggests that the behaviour and performance of deployed models can change over relatively short periods as updates are introduced. On this basis, we assume a representative model lifetime of six months for allocating training electricity consumption across users.

Total usage over this period is derived from assumptions about the scale and size of interactions with the model. A daily prompt volume of 2.5 billion is used in the low scenario and reflects publicly reported usage for ChatGPT, which is assumed to represent the largest and most widely used general-purpose language model. [22] The medium and high scenarios assume daily prompt volumes equal to 50 per cent and 10 per cent of this figure respectively, representing models with lower interaction volumes. [23]

Average prompt–response pairs are assumed to contain between 300 and 1,500 tokens, capturing variation across use cases and interaction styles. Aggregating these interaction volumes over the assumed six-month model lifetime yields an estimated total usage of between **1.37×10^{14} and 6.85×10^{15} tokens**, with a medium estimate of **1.14×10^{15} tokens**.

For comparability with some use cases in the analysis (particularly email drafting and copy-editing), token counts are later converted to words using the approximation 1 token ≈ 0.75 words. [24] This conversion is used only for reporting and does not affect the underlying allocation of shared energy across model usage.

The low scenario therefore reflects very high aggregate usage and yields the lowest per-interaction allocation of shared energy, while the high scenario reflects more limited usage and results in a higher per-interaction allocation of shared energy.

[21] Times of AI (2026) “GPT Version Timeline: How AI Got Smarter Over the Years” [online]
[22] Techcrunch (2025) “ChatGPT users send 2.5 billion prompts a day” [online]
[23] Darwin (2025) “15 Most Used Generative AI Chatbots [July 2025 Rankings & Stats]” [online]
[24] Epoch AI. (2025). How much energy does ChatGPT use? Gradient Updates (online)

Table 3.14: Total model usage assumed over six-month model lifetime

Metric	Low	Medium	High
Total tokens processed	6.85×10^{15}	1.14×10^{15}	1.37×10^{14}
Equivalent words (1 token ≈ 0.75 words)	5.14×10^{15}	8.56×10^{14}	1.03×10^{14}

Source: Europe Economics analysis.

AI model size and compute assumptions

The AI-assisted use cases are assumed to rely on a frontier large language model. Model size at inference is represented using a range of parameter counts drawn from publicly known large models: 70 billion parameters as a credible lower bound for a strong chat-capable model, 175 billion parameters as a central reference scale, and 1.6 trillion parameters as a large upper benchmark.[25] Following a parameter-to-memory relationship,[26] model size is estimated by multiplying the number of parameters by the number of bytes required per parameter. Assuming 16-bit precision (approximately 2 bytes per parameter), this corresponds to approximate **model sizes of 140 GB, 350 GB and 3,200 GB** respectively.

Training compute is assumed to span three orders of magnitude, from approximately 3×10^{24} to 3×10^{26} FLOPs, reflecting the wide dispersion in training compute observed across observed large language models. We use a value of 3×10^{25} FLOPs in the medium scenario. These values are informed by published estimates for recent frontier models reported by Our World in Data and Epoch AI,[27] and are intended to capture both uncertainty in published figures and genuine variation across model size, architecture and generation.

Inference compute per token is assumed to range from 5×10^{10} to 2×10^{11} FLOPs per token. The upper bound is informed by Epoch AI[28] estimates indicating that a GPT-4-class query generating around 500 output tokens requires on the order of 10^{14} FLOPs, equivalent to approximately 2×10^{11} FLOPs per token. Lower values reflect more efficient inference configurations, including model optimisation and hardware utilisation improvements.

---

[25] Abacha, A. B., Yim, W. W., Fu, Y., Sun, Z., Yetisgen-Yildiz, M., Xia, F., & Lin, T. (2025, July). Medec: A benchmark for medical error detection and correction in clinical notes. In *Findings of the Association for Computational Linguistics: ACL 2025* (pp. 22539-22550). [online]
[26] Github (2025) “Train Big, Plan Smart - How to Calculate Memory and Estimate GPUs for LLMs” [online]
[27] Our World in Data. (2024). Training computation of notable artificial intelligence systems. (online)
[28] Epoch AI. (2025). How much energy does ChatGPT use? Gradient Updates (online)

Table 3.15: Assumptions for large language model size and compute

	Units	Low	Medium	High
Size of model	GB	140	350	3,200
FLOPs used in training of AI model	FLOP	3.0×10^{24}	3.0×10^{25}	3.0×10^{26}
Inference compute per token	FLOPs/token	5.0×10^{10}	1.0×10^{11}	2.0×10^{11}

Source: Europe Economics analysis.

In the following sections, we discuss assumptions which are specific to each individual AI-assisted use case.

3.3.6 AI-assisted email drafting

This use case reflects the use of a generative-AI email assistant to draft routine professional emails, compared with manual drafting carried out entirely by a human. In both pathways, the underlying activity is identical: producing a standard business email of comparable quality and length. The difference lies in how the text is generated. In the AI-assisted pathway, the user submits a short instruction-based prompt to an AI model, receives a draft response, and then reviews or lightly edits it. In the manual pathway, the user drafts the email without AI support.

The AI-assisted pathway introduces additional data centre activity, including AI inference and the shared energy associated with training the underlying model, while reducing the time the user spends drafting the email.

Representative Email-drafting Activity

The representative scenario assumes that a knowledge worker drafts 50 emails per week, consistent with typical estimates of 8–12 emails per working day for office-based roles. Each email is assumed to contain approximately 200 words of text, excluding attachments or long quoted threads. This results in a total of 10,000 words drafted per week.

In the AI-assisted pathway, each email is assumed to involve a single prompt–response interaction with the AI model. The AI-generated draft text is assumed to have a data size of 0.000005 GB (5 kB) per email. This figure reflects not only the raw text content, but also system and message-wrapper overheads (such as metadata and formatting) and slightly longer AI-generated drafts relative to the final edited email.

The user's prompt is assumed to be relatively short and instruction-based (for example, "Draft a polite response declining the meeting"), rather than consisting of a pre-written draft. We therefore assume that the uploaded prompt text is equivalent to 25 per cent of the length of the generated email text. This aligns with observed prompt patterns in email-drafting tools, where prompts are typically concise and directive.

Inference workload

Email drafting is assumed to generate approximately 16,250 tokens of inference per week, based on the assumed number of emails drafted and the average word count per email. Words are converted to tokens using the standard approximation 1 word ≈ 1.3 tokens.

Total inference compute is calculated by multiplying the number of tokens generated by the assumed inference compute per token, as reported in Table 3.15. Applying this assumption yields weekly inference requirements in the range of approximately 8×10^{14} to 3×10^{15} FLOPs.

Table 3.16: Input assumptions for digital scenario – AI-assisted email drafting

	Units	Low	Medium	High
Tokens used for emails per week	Tokens	16,250	16,250	16,250
FLOPs per week for inference	FLOP	8.1×10^{14}	1.6×10^{15}	3.3×10^{15}

Source: Europe Economics analysis.

Time savings from AI assistance

Evidence from recent studies suggests that generative AI tools reduce the time spent drafting emails. Observational field evidence[29] indicates reductions of around 12–17 per cent in time spent on email writing. Controlled experimental evidence[30] on writing tasks more broadly suggests larger reductions, of around 40 per cent. Industry surveys of users of AI email tools report self-assessed[31] time savings of 60–75 per cent, though these figures are likely to be subject to upward bias.

Reflecting this range, we adopt a conservative medium estimate of a 40 per cent time saving, with low and high scenarios of 60 per cent and 15 per cent respectively. These values are consistent with the evidence base while avoiding reliance on the most optimistic self-reported estimates.

Time to email and office electricity use

In the manual drafting pathway, the speed of human emailing is assumed to range from 300 to 800 words per hour, with a medium estimate of 500 words per hour. These values reflect the drafting of professional-quality business emails, including planning, structuring and revision.

Applying the assumed time savings yields higher effective writing speeds in the AI-assisted pathway. Office electricity consumption is assumed to scale linearly with the time spent drafting emails and is modelled using standard office electricity-intensity benchmarks, consistent with the approach used elsewhere in the report.

[29] Dillon et al. (2025, arXiv:2504.11436) – Shifting Work Patterns with Generative AI (online)
[30] Noy & Zhang (2023, Science) – Experimental evidence on the productivity effects of generative AI (online)
[31] GetReadyToSend (2024) report – Self-reported time savings for AI email writers (online)

Table 3.17: Time assumptions for email drafting

	Units	Low	Medium	High
Time saving with AI assistance				
Time saving using AI tools	%	60	40	15
Time to email (digital)				
Speed of human emailing	Words/hour	800	500	300
Speed of human emailing with AI efficiencies	Words/hour	2,000	833	353

Source: Europe Economics analysis.

3.3.7 AI-assisted code development

This use case reflects the use of a coding-specialised generative AI model to assist with the development of a small piece of software, compared with manual coding carried out entirely by a human developer. In both pathways, the underlying activity is identical: producing a working block of source code of comparable scope and quality. The difference lies in how the code is generated. In the AI-assisted pathway, the developer interacts iteratively with an AI coding model through a sequence of prompts, receiving suggested code which is then reviewed, adapted and integrated. In the manual pathway, all code is written directly by the developer without AI support.

The AI-assisted pathway introduces additional data centre activity, including AI inference and the shared energy associated with training the AI coding model, while reducing the time the developer spends writing and refining code.

Representative Coding Activity

The representative scenario assumes a small software development task comprising 500 lines of code.[32] This is consistent with published benchmarks for small software projects and discrete coding tasks. Average line length is assumed to be 40 characters per line.

This results in a total of 20,000 characters of generated source code. Source code is assumed to be ASCII-dominant text encoded using UTF-8, such that one character corresponds approximately to one byte. Converting bytes to gigabytes using binary units (1 GB = 1,073,741,824 bytes) yields a total generated code size of approximately 0.00002 GB.

[32] Loreley. Lines of Code (LoC). (online)

Each AI interaction is assumed to involve a relatively short prompt, such as a request to generate, modify or explain a block of code. The uploaded prompt text per interaction is therefore assumed to be 5 per cent of the size of the generated code. On this basis, the total uploaded prompt data amounts to approximately 0.000001 GB per coding task.

User-base Assumptions for Shared Allocation

For the purpose of allocating electricity associated with shared processes, particularly AI model training, the user base is defined in terms of the total volume of source-code characters generated by the deployed AI coding model over its assumed lifetime of one year. This reflects the longer useful life of coding-specialised models relative to general-purpose large language models.[33]

The medium scenario assumes that the model generates approximately 400 billion source-code characters over one year, with low and high scenarios of 800 billion and 200 billion characters respectively. These figures are AI-imputed and represent plausible aggregate usage across many developers and tasks, based on bottom-up assumptions about model adoption, frequency of coding tasks, and task size. Source code is assumed to be plain text, with approximately one byte per character.

The low scenario involves a larger user base and a lower per-task allocation of training energy, while conversely the high scenario reflects a more limited deployment of the model.

AI Model Size and Computational Requirements

The AI-assisted coding use case is assumed to rely on a coding-specialised large language model with approximately 34 billion parameters.[34] Model size at inference is assumed to range from 70 GB to 90 GB, with a medium estimate of 80 GB. This reflects storage of model parameters at FP16 precision (approximately 2 bytes per parameter), plus additional overhead for embeddings and runtime deployment.

Training compute is assumed to lie between 1×10^{21} and 3×10^{23} floating-point operations (FLOPs). These values are informed by published estimates for coding-specialised models compiled by Our World in Data using Epoch AI data.

Inference compute per token is assumed to range from 5×10^{10} to 2×10^{11} FLOPs per token. The upper bound is informed by Epoch AI (2025) estimates, while lower values reflect more efficient inference configurations and the reduced complexity of coding-specialised models.

Each coding task is assumed to involve approximately 5,250 tokens of inference, based on one token being equivalent to roughly four characters. This yields total inference requirements per task in the range of approximately 2.6×10^{14} to 1.05×10^{15} FLOPs.

[33] DeepMind. (2022). Competitive programming with AlphaCode. Google DeepMind. (online)
[34] Meta AI. (2023). Code Llama: Large language model for coding. Meta Platforms, Inc. (online)

3.18: Input assumptions for digital scenario - AI-assisted coding

	Units	Low	Medium	High
User base for shared processes				
User base for shared processes	Code characters	8.0×10^{11}	4.0×10^{11}	2.0×10^{11}
AI coding model				
Size of model	GB	70	80	90
FLOPs used in training of AI model	FLOP	1.0×10^{21}	1.0×10^{22}	3.0×10^{23}
Inference compute per token	FLOPs/token	5.0×10^{10}	1.0×10^{11}	2.0×10^{11}
Tokens in representative scenario	Tokens	5,250	5,250	5,250
FLOPs for inference	FLOP	2.63×10^{14}	5.25×10^{14}	1.05×10^{15}

Source: Europe Economics analysis.

Number of prompts

Coding is typically an iterative process when using AI assistance, involving multiple prompt–response interactions. Evidence from real GitHub pull requests and issue discussions that include ChatGPT conversations suggests that around four prompts per coding task is typical.[35]

We therefore assume an average of four prompts per task in the medium scenario, with low and high values of two and ten prompts respectively, reflecting variation in coding style, task complexity and developer familiarity with AI tools.

Time savings from AI assistance

Empirical evidence suggests that AI-assisted coding can substantially reduce development time. Published studies report time savings of around 30-40 per cent for debugging and repetitive coding tasks,[36] with controlled experiments reporting reductions of up to 55.8 per cent when using tools such as GitHub Copilot.[37]

---

[35] Al Awad, M. N., Ivanov, S., & Tikhonova, O. (2025). Pre-Filtering Code Suggestions using Developer Behavioral Telemetry to Optimize LLM-Assisted Programming. In *2025 40th IEEE/ACM International Conference on Automated Software Engineering Workshops (ASEW)* (pp. 113-120). IEEE. [online]

[36] Pandey, R., Singh, P., Wei, R., & Shankar, S. (2024). Transforming software development: Evaluating the efficiency and challenges of GitHub Copilot in real-world projects (online)

[37] Peng, S., Kalliamvakou, E., Cihon, P., & Demirer, M. (2023). The impact of AI on developer productivity: Evidence from GitHub Copilot (online)

Reflecting this evidence, we adopt a medium time-saving assumption of 40 per cent. As a large percentage time saving will lead to lower electricity consumption (and vice versa), we use a value of 55 per cent in the low scenario and 30 per cent in the high scenario.

Time to code and office electricity use

In the manual coding pathway, developers are assumed to write between 25 and 125 lines of code per day, with a medium estimate of 75 lines per day. Combined with the assumed average of 40 characters per line and a 7.5-hour productive workday, this corresponds to human coding speeds of approximately 133 to 667 characters per hour.

Applying the assumed time savings yields higher effective coding speeds in the AI-assisted pathway. Office electricity consumption is assumed to scale linearly with the time spent coding and is modelled using standard office electricity-intensity benchmarks.

Table 3.19: Time assumptions for coding

	Units	Low	Medium	High
Time saving with AI assistance				
Time saving using AI tools	%	55%	40%	30%
Time to code (digital)				
Lines written per day	Lines/day	125	75	25
Speed of human coder	Characters/hour	667	400	133
Speed of human coder with AI efficiencies	Characters/hour	1,481	667	190

Source: Europe Economics analysis.

3.3.8 AI-assisted copy-editing

This use case reflects the use of a generative AI model to review and refine an existing piece of professional text, compared with manual copy-editing carried out entirely by a human. In both pathways, the underlying activity is identical: improving a given text to a comparable standard of quality, clarity and readability. The difference lies in how the editing task is performed.

In the AI-assisted pathway, the user pastes the full relevant text into the AI system, which provides suggested edits and improvements. The user then reviews and accepts or rejects these suggestions and completes the final copy-editing. In the manual pathway, the user reviews and edits the same text without AI support.

The AI-assisted pathway introduces additional data centre activity, including AI inference and the shared electricity associated with training the underlying model, while reducing the time the user spends on copy-editing.

Representative copy-editing activity

The representative scenario assumes that a knowledge worker reviews and copy-edits a single piece of professional text containing 10,000 words. Based on an average sentence length of 15 words and 500 words per page, this corresponds to around 667 sentences or 20 pages of text to be reviewed.

In the AI-assisted pathway, the AI copy-editing tool supports the user by suggesting edits to the text, while the user remains responsible for reviewing, accepting or rejecting these suggestions and completing the copy-editing task. Using a page-to-data conversion of 0.000015 GB per page, the total data size of the uploaded and returned text is estimated at 0.00030 GB. This figure reflects the textual content itself and modest system overheads (such as formatting and metadata).

In the manual pathway, the same volume of text is reviewed and edited by the user without AI support. The comparison therefore isolates the impact of AI assistance on data centre activity and time spent copy-editing, holding the volume and nature of the text constant across both pathways.

Inference workload

For the representative copy-editing scenario, a document length of 10,000 words is assumed. Using the standard conversion 1 word ≈ 1.3 tokens, and accounting for both uploaded input text and downloaded AI-edited output, this corresponds to 26,000 tokens processed by the model.

Applying the per-token inference compute assumptions reported in Table 3.15 yields total inference requirements of approximately 1.3×10^{15}, 2.6×10^{15}, and 5.2×10^{15} FLOPs in the low, medium, and high scenarios respectively..

Table 3.20: Input assumptions for digital scenario – AI-assisted copy-editing

	Units	Low	Medium	High
Tokens in representative scenario	Tokens	26,000	26,000	26,000
FLOPs for inference	FLOP	1.3×10^{15}	2.6×10^{15}	5.2×10^{15}

Source: Europe Economics analysis.

Time savings from AI assistance

Evidence from recent studies suggests that generative-AI tools can reduce the time spent on professional copy-editing tasks, though the estimated magnitude of savings varies by evidence type. Observational and timed evidence from professional editing contexts indicates relatively modest reductions in editing time, of around 10 per cent.[38] Broader experimental studies of AI-assisted writing tasks suggest larger time savings, in some cases approaching 40 per cent, based on measured task-completion times in controlled experiments.[39]

Reflecting this evidence, we adopt a conservative medium estimate of a 15 per cent time saving, placing greater weight on observed professional editing evidence of around 10 per cent. A low scenario of 40 per cent reflects higher, self-reported workflow savings, while a high scenario of 10 per cent reflects only the observed evidence. This approach prioritises measured behaviour over self-reported outcomes, while considering more optimistic assumptions as a sensitivity in the low scenario.

Time to copy-edit and office electricity use

In the digital (AI-assisted) pathway, copy-editing speed is modelled at the sentence level. The medium estimate assumes that copy-editing a single sentence takes 31 seconds, based on observational evidence from professional sentence-level editing tasks reported by Vadhera (2025).[40] A ±20 per cent range is then applied around this central estimate to reflect variation in editing complexity and reviewer speed, yielding a range of 25 to 37 seconds per sentence. These timings reflect the time required for careful review and correction of existing text, rather than initial drafting.

These assumptions imply a baseline human editing speed of between 96 and 144 sentences per hour, with a medium estimate of around 116 sentences per hour. Applying the assumed AI-related efficiency gains increases effective editing speed in the AI-assisted pathway. Office electricity consumption is assumed to scale linearly with the time spent editing and is modelled using standard office electricity-intensity benchmarks, consistent with the approach used elsewhere in the report.

Table 3.21: Time assumptions for copy-editing

[38] Vadehra, A., Johnson, B., Saunders, G., & Poupart, P. (2025). Time Is Effort: Estimating Human Post-Editing Time for Grammar Error Correction Tool Evaluation. In *Proceedings of the Fourth Workshop on Bridging Human-Computer Interaction and Natural Language Processing (HCI+ NLP)* (pp. 178-196). [online]
[39] Forrester (2024) "The Total Economic Impact of Grammarly" [online]
[40] *Ibid.*

	Units	Low	Medium	High
Time saving with AI assistance				
Time saving using AI tools	%	40%	15%	10%
Time to copy-edit (digital)				
Speed of human editor	Sentences edited /hour	144	116	96
Speed of human editor with AI efficiencies	Sentences edited /hour	241	136	107

Source: Europe Economics analysis.

3.3.9 AI-assisted data extraction

This use case reflects the use of a generative-AI data-extraction assistant to identify and extract structured data elements from an unstructured document, followed by human checking, compared with manual data extraction carried out entirely by a human user. In both pathways, the underlying activity is identical: extracting a defined set of data elements from the same source document to a comparable level of completeness and accuracy. For comparability, the representative document is assumed to be a dense, full-text clinical study, consistent with the document type used in the study that informs our estimates of AI time savings. The difference lies in how the extraction task is performed.

In the AI-assisted pathway, the user uploads the full source document together with a short instruction specifying the data elements to be extracted. The AI processes the document and returns the extracted data, which the user then reviews and, where necessary, corrects or completes. In the manual pathway, the user reads through the same document and manually identifies and records the required data elements without AI support.[41]

The AI-assisted pathway introduces additional data centre activity, including document upload, AI inference, and the shared energy associated with training the underlying model, while reducing the time the user spends on manual review and extraction. This makes the use case well suited to assessing the trade-off between increased data centre electricity use and reduced office electricity consumption in document-intensive analytical tasks.

The reduced time spent extracting data represents an increase in the productivity of workers, allowing them to perform more data extraction tasks in any given period of time.

Representative data-extraction activity

The representative scenario assumes that a knowledge worker extracts structured data from a single unstructured document, assumed to be a dense, full-text clinical study. For comparability

[41] Gartlehner, G., Kugley, S., Crotty, K., Viswanathan, M., Dobrescu, A., Nussbaumer-Streit, B., ... & Kahwati, L. C. (2025). Artificial Intelligence–Assisted Data Extraction With a Large Language Model: A Study Within Reviews. *Annals of Internal Medicine*, *178*(12), 1763-1771. [online]

with the evidence base used to inform AI time-saving assumptions, the document is assumed to contain 15 pages, with approximately 1,000 words per page, resulting in a total document length of around 15,000 words. The number of data elements to be extracted is assumed to range from 120 to 180, with a medium estimate of 150, based on empirical evidence reported in *Gartlehner et al. (2025)* on data extraction in systematic reviews.

In the AI-assisted pathway, the user uploads the full document to the AI data-extraction tool together with a short instruction specifying the required data elements. Using a page-to-data conversion of 0.000015 GB per page and adjusting for higher page density relative to a standard word document, the total size of the uploaded document is estimated at 0.00045 GB. The AI processes the full text and returns the extracted data, which the user then reviews and amends if necessary. Output size is determined by the number of extracted data elements and the length of each extracted value. The number of output tokens per data element is assumed to range from 3 to 10 tokens, reflecting variation between simple numeric values and more complex text-numeric entries, resulting in between 360 and 1,800 output tokens per study.

In the manual pathway, the same document is reviewed by the user without AI support, and the same set of data elements is identified and recorded manually. The comparison therefore isolates the impact of AI assistance on data centre activity and time spent on data extraction, holding the document type, length and extraction task constant across both pathways.

Inference workload

For the representative data-extraction scenario, a single dense, full-text clinical study is assumed, comprising 15 pages with approximately 1,000 words per page. The full document is uploaded to the AI system, together with a short instruction-based prompt of around 50 tokens specifying the data elements to be extracted.

The model is assumed to extract between 120 and 180 data elements, with outputs ranging from simple numeric values to short text–numeric entries. Taken together, the document upload, prompt overhead, and extracted outputs imply total token volumes of approximately 20,000 to 21,500 tokens per document in the low, medium, and high scenarios.

Applying the per-token inference compute assumptions reported in Table 3.15 yields total inference requirements of approximately 1.0×10^{15}, 2.1×10^{15}, and 4.3×10^{15} FLOPs respectively.

Table 3.22: Input assumptions for digital scenario – AI-assisted data extraction

	Units	Low	Medium	High
Tokens used for data extraction	Tokens	19,910	20,450	21,350
FLOPs for inference	FLOP	9.96×10^{14}	2.05×10^{15}	4.27×10^{15}

Source: Europe Economics analysis.

Time savings from AI assistance and office electricity use

Evidence from empirical studies indicates that generative-AI tools can materially reduce the time required for professional data-extraction tasks, though the magnitude of savings varies across documents and workflows. The primary study used to inform this use case (Gartlehner et al., 2025) reports a median time saving of approximately 33 per cent per study when using AI assistance relative to a fully manual approach. Across the six systematic reviews analysed, observed time savings ranged from a minimum of around 19 per cent to a maximum of approximately 42 per cent in five cases, while one review experienced an increase in total time.

Reflecting this evidence, we adopt a range of AI-related time savings of 19 per cent (high case), 33 per cent (medium case) and 42 per cent (low case). These values are taken directly from the minimum, median and maximum time efficiencies observed in the primary study and therefore reflect measured performance rather than self-reported estimates. The reported time savings include the full extraction workflow, including verification and, in the AI-assisted pathway, time spent on prompt specification and refinement.

In the manual pathway, the time taken to extract data from a single document is assumed to range from 83 to 184 minutes, with a medium estimate of 125 minutes, based on the same study (Gartlehner et al., 2025). Applying the assumed AI-related time savings yields proportionally lower effective extraction times in the AI-assisted pathway. Office electricity consumption is assumed to scale linearly with the time spent on data extraction and is modelled using standard office electricity-intensity benchmarks, consistent with the approach applied elsewhere in the report.

Table 3.23: Time assumptions for data extraction

	Units	Low	Medium	High
Time saving with AI assistance				
Time saving using AI tools	%	42	33	19
Time for data extraction (digital)				
Speed of human data extractor	Hours / document reviewed	1.38	2.08	3.07
Speed of human data extractor with AI efficiencies	Hours / document reviewed	0.80	1.40	2.48

Source: Europe Economics analysis.

3.3.10 AI-assisted transcription

This use case reflects the use of a generative-AI transcription assistant to convert spoken audio into written text, compared with manual transcription carried out entirely by a human. In both pathways, the underlying activity is identical: producing an accurate written transcript of the same audio content to a comparable standard of completeness and accuracy. The difference lies in how the transcription task is performed.

In the AI-assisted pathway, the user uploads an audio recording to the AI transcription tool, which processes the file and returns a text transcript. The user then reviews the output and makes any necessary corrections. In the manual pathway, the user listens to the same audio recording and manually transcribes the content without AI support.

The AI-assisted pathway introduces additional data centre activity, including audio file upload, AI inference during transcription, and the shared electricity associated with training the underlying model, while reducing the time the user spends on transcription. This makes the use case well suited to assessing the trade-off between increased data centre electricity use and reduced office electricity consumption for audio-based information processing tasks.

The time savings from AI assistance represent a productivity increase for workers, allowing them to extract more data during their working hours.

Representative transcription activity

The representative scenario assumes that a knowledge worker transcribes a single audio recording of one hour in length. The recording is assumed to be a standard spoken recording (for example, a meeting, interview or lecture), with an average speech rate resulting in approximately 10,500 words in the final transcript.[42] The worker is assumed to be office-based.

In the AI-assisted pathway, the user uploads the one-hour audio file to the AI transcription tool. Based on medium estimates for common audio formats, the file size of the uploaded audio is assumed to be approximately 0.10 GB.[43] The AI processes the audio and returns a text transcript, which the user then reviews and corrects if necessary. The downloaded transcript is assumed to have a file size of approximately 0.00002 GB, reflecting the small data footprint of text relative to audio. We assume a model lifetime of one year, reflecting the observed update cycle of specialised speech-to-text models; for example, OpenAI's Whisper Large V2 was released in December 2022 and replaced by Whisper Large V3 in November 2023.

In the manual pathway, the same audio recording is transcribed by the user without AI support. The comparison therefore isolates the impact of AI assistance on data centre activity and time spent on transcription, holding the audio length, content and output volume constant across both pathways.

User-base assumptions for shared allocation

For the purpose of allocating electricity from shared processes, particularly AI model training, the user base for the AI-assisted transcription use case is defined in terms of the total number of hours of audio processed by the model over its assumed operational lifetime of one year.

[42] GMR Transcription (2025) "How Long Does It Take to Transcribe a 1-Hour Interview?" [online]
[43] Soundcy (2026) "Understanding The Storage Space Required For One Hour Of Sound" [online]

Hours of audio transcribed are used as the allocation unit to reflect the fact that computational demand in speech-to-text systems scales primarily with audio duration.

In the absence of a single published figure for total transcription volumes, a top-down approach is applied using publicly available pricing information and estimates of market size. Public pricing for Whisper-class transcription services is approximately USD 0.36 per hour of audio processed.[44] Estimates of the global AI transcription market indicate an annual market value of around USD 5 billion.[45] Taken together, these figures imply total global transcription volumes of the order of 14 billion hours per year.

To allocate shared electricity across users, three scenarios are defined based on illustrative market shares for a single transcription model provider. The high-usage case assumes a large, market-leading model capturing 10 per cent of total market activity, corresponding to approximately 1.4×10^9 hours of transcription per year. The medium and low usage assume market shares of 10 per cent and 1 per cent respectively, corresponding to around $6.9 \times 10^8$and 1.4×10^8 hours per year. These values are intended to span a plausible range of deployment scales, from dominant models to more specialised or niche models.

AI Model Size

The AI-assisted transcription use case is assumed to rely on a specialised speech-to-text model rather than a general-purpose language model. Model size at inference is therefore substantially smaller than for large text-generation models. We assume a representative speech-to-text model with around 1.5 billion parameters.[46] The digital size of stored model weights is estimated to range from approximately 1.5 GB to 6 GB, depending on numerical precision and deployment configuration. The low estimate reflects INT8 quantisation, the medium estimate assumes FP16 or BF16 precision commonly used in production deployments, and the high estimate reflects FP32 storage.[47]

Training compute for speech-to-text models is assumed to span a wide range, reflecting substantial variation in model architecture, dataset size, and training methodology. Based on published estimates compiled by Epoch AI and reported by Our World in Data, training compute is assumed to range from approximately 3.1×10^{17}FLOPs in the low case to 7.6×10^{21}FLOPs in the high case, with a medium estimate of 3.8×10^{21}FLOPs. The low estimate corresponds to earlier large-scale speech recognition systems, while the high estimate reflects modern end-to-end automatic speech recognition architectures combining pretraining with large-scale semi-supervised learning.

Table 3.24: Input assumptions for digital scenario (user base and AI model) – AI-assisted transcription

[44] OpenAI (nd.) “Whisper” [online]
[45] Sonix (2026) “25 Meeting Transcription Adoption Statistics Every Professional Should Know in 2026” [online]
[46] Nvidia (n.d.) “whisper-large-v3” [online]
[47] Github (2025) “Train Big, Plan Smart - How to Calculate Memory and Estimate GPUs for LLMs” [online]

	Units	Low	Medium	High
User base for shared processes				
User base for shared processes	Hours	1.4×10^{9}	6.9×10^{8}	1.4×10^{8}
AI transcription model				
Size of model	GB	1.5	3	6
FLOPs used in training AI model	FLOP	3.1×10^{17}	3.8×10^{21}	7.6×10^{21}

Source: Europe Economics analysis.

AI computational requirements

Compute requirements for AI-assisted transcription are estimated using a number of steps, with the unit of activity defined as the transcription of one hour of recorded audio.

The first step is to estimate the runtime of the AI system. This is captured using the real-time factor (RTF), defined as the ratio of processing time to audio duration. Evidence from service-level deployments of speech-to-text systems (such as cloud or API-based transcription) suggests RTF values typically in the range 0.2–0.6,[48] reflecting end-to-end operation including decoding, batching, orchestration and other overheads. By contrast, hardware-focused benchmarks under ideal conditions indicate that substantially faster performance is technically possible, with reported RTF values as low as around 0.025[49] when high-end GPUs are fully utilised. These benchmark results are interpreted as upper bounds rather than representative of typical single-file transcription workloads. Reflecting this evidence, an RTF range of 0.05 to 0.6 is adopted, with a medium estimate of 0.4. Under these assumptions, transcribing a one-hour audio file requires approximately 3 minutes (low), 24 minutes (central), or 36 minutes (high) of processing time.

The second step anchors the calculation to representative data centre hardware. Peak compute capability is based on published specifications for an NVIDIA A100 GPU. The medium estimate uses the reported peak throughput of 3.12×10^{14} FLOP/s,[50] with the low and high values defined as ±20 per cent around this figure to reflect uncertainty across deployment configurations, clock speeds and precision modes.

As inference workloads do not continuously operate at peak performance, peak throughput is scaled by an assumed utilisation rate. Utilisation is assumed to range from 30 per cent in the low case to 50 per cent in the high case, with a medium estimate of 40 per cent.[51] These

48 OpenVoice Tech (2024) "Real-time-factor" [online]
49 Salad (2024) "Whisper Large V3 Speech Recognition Benchmark: 1 Million hours of audio transcription for just $5110" [online]
50 Nvidia (n.d.) ""NVIDIA A100 TENSOR CORE GPU [online]
51 OPeninfer (2025) "Unlocking the Full Potential of GPUs for AI Inference" [online]

values are informed by public evidence on realised GPU utilisation in inference workloads and reflect a conservative view of achievable performance in shared production environments.

In addition, transcription jobs are assumed to share GPU resources with other concurrent workloads rather than occupying a full GPU in isolation. The share of GPU effectively allocated to a single transcription task is assumed to be 14 per cent in the low case, corresponding to one-seventh of an NVIDIA A100 GPU, consistent with common multi-instance GPU partitioning configurations.[52] The medium and high cases assume higher effective allocations of 25 per cent and 50 per cent respectively, reflecting increasingly conservative assumptions about workload isolation.

Total inference compute is calculated as the product of runtime, peak throughput, utilisation rate and GPU allocation share. Under these assumptions, total compute requirements for transcribing one hour of audio are estimated at approximately 1.9×10^{15} FLOPs in the low case, 4.5×10^{16} FLOPs in the medium case, and 1.7×10^{17} FLOPs in the high case.

Table 3.25: FLOP used in inference of audio recording

	Units	Low	Medium	High
Audio duration	Seconds	3600	3600	3600
Real-time factor		0.05	0.10	0.20
GPU runtime	Seconds	180	1440	2160
Peak GPU throughput	FLOP/s	2.50×10^{14}	3.12×10^{14}	3.74×10^{14}
GPU utilisation	Percentage	30%	40%	50%
Achieved throughput	FLOP/s	7.50×10^{13}	1.25×10^{14}	1.56×10^{14}
Share of GPU used	Percentage	0.14	0.25	0.50
Total compute	FLOPs	**1.9×10^{15}**	**4.5×10^{16}**	**1.7×10^{17}**

Source: Europe Economics analysis.

Time savings from AI assistance and office electricity use

Evidence from recent studies indicates that generative-AI tools can substantially reduce the time required for transcription tasks, though the magnitude of savings varies by use case and workflow. We use a study by Sponholz et al. (2025) to inform this use case, which reports large reductions in transcription time when AI assistance is used, with observed time savings ranging from around 50 per cent to 75 per cent relative to fully manual transcription.[53]

Reflecting this evidence, we adopt a range of AI-related time savings of 50 per cent (high case), 62.5 per cent (medium case) and 75 per cent (low case). These values are derived

[52] *Ibid.*
[53] Sponholz, J., Weilinghoff, A., & Schopf, J. (2025). Halving transcription time: A fast, user-friendly and GDPR-compliant workflow to create AI-assisted transcripts for content analysis. *arXiv preprint arXiv:2503.13031*. [online]

directly from the minimum and maximum time efficiencies reported in the study, with the medium estimate taken as the midpoint of the observed range. The reported time savings reflect the full transcription workflow, including review and correction of AI-generated transcripts.

In the manual pathway, the time taken to transcribe a one-hour audio recording is assumed to range from 3 to 5 hours, with a medium estimate of 4 hours, consistent with published benchmarks for human transcription performance.[54] Applying the assumed AI-related time savings yields effective transcription times of approximately 2.5 hours, 1.5 hours, and 0.75 hours in the high, medium and low AI-efficiency cases respectively. Office electricity consumption is assumed to scale linearly with the time spent on transcription and is modelled using standard office electricity-intensity benchmarks, consistent with the approach applied elsewhere in the report.

Table 3.26: Time assumptions for transcription

	Units	Low	Medium	High
Time saving with AI assistance				
Time saving using AI tools	%	75.0	62.5	50.0
Time to transcribe (digital)				
Speed of human transcription (time taken to transcribe one hour of audio)	Hours	3.00	4.00	5.00
Speed of human data extractor with AI efficiencies	Hours	0.75	1.50	2.50

Source: Europe Economics analysis.

3.3.11 AI-assisted demand forecasting model development

This use case examines the development of a short-term electricity demand forecasting model using an automated machine learning (AutoML) system, compared with a conventional manual modelling approach carried out by analysts. In both pathways, the objective is the same: to produce a statistical model that forecasts electricity demand over a short time horizon using the same underlying dataset and achieving a comparable analytical outcome. The distinction lies in the method used to build and evaluate the model.

The model is assumed to be estimated using a historical dataset containing 20 years of half-hourly electricity consumption data, along with multiple weather variables (potentially across locations), calendar effects, and a wide set of engineered features (e.g. lags, rolling statistics and interaction terms). The precise model specification and choice of variables is determined during the modelling process. We assume that the model is designed so as to be able to project half-hourly electricity demand for several days ahead.

[54] GoTranscript (2025) “How Long Does It Take to Transcribe an Hour of Audio?” [online]

Under the AI-assisted pathway, analysts run an AutoML workflow that automatically estimates and evaluates multiple candidate models, explores different hyperparameter configurations and ranks the resulting models based on performance. The analysts review the outputs and select the final model from the set generated by the system. Under the manual pathway, analysts carry out these tasks themselves, developing and testing alternative model specifications iteratively using standard statistical tools.

The AI-assisted pathway therefore introduces additional data centre activity associated with running the automated model-search process, as well as the shared electricity associated with developing the AutoML system itself. In contrast, the manual pathway relies primarily on analyst time and local computing resources.

AI computational requirements

Compute requirements for AI-assisted electricity demand forecasting model development are estimated using a number of steps.

The first step is to estimate the runtime of the AutoML system when developing and evaluating the electricity demand forecasting model. The use of AutoML systems is typically described in terms of the time budget allocated to the automated model search or the number of candidate models trained during the process. For example, H2O AutoML allows the user to specify either a maximum runtime or a maximum number of models, and defaults to a one-hour run if neither is specified.[55] Similarly, AutoGluon TimeSeries likewise allows the user to set a time limit for the model search, and a single fit call trains multiple models as part of that process.[56]

Reflecting this evidence, the AutoML runtime for the representative forecasting use case is assumed to range between **1 hour and 12 hours**, with a central estimate of **4 hours**. This runtime refers to the total model development process, which may involve multiple AutoML searches. Individual searches are typically shorter (e.g. 1–4 hours), with users often running several iterations to refine model performance.

These assumptions represent three plausible levels of search effort: a short bounded run, a more substantive search across multiple candidate models, and a more exhaustive search exploring a larger space of model specifications. Hour-scale budgets are common in AutoML benchmarking studies, and forecasting experiments using AutoGluon TimeSeries report results for training budgets of 10 minutes, 1 hour and 4 hours, indicating that hour-scale search times are realistic for forecasting workflows.[57]

The second step anchors the calculation to representative data centre hardware. To reflect variation in how AutoML workflows are deployed in practice, compute capability is estimated using a range of plausible hardware configurations rather than a single GPU accelerator

[55] H20 (n.d.) “H2O AutoML: Automatic machine learning” [online]
[56] AutoGluon (n.d.) “Time Series Forecasting” [online]
[57] Shchur, O., Turkmen, A. C., Erickson, N., Shen, H., Shirkov, A., Hu, T., & Wang, B. (2023). AutoGluon–TimeSeries: AutoML for probabilistic time series forecasting. In *International Conference on Automated Machine Learning* (pp. 9-1). PMLR. [online]

benchmark. The high case is based on published specifications for an NVIDIA A100 GPU, which has been widely used in both research and industry for machine-learning workloads. NVIDIA's datasheet reports peak Tensor Float-32 (TF32) performance of approximately **1.56 × 10^{14} FLOPs per second**, which is used as the reference throughput for the high scenario.[58] The low case reflects a CPU-only deployment. For this, peak compute capability is approximated using a representative CPU-only server configuration, based on a dual-socket set-up with 56 cores per socket running at 2.0 GHz.[59] This gives an estimated throughput of **1.43 × 10^{13} FLOPs per second** for the CPU-only case.[60] The medium case is set at the midpoint between this value and the high-end GPU benchmark.

As AutoML workflows do not operate continuously at peak hardware performance, peak throughput is scaled by an assumed utilisation rate. Utilisation is assumed to range between **30 per cent and 50 per cent**, with a central estimate of **40 per cent,**[61] to translate theoretical peak throughput into realised throughput in shared production environments. These values should not be interpreted as literal processor occupancy, but as an adjustment for the gap between theoretical peak and achievable performance in practical AutoML workflows.

Total compute requirements are then calculated as the product of runtime, peak hardware throughput and the assumed utilisation rate. Under these assumptions, the compute required to develop a forecasting model using an AutoML workflow is estimated to be approximately 1.5 × 10^{16} FLOPs in the low case, 4.9 × 10^{17} FLOPs in the medium case, and 3.4 × 10^{18} FLOPs in the high case.

Table 3.27: FLOPs used in AutoML forecasting model development

	Units	Low	Medium	High
AutoML runtime	Seconds	3,600	14,400	43,200
Reference GPU throughput	FLOPs/s	1.43 × 10^{13}	8.52 × 10^{13}	1.56 × 10^{14}
GPU utilisation	Percentage	30	40	50
Achieved throughput	FLOPs/s	4.29 × 10^{12}	3.41 × 10^{13}	7.80 × 10^{13}
Total compute	FLOPs	1.5 × 10^{16}	4.9 × 10^{17}	3.4 × 10^{18}

Source: Europe Economics analysis.

AutoML platform development electricity requirements

[58] Nvidia "NVIDIA A100" [online]
[59] HPE "Intel Processors: Intel Xeon-Platinum 8480+ 2.0GHz 56-core 350W Processor for HPE(P49607-B21)" [online]
[60] The calculation estimates peak compute throughput by combining the number of processors, the number of cores, the processor speed, and the amount of computation each core can perform per cycle. This gives an estimated peak throughput of 1.43 × 10^{13} FLOPs/s, calculated as 2 sockets × 56 cores per socket × 2.0 × 10^{9} cycles per second × 64 FLOPs per cycle per core.
[61] OpenInfer (2025) "Unlocking the Full Potential of GPUs for AI Inference" [online]

The development of an AutoML platform requires substantial electricity, reflecting the need to develop a platform which can produce, test and benchmark models across a range of datasets and configurations. Electricity requirements are therefore estimated using a top-down approach, with the unit of activity defined as the development of a single AutoML system.

The starting point for the calculation is a published estimate of compute usage associated with AutoML research and development. Tornede et al. (2023) report that an AutoML research group consumed approximately 5.5 million CPU-hours, which they state is equivalent to around **92.8 MWh of CPU-related electricity use over one year**. This reflects the electricity required to design, test and benchmark AutoML systems across multiple datasets and configurations.[62]

To convert this into the development of a single AutoML platform, the annual figure is scaled by an assumed development period. AutoML systems are typically developed iteratively over multiple years, reflecting ongoing experimentation, benchmarking and system refinement. For example, widely used frameworks such as auto-sklearn has evolved over several years through successive releases.[63] On this basis, a 2-year development period is adopted as a representative assumption.

The published estimate primarily captures CPU-related electricity use and does not fully reflect the contribution of other IT components, including memory, storage and networking. As highlighted in the Green AutoML literature (Tornede et al., 2023), these components play a material role in supporting data handling, orchestration and model training workflows. To account for this, a modest uplift of **5 to 20 per cent** is applied, with a central estimate of **15 per cent**.

Combining these assumptions, total IT electricity energy consumption associated with developing a single AutoML platform is estimated to be:

- Low: 92.8 × 2 × 1.05 = 194.9 MWh
- Central: 92.8 × 2 × 1.15 = 213.4 MWh
- High: 92.8 × 2 × 1.20 = 222.7 MWh

Size of AutoML platform

The size of the AutoML platform is estimated based on the storage footprint of representative AutoML software environments.

AutoML systems are typically deployed as containerised environments that bundle machine learning libraries, dependencies and supporting tools. Published examples indicate that these environments are substantially larger than source-code packages alone. For example, the

[62] Tornede, T., Tornede, A., Hanselle, J., Mohr, F., Wever, M., & Hüllermeier, E. (2023). Towards green automated machine learning: Status quo and future directions. *Journal of Artificial Intelligence Research*, *77*, 427-457. [online]
[63] Feurer, M., Klein, A., Eggensperger, K., Springenberg, J., Blum, M., & Hutter, F. (2015). Efficient and robust automated machine learning. *Advances in neural information processing systems*, *28*. [online]

auto-sklearn Docker image provides a relatively lightweight AutoML runtime (around 0.5GB),[64] while the AutoGluon framework distributes larger containerised environments (around 2.6 GB) for broader machine learning workloads.[65]

Reflecting this evidence, the size of the AutoML platform is assumed to range between 0.5 GB and 2.6 GB, with a central estimate of 1.55 GB, representing lightweight, standard and more comprehensive deployment configurations.

Size of data

A dataset **size of 2 GB** is assumed to reflect a rich modelling environment, including approximately 20 years of half-hourly electricity demand data, combined with multiple weather variables (potentially across locations), calendar effects, and a wide set of engineered features (e.g. lags, rolling statistics and interaction terms).

Data downloads from the system are assumed to be limited to summary outputs, diagnostics and model results, and are therefore set at **1 per cent of the uploaded data volume**.

User-base assumptions for shared allocation

The user base for AutoML-assisted statistical model development is estimated by combining evidence on platform adoption with assumptions on user activity. The unit of activity is defined as the total number of AutoML search-hours per year, representing the aggregate compute workload generated by users of the platform.

Published evidence indicates that AutoML platforms are used across a large number of organisations and users. For example, H2O reports usage by over 100,000 data scientists across more than 10,000 organisations, implying an average of around 10.8 users per organisation.[66] Similarly, DataRobot reports usage across more than 1,000 enterprise customers, indicating a lower bound for platform adoption.[67] Taken together, this evidence suggests a plausible range for the number of organisations using a representative AutoML platform of 1,000 to 12,000, with a central estimate of 5,000 organisations. Based on the H2O evidence, we assume an average of around 10.8 AutoML users per organisation.

The frequency of AutoML use is characterised by the number of searches undertaken per user each year. This is highly uncertain, as public evidence on actual usage intensity is limited and likely to vary substantially across organisations, teams and use cases. AutoML workflows are typically iterative, with users running multiple searches to refine model performance, but the extent of this iteration is not well observed in public sources. To reflect this uncertainty, average search frequency is assumed to range from 12 runs per year (approximately monthly

[64] https://hub.docker.com/r/mfeurer/auto-sklearn/
[65] [online]
[66] H20 [online]
[67] Datarobot [online]

use) to 52 runs per year (approximately weekly use), with a central estimate of 24 runs per year.

Each AutoML search is defined in terms of a time budget for model training and evaluation. Hour-scale runtimes are common in the AutoML literature and benchmarking studies, where workflows are typically defined using explicit time budgets for model search. To reflect this, the duration of each search is assumed to range from 1 hour to 4 hours, with a central estimate of 2 hours, capturing light, moderate and more exhaustive search processes. This refers to the duration of a single AutoML search, rather than the total runtime of the representative forecasting-model development process, which may involve multiple searches.

Combining these assumptions yields estimates of total AutoML activity expressed in annual search-hours. The resulting estimates are summarised below. These figures should be interpreted as indicative rather than definitive, as they are sensitive to our assumptions regarding the average number of searches undertaken per user each year, and as discussed earlier there is material uncertainty around this input.

Table 3.28: AutoML user base and activity

Parameter	Low	Medium	High
Organisations using platform	12,000	5,000	1,000
Users per organisation	10.8	10.8	10.8
Total users	**129,600**	**54,000**	**10,800**
Searches per user per year	52	24	12
Hours per search	4	2	1
Total AutoML search-hours per year	**26,956,800**	**2,592,000**	**129,600**

Source: Europe Economics analysis.

Time savings from AutoML assistance

Experimental evidence suggests that AI assistance can materially reduce the human time required for statistical modelling and machine-learning tasks. For example, Wang et al. (2021) find that tasks completed using an AutoML-style system took around five minutes compared with roughly fifteen minutes using a conventional workflow, implying a reduction of approximately 67 per cent.[68] Similarly, Yao et al. (2025) report task completion times of 8.5

[68] Wang, D., Andres, J., Weisz, J. D., Oduor, E., & Dugan, C. (2021, May). Autods: Towards human-centered automation of data science. In *Proceedings of the 2021 CHI conference on human factors in computing systems* (pp. 1-12). [online]

minutes versus 17.3 minutes for an LLM-driven AutoML approach, corresponding to a reduction of approximately 51 per cent.[69]

Taken together, these studies suggest time savings in the range of 50–70 per cent in controlled settings. A central estimate of 60 per cent is adopted as a midpoint within this range, reflecting typical reductions while recognising that some elements of the modelling process continue to require human judgement.

[69] Yao, J., Zhang, L., & Huang, J. (2025). Evaluation of large language model-driven AutoML in data and model management from human-centered perspective. *Frontiers in Artificial Intelligence*, *8*, 1590105. [online]

4 Methodology for Modelling Electricity Consumption in Counterfactuals

4.1 Counterfactuals involving office-based work

4.1.1 Overview of approach

This section outlines our approach to estimating electricity consumption for use cases in which the service would, under a physical alternative, be delivered by a professional working in an office environment.

The calculation draws on the electricity consumption attributable to the office space occupied by the service provider. The analysis is grounded in the same representative scenario used for the digital use case, which sets the parameters of the activity under consideration.

Similarly to the methodology highlighted in section 3.1.2, the primary electricity contribution in this scenario is the electricity required to support the working environment and computer equipment of the person delivering the service. This is estimated using a bottom-up approach:

We apply a **standard electricity energy use intensity (EUI)** for commercial office buildings, measured in kilowatt-hours per square metre per year (kWh/m²/year).

This is multiplied by an assumption for **average office space per worker** (in m²), yielding an estimate of **annual electricity use per office worker**.

The EUI figures used are all-inclusive: they represent the total annual metered electricity demand of the premises and therefore encompass all forms of electricity use within the building. This includes electricity for IT and computing equipment, lighting, heating and cooling, ventilation, and other building services. As such, they already incorporate the electricity used by the office-based worker’s device and supporting infrastructure.

We then scale this annual figure based on the **time required to perform the specific task,** without any AI assistance. This provides an estimate of the **electricity consumed during the performance of the activity**.

4.1.2 Input assumptions for relevant use cases

We apply the same input assumptions for electricity use per office worker per working day as we used to calculated office electricity consumption in the AI use cases, as set out in Section 3.2.4.

Office electricity consumption is assumed to scale linearly with the time taken to complete the relevant task in each use case below.

Manual email drafting - In the manual drafting pathway, the speed of human emailing is assumed to range from 300 to 800 words per hour, with a medium estimate of 500 words per hour. These values reflect the drafting of professional-quality business emails, including planning, structuring and revision.

Manual Code drafting - In the manual coding pathway, developers are assumed to write between 25 and 125 lines of code per day, with a medium estimate of 75 lines per day. For the representative scenario, we combine this with the assumed average of 40 characters per line and a 7.5-hour productive workday, giving human coding speeds of approximately 133 to 667 characters per hour.

Manual copy-editing - In the manual copy-editing pathway, editing speed is modelled at the sentence level. The time taken to edit a single sentence is assumed to range from 25 to 37 seconds, with a medium estimate of 31 seconds, based on observational evidence from professional editing contexts. These assumptions imply copy-editing speeds of approximately 96 to 144 sentences per hour, with a medium estimate of around 116 sentences per hour. These figures reflect the time required for careful review and correction of existing text rather than initial drafting.

Manual data extraction - In the manual data-extraction pathway, the time taken to extract structured data from a single document is assumed to range from 1.38 to 3.07 hours, with a medium estimate of 2.08 hours. These values are drawn from the primary study used to inform estimates of AI time savings and reflect the full extraction workflow, including document review and recording of data elements.

Manual demand forecasting model – In the manual pathway, developing a statistical demand forecasting model is assumed to take around 15-25 analyst-days. This reflects the full workflow, including model specification, calibration, and validation. This assumption is representative and intended to capture a typical short-term demand forecasting exercise under standard working conditions.

Table 4.1: Time assumptions for office-based work

	Units	Low	Medium	High
Speed of human emailing	Words drafted per hour	800	500	300
Speed of human coder	Characters coded per hour	667	400	133
Speed of copy-editing	Sentences edited per hour	144	116	96
Speed of data extraction	Time taken to extract data from one document (hours)	1.38	2.08	3.07

	Units	Low	Medium	High
Speed of transcription	Time taken to transcribe one-hour audio recording (hours)	3	4	5
Speed of model development	Number of analyst-days required to build model per person	15	20	25

Source: Europe Economics analysis.

4.2 Other counterfactuals

4.2.1 TV broadcasting

Overview of approach

This section outlines our approach to estimating electricity consumption for the use case in which live television is delivered, under the counterfactual, via the UK's digital terrestrial television (DTT) platform and viewed on a television set in the home. The calculation is grounded in the same representative viewing session (four hours of live viewing). Electricity estimates are constructed from three components: electricity attributable to the terrestrial broadcast infrastructure required to deliver DTT services; and electricity consumed by the end-user television during viewing.

Terrestrial infrastructure (transmission network)

The terrestrial infrastructure component captures the electricity required to deliver DTT services across the UK's broadcast distribution and transmission chain – i.e. the full end-to-end terrestrial delivery system that takes live TV from the broadcaster's distribution systems through the national network of transmitters/relays and "radiates" it for reception via household aerials. Importantly, this is not just a single mast or local link: it reflects the whole shared DTT transmission network, together with the centralised broadcast processing and distribution facilities (server/computing infrastructure) needed to originate and route the signal into that network.

We estimate this using an electricity intensity expressed per device viewing hour (kWh per viewer-hour). Rather than modelling individual transmitters bottom-up, we take the UK-specific "delivery system" electricity intensity reported in the Ofcom-commissioned Carnstone analysis and apply it directly as the terrestrial infrastructure intensity.[70] For this counterfactual, we take the Ofcom/Carnstone "data centres and network transmission" figure for DTT as 2.6 Wh per device hour. We then construct uncertainty bounds by applying a symmetric variation of ±20% around this medium value to generate low and high cases. This provides a UK-specific,

[70] Carnstone (2024) "Carbon emissions of streaming and digital terrestrial television" [Online]

empirically grounded estimate that already reflects the allocation of shared broadcast infrastructure to viewing activity on a per-hour basis.

End-user devices

The end-user device component captures the electricity consumed by the television set while live TV is being watched. This reflects the fact that, once the DTT signal reaches the household, the TV must be powered to decode and display the content.

This component is estimated by applying an average television electricity consumption rate during viewing, expressed as kWh per viewer-hour. In the modelling, this rate is specified as a range with low, medium, and high figures to reflect variation in TV technologies, sizes, and operating conditions. The TV electricity is summed with the terrestrial infrastructure component to obtain total electricity per viewer-hour.

4.2.2 Gaming on own computer

This counterfactual considers the case in which a user plays a video game locally on their own device using a physical game disc. The analysis incorporates three broad components of electricity use: the production and storage of the disc, transportation through the supply chain, and end-user device electricity consumption associated with installation and gameplay.

The production and storage component captures the main stages involved in supplying a physical game disc, including manufacture of the disc and its packaging, warehousing, and retail store operations. The transport component reflects movement from the manufacturing site to the warehouse and onwards to retail outlets, with separate assumptions applied for UK and overseas production. In the latter case, the supply chain also encompasses ferry transport to the UK. Alternative assumptions are also applied to reflect decarbonised transport pathways.

End-user device electricity consumption is modelled separately using a range of representative devices. The low estimate is based on a gaming console and the high estimate on a gaming PC, with the medium estimate defined as the midpoint between the two. The installation stage is also reflected within this component through an allowance for the electricity associated with data transfer during installation.

Low, medium and high estimates are developed for each component in order to reflect uncertainty in the underlying assumptions. Full details of the assumptions and calculation steps for disc production, storage and transportation are provided in the Appendix.

4.2.3 Short in-person course

Overview of approach

This counterfactual involves a person attending a short in-person course in an educational institution, rather than viewing pre-recorded training content online. We assume that total energy consumed in the process is a combination of the following: energy consumed in travelling to the place where the course takes place and electricity consumed in the building in which the course takes place.

In considering travel-related energy use, we account for both current and potential future patterns of energy consumption. In the baseline case, travel is assumed to take place using a conventional diesel vehicle, reflecting typical current behaviour. In practice, however, the same journey could also be undertaken using an electric vehicle. The implications of this are reported under the decarbonised scenario, in which transport is electrified, consistent with a net zero pathway. This allows us to capture how the relative importance of travel changes as energy systems decarbonise, and to ensure comparability with electricity-based energy use in the digital scenario.

Distance travelled

We have taken our assumption on the distance travelled to the educational building from a dataset on distances travelled for different purposes.[71] In our medium scenario, we assume the distance travelled one-way to be 3.2 miles, the distance travelled by a typical person for educational purposes. We apply a ±20 per cent adjustment to estimate low and high values, resulting in one-way distances of 2.56 miles and 3.84 miles respectively. We convert these figures to kilometres for use in our calculations (1 mile = 1.6 kilometres).

Car energy consumption

In our medium scenario we assume that a representative car travels 46 miles per gallon.[72] We use values of 53 and 40 miles/gallon in our low and high scenarios respectively. We converted these values to litres per kilometre using the following: 1 gallon = 4.546 litres and 1 mile = 1.6 km, so litres per kilometre = 2.825 ÷ miles per gallon. Our calculated values for litres per kilometre are 0.05, 0.06 and 0.07 across the low, medium and high scenarios.

In addition, we consider a decarbonised case in which travel is undertaken using an electric vehicle, with electricity consumption expressed directly in electricity terms based on representative efficiency assumptions. We assume electricity use of 0.125, 0.15 and 0.175 kWh per km across the low, medium and high scenarios.[73]

[71] DfT (2025) “Purpose of travel” [online].
[72] Ageas (2024) [online].
[73] Go-e (2025 “EV Consumption: How Much Energy Does an Electric Car Consume?” [online]

Educational building electricity consumption

We apply a standard electricity energy use intensity (EUI) for educational buildings, measured in kilowatt-hours per square metre per year (kWh/m²/year). In our medium scenario we assume 56.2 kWh/m²/year, scaling it to 47.3 kWh/m²/year in our low scenario and 65.0 kWh/m²/year in our high scenarios.[74] This is multiplied by an assumption for average educational space per student (in m²), yielding an estimate of annual electricity use per student. We assume 10 m², 12.5 m² and 15 m² per student in our low, medium and high scenarios respectively.[75]

The EUI figures used are all-inclusive: they represent the total annual metered electricity demand of the premises and therefore encompass all forms of electricity use within the building. This includes electricity for IT and computing equipment, lighting, heating and cooling, ventilation, and other building services. As such, they already incorporate the electricity used by the student's and the teacher's devices and supporting infrastructure.

We then scale this annual figure per student to obtain the electricity consumed per day, assuming electricity is only consumed during working days, and that there are on average 253.25 working/educational days in a year after accounting for weekends and Bank Holidays (and taking a weighted average of leap years and non-leap years).

The table below summarises our inputs for the short in-person course use-case.

Table 4.2: Assumptions for the short in-person course use case

Category	Low	Medium	High
One-way distance to place of education (miles)	2.56	3.2	3.84
Car energy consumption (miles/gallon)	53	46	40
EUI for educational buildings	47.3	56.2	65.0
Educational space per person	10	12.5	15

Source: Europe Economics analysis.

4.2.4 Back-up onto local server

Overview of approach

This section sets out the counterfactual electricity use associated with backing up data in an server room located in another building. The representative scenario is a 100 GB back-up copied from a primary server to a dedicated back-up server, where it is stored for one year. Electricity consumption within the server rooms is estimated across three components: downloading from the primary server, uploading to the back-up server and long-term storage. Results are calculated for low, medium and high scenarios and are then adjusted for the higher

[74] See Annex 2 at DESNZ (2024) "Non-Domestic Building Stock in England and Wales" [online
[75] Low figure from University of Dundee estimates [online]. High figure from MabelSpce estimates [online]. The medium scenario is their average.

electricity overheads of a local server room using a PUE uplift. Electricity used in internet transmission between the two buildings is also taken into account.

The modelling uses the same processes as for data centres (“uploading”, “long-term storage”, “downloading”), but converts them to a local-server setting using: (i) an uplift to the energy value used for long term storage to reflect lower utilisation in local server rooms and less energy-efficient hardware; (ii) estimates for uploading and downloading energy reflecting lower capacity utilisation in local server rooms; and (iii) a higher PUE for the local server room.

Long-term storage

The long-term storage component captures the electricity required to store data on persistent storage so that it remains available for later retrieval. As set out in Section 3.2.1, our model already includes values for the electricity used in long-term storage at data centres, specified in kWh per GB stored per year. Two percentage uplifts are then applied to reflect a local server room context.

First, we apply an uplift to reflect lower capacity utilisation (i.e. more spare capacity) in local server rooms. Data centres are assumed to operate storage closer to an efficient fill level (around 80 per cent full), whereas local server rooms are assumed to run at materially lower utilisation (around 45 per cent full).[76,77] Lower utilisation increases electricity per GB stored because similar always-on storage electricity consumption is spread over fewer stored GB. To calculate the uplift, we use range estimates for storage utilisation at data centres and local servers. By comparing these utilisation rates, we derive low, medium and high uplifts to electricity consumption of 42 per cent, 78 per cent and 150 per cent.

Second, we apply an uplift to reflect less energy-efficient storage hardware in local server rooms relative to the specialised, high-capacity drives commonly used in data centres. For example, a typical HDD in a data centre might have 20TB of storage while consuming idle power 4W, while an HDD in a local server is assumed to consume around 8W of power with 10TB of capacity.[78] For robustness, our low, medium, and high figures are computed by comparing the W/TB ratio of a 21TB data centre HDD from the previous report with the minimum, average, and maximum W/TB ratios from a selection of standard 10TB HDDs.

Applying these uplifts yields an adjusted local-server long-term storage intensity.

[76] VMware (2019) “vSAN Capacity Management and Monitoring Part 2” [Online]
[77] Computerworld (2010) “A waste of space: Bulk of drive capacity still underutilized” [Online]
[78] [Data centre HDD] [Local server HDD]

Table 4.3: Electricity consumption of storage devices in server rooms

	Units	Low	Medium	High
Electricity used in long-term storage at data centres	kWh per GB stored per year	0.00294	0.00625	0.010667
Percentage uplift due to lower capacity utilisation	Percentage	42	78	150
Percentage uplift due to less energy-efficient hardware in local server rooms	Percentage	97	139	168
Electricity used in long-term storage in local server rooms	kWh per GB stored per year	**0.0082**	**0.0266**	**0.0715**

Source: Europe Economics analysis.

Internet transmission

As the back-up server counterfactual assumes that the back-up server is located off-site and that the data is transferred over the internet, electricity is also consumed in internet transmission. We apply the same per-GB energy intensities as we use for internet transmission in the digital use cases. This approach reflects that, in both cases, the dominant drivers of internet transmission energy are the wide-area network and associated network equipment.

Uploading and downloading

For the server back-up counterfactual, we estimate uploading and downloading electricity consumption using a high-level allocated-average, IT-only approach. The scope includes only internal networking equipment involved in transferring data to the server, such as access switches and any router or firewall performing internal routing. Server and client compute, WAN or internet transmission outside the office, facility overheads (i.e. no PUE uplift), and electricity associated with storing the data once uploaded are excluded.

The estimate combines two inputs over the same period: (i) the average power draw of the in-scope networking equipment (from metering where available, otherwise from typical/standby specifications and engineering judgement[79]) and (ii) the realised average network throughput. Electricity intensity is calculated as average power divided by the volume of data transferred. The main driver of the result is utilisation: office networks often draw much of their power regardless of load, so low realised throughput can yield a higher kWh/GB even with modest equipment power.

Based on these assumptions, we adopt a central estimate of **0.0003 kWh/GB** with a plausible range. The lower bound (**0.00001 kWh/GB**) reflects a short wired path with high utilisation (around 4 Gbit/s realised throughput). The upper bound (**0.005 kWh/GB**) reflects a longer path with low utilisation and less efficient equipment (around 13 Mbit/s realised throughput).

[79] An example of in-scope networking equipment includes access switches [online], and routers [online]

PUE

The upload, storage, and download elements represent IT electricity consumed within server rooms. To estimate total electricity attributable to the server-room environment, we apply a PUE factor to capture supporting overheads such as cooling, lighting, and power distribution. The literature indicates that local server room have a higher PUE than a data centre, consistent with smaller-scale and less optimised facilities.[80] The PUE of the physical server room is assumed to be 1.4 in the low scenario, 1.75 in the medium scenario and 2.1 in the high scenario.[81]

Table 4.4: PUE of a physical server room

	Low	Medium	High
PUE of a physical server room	1.4	1.75	2.1

Source: Berkeley Lab (2013)

[80] Berkeley Lab (2013) "Energy Efficiency of Small Data Centers" [Online]
[81] Berkeley Lab (2013) "Energy Efficiency of Small Data Centers" [Online]

5 Results for Use Cases

In this section of the report, we present the results of our calculations for our ten use cases, in each case comparing electricity consumption under the digital approach with electricity consumption under the counterfactual.

5.1 Video streaming versus TV broadcasting

This use case assesses the electricity consumption involved in watching a video either through a streaming platform (digital scenario) or a standard DTT broadcasting service (physical scenario).

The representative scenario for this use case assumes a two-hour video, with a streaming-optimised data rate of 3.5 GB per hour,[82] resulting in a total file size of 7 GB per viewing. The video is assumed to be watched twice over a five-year period, meaning that a total of 14 GB of data is transmitted. The viewer is assumed to access the content using a smart TV connected to a fixed-line broadband network.

5.1.1 Digital scenario

This section sets out the estimated electricity use associated with video streaming using a digital platform. The analysis includes electricity use from three main sources: (i) data centre electricity consumption, including both IT and non-IT processes, (ii) internet transmission electricity consumption, and (iii) electricity consumed by the end-user's device.

Data centre

The delivery of streaming content involves several IT processes within data centres, including the uploading and processing of video content, long-term and RAM storage, and final downloading/streaming to the user. Each of these processes is characterised by a defined activity level and electricity intensity, with some processes classified as shared across users (e.g. uploading, processing, long-term storage) and others as user-specific (e.g. downloading/streaming, RAM usage).

The table below presents the activity volumes for each process under low, medium and high scenarios, along with their classification as shared or specific.

Table 5.1: Activity volumes and process classification for IT processes

[82] GoBrolly (2025) "The Amount of Data and Bandwidth Required for Streaming Video" [online]

Data process	Units	Low	Medium	High	Shared/ Specific	Basic of Estimate
Uploading to data centres	GB Uploaded	7	7	7	Shared	*File size of a 2hr movie (from representative scenario)*
Video processing	GB processed	7	7	7	Shared	*File size of a 2hr movie (from representative scenario)*
Long-term storage	GB stored	7	7	7	Shared	*File size of a 2hr movie (from representative scenario)*
RAM storage	GB-hours	2.8	2.8	2.8	Specific	*Size of video × % cached in RAM × number of viewings × video length*
General processing	CPU-hours	0	0	0	-	-
GPU processing	GPU-hours	0	0	0	-	-
AI training	FLOP	0	0	0	-	-
AI interface	FLOP	0	0	0	-	-
Downloading/Streaming	GB downloaded/ streamed	14	14	14	Specific	*Total file size of movie streamed over representative scenario (file size × number of viewings)*

Source: Europe Economics analysis.

Using these activity volumes and corresponding electricity intensities, the total electricity consumption associated with each IT process is calculated. These figures represent the total electricity usage, prior to allocating shared electricity across the user base.

Table 5.2: Electricity used in each IT process – TV streaming (kWh)

Data process	Low	Medium	High
Uploading to data centres	0.000035	0.000350	0.003500
Video processing	0.7	2.8	7.0
Long-term storage	0.1029	0.2188	0.3733
RAM storage	0.0011	0.0011	0.0011
Downloading/streaming	0.00007	0.00070	0.00700
Sub-total	**0.8041**	**3.0209**	**7.3850**

Source: Europe Economics analysis.

To estimate the electricity attributable to the representative user, we allocate shared electricity across the user base. For specific processes, the full electricity is allocated to the user. The table below presents electricity consumption attributable to the representative user.

Table 5.3: IT electricity attributable to each user (kWh)

IT process (attributable to user)	Low	Medium	High
Uploading to cata centres	2.8×10^{-11}	1.4×10^{-9}	2.8×10^{-8}
Video processing	5.6×10^{-7}	1.12×10^{-5}	5.6×10^{-5}
Long-term storage	8.2×10^{-8}	8.8×10^{-7}	3.0×10^{-6}
RAM storage	0.00105	0.001098	0.001142
Downloading/streaming	0.00007	0.00070	0.00700
Total IT electricity	**0.0011**	**0.0018**	**0.0082**

Source: Europe Economics analysis.

In addition to the IT electricity, we include an allocation for non-IT electricity used to support data centre operations. This is calculated by applying an uplift to the energy used in each IT process using PUE values. The table below summarises both IT and non-IT electricity components to give the total data centre electricity attributable to the user.

Table 5.4: Summary of data centre electricity consumption (kWh)

Component	Low	Medium	High
IT electricity	0.0011	0.0018	0.0082
Non-IT electricity	0.0001	0.0005	0.0041
Total data centre electricity	**0.0012**	**0.0024**	**0.0123**

Source: Europe Economics analysis.

Internet transmission

Electricity is consumed in transmitting the video content between the data centre and the end-user. This includes both the initial upload of content to the platform (e.g. by a content producer) and the streaming of the video to the user. Upload transmission is considered a shared process, while download transmission is specific to the user. Therefore, only a small proportion of the upload energy is attributed to the representative user, while the full energy for download is allocated to the representative user.

Table 5.5: Transmission electricity attributable to each user (kWh)

Component	Low	Medium	High
Upload transmission	1.17×10^{-7}	7.34×10^{-7}	1.76×10^{-6}
Download transmission	0.293	0.367	0.440
Total transmission	**0.2934**	**0.3668**	**0.4402**

Source: Europe Economics analysis.

End-user device

The end-user is assumed to watch the video using a smart TV connected via a fixed-line broadband connection. Electricity consumption is estimated by multiplying the smart TV’s assumed power draw by the total hours of use – in this case, 4 hours (2 viewings of a 2-hour video).

Table 5.6: Electricity use in end-user device (kWh)

Units	Low	Medium	High
Electricity used in device	0.18	0.40	0.60
Electricity used in set-top box (if used)	0	0	0
Total device electricity	**0.18**	**0.40**	**0.60**

Source: Europe Economics analysis.

Summary of electricity use

This section consolidates the estimated energy consumption associated with streaming a 2-hour video twice over a five-year period. The total electricity use accounts for data centre operations (both IT and non-IT), internet transmission, and electricity consumed by the end-user device.

As shown in the table below, internet transmission and end-user device usage each account for a sizeable share of the overall electricity consumption.

Table 5.7: Total electricity consumption for representative scenario (kWh)

Component	Low	Medium	High
IT electricity (data centres)	0.001	0.002	0.008
Non-IT electricity (data centres)	0.000	0.001	0.004
Internet transmission	0.293	0.367	0.440
End-user device	0.180	0.400	0.600
Total	**0.475**	**0.769**	**1.052**

Source: Europe Economics analysis.

It is also useful to express electricity consumption on a per-hour-of-viewing basis. Given that the user watches a total of 4 hours of video in this scenario, the following table shows the electricity footprint per viewing hour.

Table 5.8: Total electricity consumption for representative scenario (kWh)

Component	Low	Medium	High
IT electricity (data centres)	0.0003	0.0005	0.0021
Non-IT electricity (data centres)	0.0000	0.0001	0.0010
Internet transmission	0.0734	0.0917	0.1100
End-user device	0.0450	0.1000	0.1500
Total	**0.1187**	**0.1923**	**0.2631**

Source: Europe Economics analysis.

To illustrate the relative contributions of each component, the table below presents the percentage breakdown of total electricity consumption across the three scenarios.

Table 5.9: Breakdown of electricity consumption (%)

Component	Low	Medium	High
IT electricity (data centres)	0.24	0.24	0.78
Non-IT electricity (data centres)	0.02	0.07	0.39
Internet transmission	61.82	47.69	41.82
End-user device	37.92	52.01	57.01
Total	**100**	**100**	**100**

Source: Europe Economics analysis.

It can be more challenging for an electricity network to supply power to a single, central location than to supply power to decentralised locations. We treat electricity consumed within the data centre as centralised, while electricity used for transmission and end-user devices is decentralised. Thus, the share of electricity consumed in centralised locations ranges from 0.25 to 1 per cent.

5.1.2 Counterfactual

This section sets out the counterfactual electricity use associated with viewing live TV via digital terrestrial television (DTT). The analysis includes electricity associated with the terrestrial infrastructure (transmission network) and the end-user viewing device (TV).

Terrestrial infrastructure (transmission network)

Electricity associated with DTT terrestrial infrastructure is modelled using a fixed electricity intensity per viewing hour, taken from Ofcom/Carnstone's DTT "data centres and network transmission" estimate of 3.2 Wh per device hour. This intensity is multiplied by the representative viewing duration to obtain total terrestrial infrastructure electricity for the four-hour session.

Table 5.10: Electricity consumed by terrestrial infrastructure (kWh)

Component	Low	Medium	High
Electricity used by terrestrial infrastructure	0.0102	0.0128	0.0154

Source: Europe Economics analysis.

End-user device

Electricity associated with the TV is modelled using the assumed power draw during viewing. The assumed TV electricity use is 0.05/0.10/0.20 kWh per hour in the low/medium/high cases. These values are applied over the four-hour viewing session.

Table 5.11: Electricity consumed by end-user device (kWh)

Component	Low	Medium	High
TV	0.18	0.40	0.60

Source: Europe Economics analysis.

Total electricity use for representative scenario

Total electricity consumption for the four-hour viewing session is calculated by summing terrestrial infrastructure electricity and TV electricity.

Table 5.12: Total electricity for representative scenario (kWh) – TV broadcasting

Component	Low	Medium	High
Electricity used by terrestrial infrastructure	0.0102	0.0128	0.0154
Electricity used by end-user devices	0.1800	0.4000	0.6000
Total	**0.1902**	**0.4128**	**0.6154**

Source: Europe Economics analysis.

The component breakdown indicates that end-user equipment dominates total electricity consumption in the DTT counterfactual under the default assumptions. Terrestrial infrastructure accounts for 5.38 per cent of total electricity in the low case, 2.40 per cent in the medium case, and 1.75 per cent in the high case, with the remainder attributable to end-user devices.

Table 5.13: Breakdown of electricity consumption (%) – TV broadcasting

Component	Low	Medium	High
Electricity used by terrestrial infrastructure	5.38	3.10	2.50
Electricity used in end-user devices	94.62	96.00	97.50
Total	100	100	100

Source: Europe Economics analysis.

In this counterfactual, electricity associated with the terrestrial infrastructure is treated as centralised, while electricity consumed by the TV is treated as decentralised. Thus, the centralised share of total UK-attributable electricity use ranges from **1.5 to 5.5 per cent**.

5.1.3 Comparison of electricity consumption

The table below summarises the total electricity consumption for the representative scenario across low, medium, and high scenarios. Our calculations show that video streaming requires substantially more electricity than its traditional broadcasting counterpart. Under medium assumptions, the digital scenario results in over 0.76 kWh of electricity consumption – close to twice the 0.41 kWh associated with DTT.

Table 5.14: Total electricity consumption – Video streaming versus TV broadcasting (kWh)

Scenario	Low	Medium	High
Video streaming	0.475	0.769	1.052
TV broadcasting	0.190	0.413	0.615

Source: Europe Economics analysis.

The figure below illustrates the same comparison graphically. The lower and upper bounds reflect the low and high scenarios, respectively, with the central line indicating the medium scenario.

Figure 5.1: Comparison of electricity consumption - Video streaming versus TV broadcasting

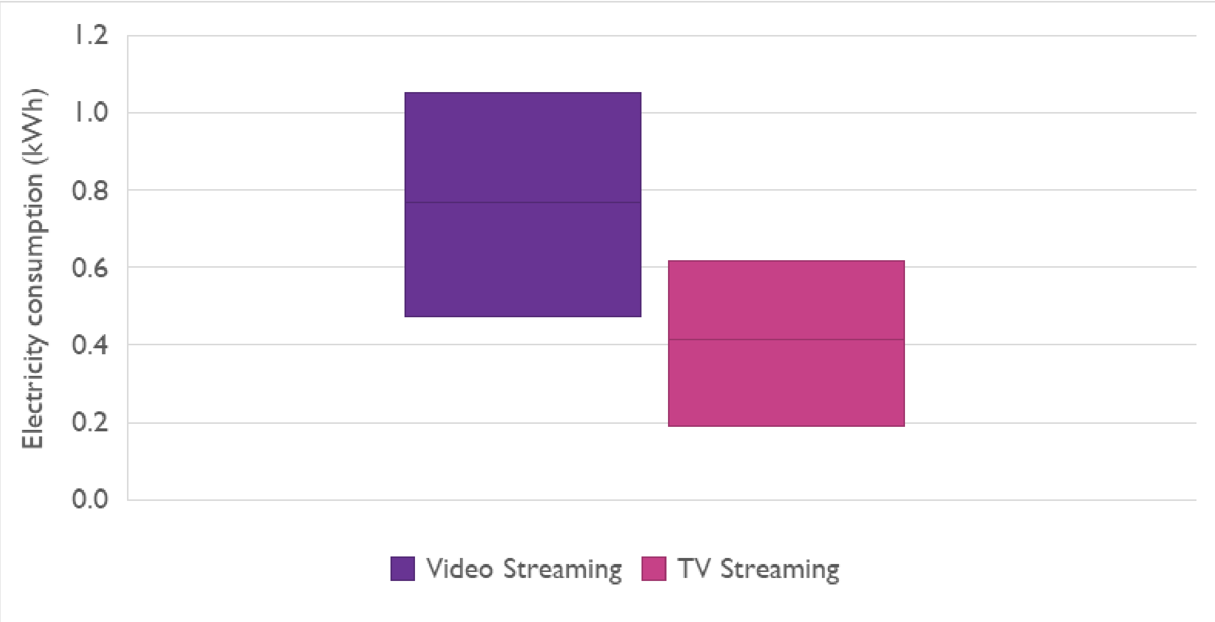

Source: Europe Economics analysis.

A caveat to this comparison is that video streaming and traditional TV broadcasting may differ in functionality of the service delivered. Streaming services often offer higher video resolutions (such as HD or 4K), on-demand viewing, and access to a wider catalogue of content, whereas DTT typically provides scheduled broadcast channels. To ensure comparability to the extent possible, the analysis focuses on a representative scenario in which the same piece of content is viewed twice over a five-year period. However, if users of video-streaming choose to watch higher-resolution video, this will typically require larger data volumes and therefore higher electricity consumption.

5.1.4 Electricity implications of a transition from DTT to OTT delivery

Given the above findings, this section explores the implications for UK electricity consumption of a transition from DTT to over-the-top (OTT) delivery. It considers both the system-wide impact of replacing current DTT viewing with video streaming and how the relative electricity intensity of DTT changes as user numbers decline. For consistency, estimates of current DTT electricity consumption and viewing patterns are drawn from the Ofcom-commissioned report on the carbon emissions of streaming and digital terrestrial television.[83] In this section, those figures are used solely to establish the scale of DTT use and associated electricity consumption; OTT electricity intensity is based on the video streaming assumptions set out earlier in this report.

Electricity impact of switching off DTT and replacing all viewing with OTT

We estimate the electricity implications of a DTT switch-off by assuming that all current DTT viewing is delivered instead via video streaming services. Electricity use by end-user devices such as television sets is excluded throughout. The comparison therefore captures only the electricity required for content delivery, including transmission networks, data centres and relevant customer equipment such as peripherals.

The starting point is the current scale of DTT use. Ofcom reports around 40.3 million DTT users in the UK, with total electricity consumption associated with delivering DTT services (excluding television sets) of approximately 240 million kWh per year. Average viewing is 3.65 hours per DTT household per day. Adjusting for average household size (2.36 people) implies around 1.55 hours of viewing per person per day, equivalent to 564 hours per person per year. In this scenario, that viewing volume is assumed to transfer in full to streaming.

Applying the video streaming electricity intensity set out above (0.09 kWh per hour, excluding the television set) to current DTT viewing volumes implies total electricity use of 2,048 million kWh per year if delivered via streaming. Against this, switching off DTT would avoid its existing 240 million kWh of annual electricity consumption. The net effect is therefore **an increase of 1,807 million kWh per year**, or around 1.8 TWh of additional electricity demand. This is equivalent to around 17 per cent of Sizewell B's 2025 electricity output of 10.4 TWh,[84] and around 0.6 per cent of total UK electricity demand, which was 319 TWh in 2024.[85]

This estimate reflects a stylised “full and immediate substitution” scenario in which all existing DTT viewing transfers directly to streaming. Under that assumption, total electricity use associated with content delivery increases materially, driven by the higher per-hour electricity intensity of video streaming relative to terrestrial broadcast transmission.

DTT user threshold at which per-viewer electricity usage matches OTT

[83] Ofcom (2022) “Carbon emissions of streaming and digital terrestrial television” [online]
[84] EDF (2026) “£1.2billion investment across 2026-28 to help boost reliability and output” [online]
[85] DESNZ (2025) “Chapter 5: Electricity” [online]

The full substitution scenario above assumes that all DTT viewing ceases at once. In practice, usage may decline gradually. It is therefore useful to consider how many DTT users would need to remain before the average electricity consumption per DTT viewer rises to match that of an OTT viewer.

To assess this, we distinguish between components of DTT electricity use that vary with user numbers and those that are largely fixed in the short to medium term:

Table 5.15: Yearly DTT electricity consumption by component (kWh)

Component	Value	Type of consumption
Peripherals	79,200,000	Variable
Customer premises equipment	88,800,000	Variable
Network transmission	72,000,000	Fixed

Source: Ofcom (2022); Europe Economics analysis.

Electricity used by peripherals and customer premises equipment is assumed to scale broadly with the number of active DTT viewers. When a household stops using DTT, the electricity associated with those devices falls away. By contrast, transmission network electricity reflects the electricity required to operate the broadcast infrastructure itself. In the absence of network closure or consolidation, this consumption does not decline proportionately with viewing and is therefore treated as fixed.

As user numbers fall, total variable electricity declines in line with those users. However, the fixed transmission electricity remains constant and must be borne by a smaller viewer base. This causes the average electricity consumption per remaining DTT viewer to rise.

Using annual viewing of 564 hours per person and an OTT electricity intensity of 0.09 kWh per hour, annual OTT electricity consumption per viewer is approximately 51 kWh. Solving for the number of DTT users at which average DTT electricity per viewer reaches this level implies a remaining DTT user base of around **1.5 million people**.

The figure below illustrates this relationship. The upward-sloping curve shows average DTT electricity per user as the user base contracts, while the horizontal line represents OTT consumption per viewer. The intersection occurs at roughly 1.5 million users. Only at this very low level of remaining DTT usage does the electricity consumption of DTT converge with the electricity consumption of streaming on a per-viewer basis; at any materially higher level, DTT remains less electricity-intensive per user.

Figure 5.2: DTT versus OTT annual electricity consumption per user (kWh)

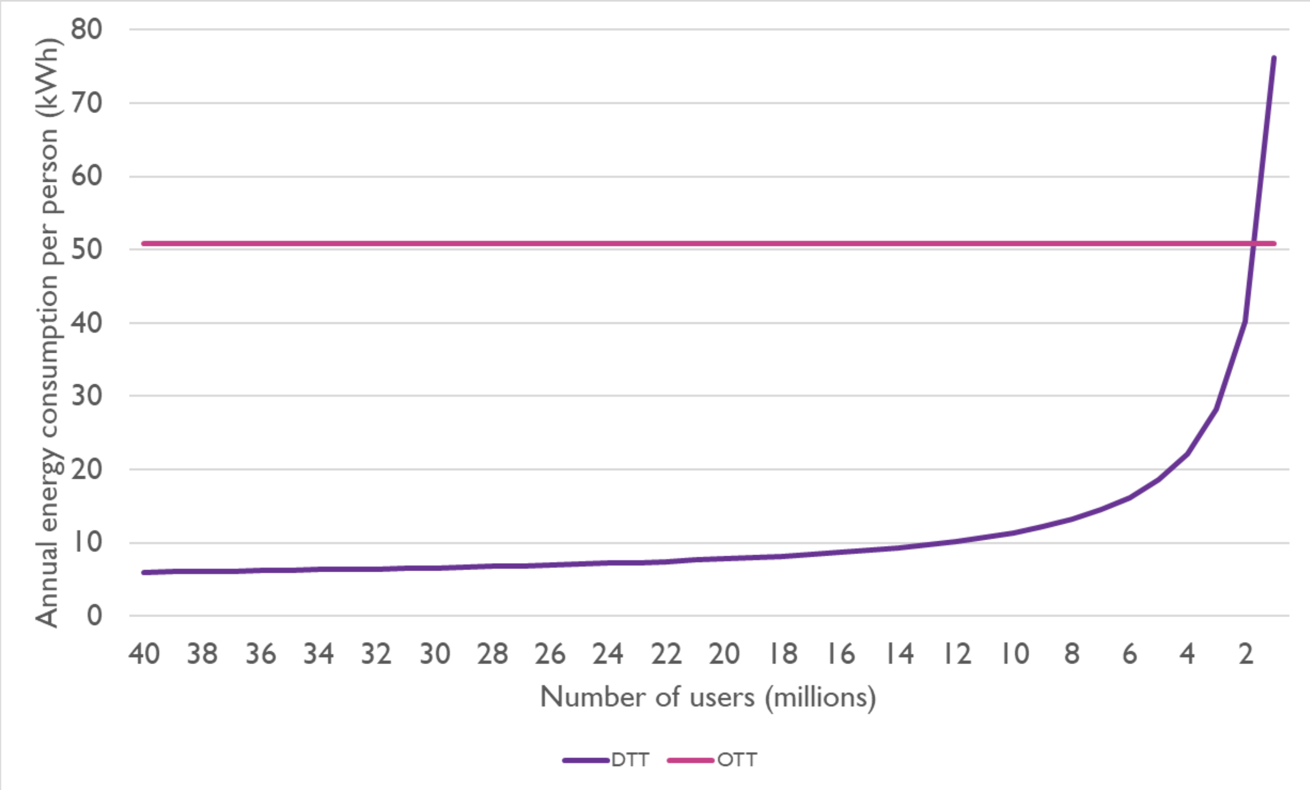

Source: Europe Economics analysis.

5.2 Cloud gaming versus gaming on own computer

This use case evaluates the electricity consumption associated with cloud gaming (digital scenario) or gaming through a CD on a gaming console (physical scenario). The representative scenario is intended to be a like-for-like comparison: we assume the **same single-player game** is played in both cases, at a **broadly comparable level of quality**, with similar resolution, frame rate and graphical fidelity. The scenario is intended to reflect a **mid-range game experience**, rather than either a very simple casual game or a highly demanding title played at the highest graphical settings.

The representative scenario that we use assumes three possibilities for the total size of the game (20 GB in the low, 50 GB in the medium, and 80 GB in the high scenario), which measures the one-off data transferred in the gaming file upload to the data centre. This one-off value is also relevant for the data registered on the gaming disc and electricity used for installation of that game on the end-user device.

We assume a two-hour gaming session, with a total of 8 GB/hour of data transferred, including both from the user to the data centre and from the data centre to the user. This results in the total size of data transferred of 16 GB per gaming session.

We assume a 10-year period for our analysis, based on the approximate lifetime of a standard gaming CD. We further assume the game is played 10 times during that period.

Cloud gaming also features two additional processes that consume electricity: general processing (responsible for running the game engine and logic process) and GPU processing (responsible for processing game visuals in the data centre, but remaining a separate process from video processing that remains responsible for streamlining the video to users). General processing and GPU processing are measured in CPUs and GPUs per hour respectively. In our representative scenario we assume 1.6 CPUs and 0.375 GPUs are used per hour of gaming,[86] resulting in 3.2 CPU hours and 0.75 GPU hours per gaming session.

For the counterfactual, the representative scenario assumes that the game does not require additional online updates after installation. In practice, incorporating updates would require assumptions on update frequency and size (for example, as a share of the original game download) and would also imply the use of data centres and network transmission for the physical scenario. To preserve a clear distinction between the digital approach and the counterfactual and to avoid introducing additional uncertainty, such updates are excluded from the representative scenario.

The user accesses the game using a smart TV with a gaming controller connected to a fixed-line broadband network. In the physical counterfactual, a range of end-user devices may be used. For the low estimate of electricity consumption, we assume the game is played on a gaming console connected to a television. For the high estimate of electricity consumption, we assume the game is played on a gaming PC. The medium estimate is calculated as the midpoint between the electricity consumption associated with these two cases.

A caveat is that electricity consumption may differ under other gaming settings. In particular, **multiplayer gaming** may involve different data-transfer requirements from single-player gaming, and games played at materially different quality settings may also have different electricity implications. These effects are not modelled here, as the representative scenario is defined to support a like-for-like comparison between cloud and local gameplay.

5.2.1 Digital scenario

This section sets out the estimated electricity use associated with cloud gaming. The analysis includes electricity use from three main sources: (i) data centre electricity consumption, including both IT and non-IT processes, (ii) internet transmission, and (iii) the end-user's device.

Data centre

[86] Based on the GPU use and estimates that each user is assigned approximately 25-50 per cent of 1 full GPU/hour found in Nvidia (2021) report.

The delivery of streaming content involves several IT processes within data centres. The table below presents the activity volumes for each process under low, medium and high scenarios, along with their classification as shared or specific processes.

Table 5.16: Activity volumes and process classification for IT processes

Data Process	Units	Low	Medium	High	Shared/ Specific	Basic of Estimate
Uploading to data centres	GB Uploaded	20	50	80	Shared	*Size of the game file*
Upload to Data Centres (during gaming)	GB Uploaded	16	16	16	Specific	*Size of gaming instructions sent over 10 gaming sessions*
Long-term storage	GB stored	20	50	80	Shared	*Size of the game file*
RAM storage	GB-hours	32	32	32	Specific	*Size of game × % cached in RAM × number of games × gameplay length*
General processing	CPU-hours	32	32	32	Specific	*CPU per gaming session × number of gaming sessions*
GPU processing	GPU-hours	7.5	7.5	7.5	Specific	*GPU per gaming session × number of games*
AI training	FLOP				-	
AI interface	FLOP				-	
Downloading/Streaming	GB downloaded/ streamed	144	144	144	Specific	*File size of 2h game accessed 10 times*

Source: Europe Economics analysis.

Using these activity volumes and corresponding energy intensities, the total electricity consumption associated with each IT process is calculated. These figures represent the total electricity usage, prior to allocating shared electricity across the user base.

Table 5.17: Electricity used in each IT process – cloud gaming

Data Process	Low	Medium	High
Uploading to data centres	0.0009	0.0105	0.1200
Long-term storage	0.0120	0.0125	0.0131

RAM storage	0.120	0.800	7.200
General processing	0.750	0.938	1.125
GPU processing	0.0009	0.0105	0.1200
Downloading/streaming	0.00072	0.00720	0.07200
Sub-total	**1.47**	**4.89**	**17.06**

Source: Europe Economics analysis.

To estimate the electricity attributable to the representative user, we allocate the electricity consumed by shared processes across the assumed user base. For specific processes, the full electricity is allocated to the user. The table below presents the electricity attributed to the representative user.

Table 5.18: IT electricity attributable to each user (kWh)

IT process (attributable to user)	Low	Medium	High
Uploading to data dentres	0.0008	0.0080	0.0800
Long-term storage	7.5×10^{-8}	4.8×10^{-7}	1.6×10^{-6}
RAM storage	0.012	0.013	0.013
General processing	0.120	0.800	7.200
GPU processing	0.750	0.938	1.125
Downloading/streaming	0.00072	0.00720	0.07200
Total IT electricity	**0.884**	**1.765**	**8.490**

Source: Europe Economics analysis.

In addition to the IT electricity, we include an allocation of non-IT electricity used to support data centre operations using PUE values. The table below summarises both IT and non-IT electric components to give the total data centre electricity consumption attributable to the user.

Table 5.19: Summary of data centre electricity consumption (kWh)

Component	Low	Medium	High
IT electricity	0.884	1.765	8.490
Non-IT electricity	0.088	0.530	4.245
Total data centre electricity	**0.972**	**2.295**	**12.735**

Source: Europe Economics analysis.

Internet transmission

Electricity is consumed in transmitting the video game content between the data centre and the end-user. This includes the initial upload of game content to the platform (e.g. by a game producer), which is a shared process, and uploads to the data centre during game play, which is specific to the user. Only a small proportion of data transfer during game play is assigned to upload transmission, with most of the user-specific internet transmission coming from the download process, given the data requirements of streaming the video game.

Table 5.20: Transmission electricity attributable to each user (kWh)

Component	Low	Medium	High
Upload transmission (uploading the game)	5.37×10^{-8}	2.01×10^{-7}	4.84×10^{-7}
Upload transmission (playing the game)	0.33536	0.4192	0.50304
Download transmission	3.01824	3.7728	4.52736
Total transmission	**3.3536**	**4.1920**	**5.0304**

Source: Europe Economics analysis.

End-user device

The end-user is assumed to play the game using a smart TV connected via a fixed-line broadband connection. Electricity consumption is estimated by multiplying the smart TV's assumed power draw by the total hour of use – in this case, 20 hours (10 gaming sessions, each lasting 2 hours).

Table 5.21: Electricity used in end-user devices (kWh)

Units	Low	Medium	High
Electricity used in device	**0.90**	**2.00**	**3.00**

Source: Europe Economics analysis.

Summary of electricity use

This section consolidates the estimated electricity consumption associated with playing a video game for two hours 10 times over a ten-year period. The total electricity use accounts for data centre operations (both IT and non-IT), internet transmission, and electricity consumed by the end-user device.

Table 5.22: Total electricity consumption for representative scenario (kWh)

Component	Low	Medium	High
IT electricity (data centres)	0.884	1.765	8.490
Non-IT electricity (data centres)	0.088	0.530	4.245
Internet transmission	3.354	4.192	5.030
End-user device	0.900	2.000	3.000
Total	**5.225**	**8.487**	**20.765**

Source: Europe Economics analysis.

It is also useful to express electricity consumption on a per-hour-of-gaming basis. Given that the user plays the game a total of 20 hours of video in our representative scenario, the following table shows the electricity footprint per gaming hour.

Table 5.23: Total electricity consumption for representative scenario (kWh)

Component	Low	Medium	High
IT electricity (data centres)	0.04	0.09	0.42
Non-IT electricity (data centres)	0.00	0.03	0.21
Internet transmission	0.17	0.21	0.25
End-user device	0.05	0.10	0.15
Total	**0.26**	**0.42**	**1.04**

Source: Europe Economics analysis.

To illustrate the relative contributions of each component, the table below presents the percentage breakdown of total electricity consumption across the three scenarios.

Table 5.5.24: Breakdown of electricity consumption (%)

Component	Low	Medium	High
IT electricity (data centres)	17	21	41
Non-IT electricity (data centres)	2	6	20
Internet transmission	64	49	24
End-user device	17	24	14
Total	**100**	**100**	**100**

Source: Europe Economics analysis.

A significant contribution to the total electricity consumption under the cloud gaming scenario comes from the electricity consumed in the data centre and the electricity consumed in transmission of data using gameplay. Thus, the share of electricity consumed in centralised locations ranges from 19 to 58 per cent, with a higher proportion of electricity consumed at data centres (a centralised location) in the high scenario.

5.2.2 Counterfactual

This section sets out the estimated energy use associated with playing a video game using a CD disc. The analysis includes energy used in the manufacturing, packaging, storage, transportation and retailing of the disc, as well as the energy used in installing and playing the game on the end-user's device. Range estimates are presented reflecting variations in production processes, packaging energy, transportation distances from manufacturing sites to end users, and the energy-intensity of the end-user device.

Given the level of detail involved in modelling the production, storage and transport stages, the discussion in this section focuses on the key results. Full details of the underlying assumptions, calculation steps and detailed results, including the treatment of UK and overseas manufacturing and decarbonised variants, are provided in Appendix A.

Summary of energy use

Electricity

We only consider electricity consumption within the UK, which will be affected by whether the gaming disc is manufacturing in the UK or abroad. We therefore consider a scenario in which the gaming disc is manufactured domestically and a scenario in which it is manufactured overseas.

Since all transportation energy is assumed to be diesel or marine diesel, the electricity consumption in this scenario is comprised of the electricity portion of the manufacturing energy consumption (if in the UK), the electricity portion of storage and retailing energy consumption and the end-user device energy consumption.

Electricity consumption for the representative scenario ranges from 3.5 kWh to 10.9 kWh, with device electricity efficiency accounting for the majority of this figure.

Table 5.25: Total energy for representative scenario (kWh) (Manufactured within the UK) – Electricity

Component	Low	Medium	High
Electricity for manufacturing, packaging, storage, and retailing	0.445	0.567	0.689
Electricity for end-user device	3.081	6.650	10.250
Total electricity for representative scenario	**3.526**	**7.217**	**10.939**

Source: Europe Economics analysis.

The table below presents the percentage breakdown of electricity consumption.

Table 5.26: Breakdown of energy consumption (%) (Manufactured within the UK) – Electricity

Component	Low	Medium	High
Electricity for manufacturing, packaging, storage, and retailing	13	8	6
Electricity for end-user device	87	92	94

Source: Europe Economics analysis.

In the UK manufacturing scenario, we classify electricity used in production, warehousing, and retail as centralised and electricity for end-user devices as decentralised. Thus, the share of electricity consumption occurring in centralised locations ranges from **6 to 13 per cent**.

All fuel

The *“all fuel”* scenario accounts for all fuel types (electricity, gas, diesel, and marine diesel) consumed across manufacturing, storage, transportation, retailing and end-user device usage. All transportation energy is assumed to be diesel or marine diesel, the manufacturing storage and retailing energy includes electricity, gas and diesel consumption, while end-user device energy remains entirely electricity-based.

The table below presents the total energy consumption, with energy from end-user devices being the dominant energy source. Gas contributes to manufacturing, storage and retailing processes, but its share is relatively small. Diesel is mainly used for road transportation, making its impact minor due to the high number of gaming CDs transported per vehicle. Marine diesel consumption is zero, as UK manufacturing does not involve maritime shipping.

Table 5.27: Total energy for representative scenario (kWh) (Manufactured within the UK) – All Fuel

Component	Low	Medium	High
Energy for manufacturing, packaging, storage, and retailing - electricity	0.445	0.5667	0.6892
Energy for manufacturing, packaging, storage, and retailing - other fuel	0.076	0.0994	0.1233
Energy for transportation - road - other fuel	0.003	0.0078	0.018
Energy for end-user device - electricity	3.081	6.650	10.250
Total energy for representative scenario	**3.605**	**7.324**	**11.080**

Source: Europe Economics analysis.

The table below provides a breakdown of energy consumption across different components. Energy from end-use devices remains the dominant energy source, accounting for over 90 per cent of total energy consumption. Electricity for manufacturing, packaging, storage and retailing is the second largest component, but is small in comparison to end-user device energy – between 6 and 12 per cent of the total. Other fuels (mainly gas) used in manufacturing and storage account for only a small fraction, highlighting their limited role. Diesel for road transport remains negligible, reinforcing the fact that logistics have a minor impact on the total energy footprint.

Table 5.28: Breakdown of energy consumption (%) (Manufactured within the UK) – All Fuel

Component	Low	Medium	High
Energy for manufacturing, packaging, storage, and retailing - electricity	12	8	6
Energy for manufacturing, packaging, storage, and retailing - other fuel	2	1	1

Energy for transportation - road - other fuel	0	0	0
Energy for end-user device- electricity	85	91	93

Source: Europe Economics analysis.

Under the UK manufacturing scenario (all fuel), we classify energy used in production, warehousing and retailing as centralised. Transport fuel and end-user device electricity are treated as decentralised. The proportion of total energy consumption that occurs in centralised locations ranges from 7 to 14 per cent.

Decarbonised scenario

In the decarbonised case, all processes, including transport and heating, are assumed to be powered by low-carbon electricity. The figures below therefore reflect the total grid-relevant energy demand under a fully electrified supply chain.

The table below presents total electricity consumption under the decarbonised scenario. End-user device electricity remains the largest contributor.

Table 5.29: Total electricity for representative scenario (kWh) (UK Manufactured) – Decarbonisation

	Low	Medium	High
Electricity for manufacturing, packaging and storage	0.521	0.666	0.812
Electricity for transportation - road	0.001	0.002	0.005
Electricity for end-user device	3.081	6.650	10.250
Total electricity for representative scenario	**3.603**	**7.318**	**11.067**

Source: Europe Economics analysis.

To illustrate the relative contributions of each component, the table below presents the percentage breakdown of total electricity consumption across the three scenarios.

Table 5.30: Breakdown of electricity consumption (%) (UK Manufactured) – Decarbonisation

	Low	Medium	High
Electricity for manufacturing, packaging and storage	14	9	7
Electricity for transportation - road	0	0	0
Electricity for end-user device	86	91	93
Total electricity for representative scenario	100	100	100

Source: Europe Economics analysis.

Under the decarbonised scenario with UK manufacturing, we classify electricity used in production, warehousing and retailing as centralised. Transport fuel and end-user device electricity are treated as decentralised. The proportion of total electricity consumption that occurs in centralised locations ranges from **7 to 14 per cent**.

Overseas Manufacturing

For products manufactured overseas, we only consider the energy consumption directly associated with UK-based activities in the analysis. This consists of warehousing, retail operations, and the domestic portion of transportation, which includes road freight from the UK port to warehouses and final retail distribution. Additionally, the energy from maritime transport is allocated between the UK and overseas. The energy used in raw material processing, manufacturing, and packaging are excluded as it remains outside the UK's energy footprint, as these processes occur abroad. The detailed modelling for this case is presented in the Appendix.

5.2.3 Comparison of energy consumption

This section compares the electricity consumption associated with cloud gaming (digital scenario) and a game being played on the user's own device (counterfactual).

Total Energy Use Across Scenarios (Electricity Only)

The table below presents electricity consumption in the UK in each scenario. For cloud gaming, this includes electricity used by data centres, internet transmission, and the end-user's device. For gaming on the user's own device, it includes electricity consumed in manufacturing, packaging, warehousing and retailing gaming CDs, and the electricity used by the end-user's device.

Electricity consumption is higher in the cloud gaming scenario across all scenarios. This reflects the additional electricity required for remote computation in data centres and continuous data transmission during gameplay. By contrast, gaming on the user's device using a gaming CD avoids these data centre workloads, with electricity consumption driven primarily by the end-user device and, to a lesser extent, by physical supply-chain activities.

For games played on an end-user's device using a CD manufactured overseas, only electricity consumed within the UK (such as warehousing, retail, and device use) is included, resulting in lower reported UK electricity consumption than for UK manufacture.

Decarbonisation slightly increases estimated electricity use for gaming on a user's own device, as fossil fuel inputs are replaced with electricity, but electricity consumption still remains below that of cloud gaming.

Table 5.31: Total electricity consumption attributable to the UK – cloud gaming versus gaming on user's device (kWh)

Scenario	Low	Medium	High
Cloud gaming (digital)	5.23	8.49	20.77
CD gaming (UK manufacture)	3.53	7.22	10.94
CD gaming (overseas manufacture)	3.16	6.75	10.37
CD gaming (UK manufactured, decarbonised)	3.60	7.32	11.07
CD gaming (overseas manufactured, decarbonised)	3.18	6.78	10.41

Source: Europe Economics analysis.

Total Energy Use Across Scenarios (All Fuels)

The table below includes electricity as well as other fuels such as gas and diesel. As with the electricity-only results, energy use in the cloud gaming scenario is entirely electricity-based, whereas gaming on a user's own device involves a broader mix of fuels, particularly in manufacturing and transport.

Including other fuels increases total energy consumption for gaming on the user's own device, especially when gaming CDs are manufactured in the UK. For overseas manufacture, a significant share of energy use occurs outside the UK and is therefore excluded, resulting in lower UK-attributable energy consumption. However, even when taking account of all fuels, cloud gaming continues to exhibit higher total UK energy consumption than gaming on the user's own device across all scenarios.

Decarbonised scenarios do not materially affect the comparison between cloud gaming and gaming on the user's device, with cloud gaming continuing to involve higher UK electricity consumption.

Table 5.32: Total energy consumption attributable to the UK (All Fuels, kWh)

Scenario	Low	Medium	High
Cloud gaming (digital)	5.225	8.487	20.765
CD gaming (UK manufacture)	3.605	7.324	11.080
CD gaming (overseas manufacture)	3.185	6.786	10.427
CD gaming (UK manufactured, decarbonised)	3.603	7.318	11.067
CD gaming (overseas manufactured, decarbonised)	3.183	6.778	10.408

Source: Europe Economics analysis.

The figure below compares electricity consumption for cloud gaming and gaming on a console or gaming laptop in the representative scenario. The results for the counterfactual represent a UK manufactured CD.

Figure 5.3: Comparison of electricity consumption - cloud gaming versus gaming on user's device

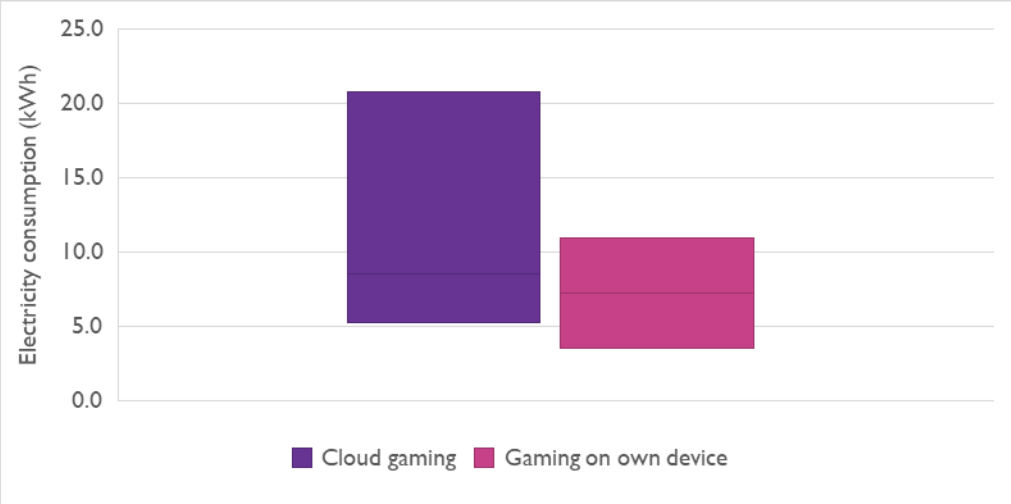

Source: Europe Economics analysis.

Conclusion

Under all scenarios, cloud gaming results in higher UK energy consumption than gaming on a user's device using a game installed from a physical disc. This is driven by the additional electricity required for data centre processing and network transmission during gameplay. While gaming using a gaming CD involves energy-intensive supply chains, particularly in manufacturing, much of this energy use occurs outside the UK when games are produced overseas. Overall, the results indicate that, for equivalent gameplay, cloud gaming involves higher UK electricity consumption perspective gaming on an owned device.

Decarbonisation Readiness

We also investigate the proportion of total energy consumption supplied in the form of electricity, which we refer to as a "Decarbonisation Readiness Index" (DRI). This index is calculated as the share of electricity within total energy use (including electricity, gas, and liquid fuels) and provides an indication of how readily each scenario could benefit from electricity grid decarbonisation. In the scenario in which gaming discs are manufactured in the UK, the DRI ranges from 97.81 to 98.73 per cent across scenarios. If manufacturing occurs overseas, UK-attributable electricity use accounts for an even higher share of total energy consumption in the UK, with a DRI ranging from 99.27 to 99.48 per cent. In both cases, electricity represents the dominant energy source, though some residual use of gas and diesel remains in UK-based manufacturing, warehousing, transportation and retailing.

By contrast, the digital gaming scenario is entirely electricity-based, resulting in a DRI of 100 per cent.

Table 5.33: DRI for gaming on own device (% of energy from electricity)

	Low	Medium	High
Digital gaming	100.00	100.00	100.00
Counterfactual (UK manufactured)	97.81	98.54	98.73
Counterfactual (overseas manufactured)	99.27	99.48	99.46

Source: Europe Economics analysis.

5.3 Online education platform versus short in-person course

This use case assesses the electricity consumption involved in participating in a short online course (digital approach) or a short course delivered in person (counterfactual). The analysis is grounded in a representative scenario that sets out the details of a short course, allowing for consistent comparison between both approaches.

For the digital use case, we estimate electricity consumption per student participating in a 12-hour online course delivered via a series of pre-recorded videos. We assume this value from estimates that most short online courses last between 5 and 20 hours of total material, split into shorter lessons[87] – we decided on an approximate average of these two values. We assume 3 GB of data transferred per hour of training video, streaming optimised. We use three years as the length of our representative period for analysis, which we take to represent a typical period after which the training material would have to be changed or updated. We assume each student takes the learning module once during this period. This results in 36 GB of data transferred (3 GB/hour × 12 hours per course) during our representative scenario. We assume the training material is accessed via a desktop computer connected to the fixed line network broadband.

For the counterfactual, we estimate the energy consumption attributable to the student and the teacher. For the teacher, the energy consumption is distributed across all students. We assume the class size to be 5 to 15 students, a typical size for a short summer course.[88]

5.3.1 Digital scenario

Here we set out the estimated electricity use associated with an online course. The analysis includes electricity use from three main sources: (i) data centre electricity consumption, including both IT and non-IT processes, (ii) internet transmission, and (iii) the end-user's device.

[87] UTeach (2025) "Course Length Guide: How to Get it Just Right" [online]
[88] University of Oxford "Inspiring Oxford Frequently Asked Questions" [online]

Data centre

The delivery of training content involves several IT processes within data centres: uploading to data centre, video processing, long-term storage, RAM storage and downloading/streaming. Some of these processes are classified as shared across users (e.g. uploading, processing, long-term storage) and others as user-specific (e.g. downloading/streaming, RAM usage).

The table below presents the activity volumes for each process under low, medium and high scenarios, along with their classification as shared or specific.

Table 5.34: Activity volumes and process classification for IT processes

Data Process	Units	Low	Medium	High	Shared/ Specific	Basic of Estimate
Uploading to data centres	GB Uploaded	36	36	36	Shared	*File size of a 12h training video*
Video processing	GB processed	36	36	36	Shared	*File size of a 12h training video*
Long-term storage	GB stored	36	36	36	Shared	*File size of a 12h training video (representative scenario)*
RAM storage	GB-hours	43.2	43.2	43.2	Specific	*Size of video × % cached in RAM × number of viewings × video length*
General processing	CPU-hours	0	0	0	-	-
GPU processing	GPU-hours	0	0	0	-	-
AI training	FLOP	0	0	0	-	-
AI interface	FLOP	0	0	0	-	-
Downloading/Streaming	GB downloaded/ streamed	36	36	36	Specific	*File size of a 12h training video (file size × times watched)*

Source: Europe Economics analysis.

Using these activity volumes and corresponding energy intensities, the total electricity consumption associated with each IT process is calculated. These figures represent the total electricity usage, prior to allocating the electricity associated with shared processes across the user base.

Table 5.35: Electricity used in each IT process – short online course (kWh)

Data Process	Low	Medium	High
Uploading to data centres	0.00018	0.00180	0.01800
Video processing	3.6	14.4	36
Long-term storage	0.318	0.675	1.152
RAM storage	0.016	0.017	0.018
Downloading/streaming	0.00018	0.00180	0.01800
Sub-total	**3.934**	**15.096**	**37.206**

Source: Europe Economics analysis.

To estimate the electricity attributable to the representative user, we apply allocation factors for shared processes (e.g. uploading to data centres, video processing, and long-term storage), based on the assumed user base. For specific processes such as RAM storage and downloading/streaming, the full electricity is allocated to the user. The table below presents these user-level allocations.

Table 5.36: IT electricity attributable to each user (kWh)

IT process (attributable to user)	Low	Medium	High
Uploading to data centres	1.2×10-8	6.0×10-7	7.5×10-6
Video processing	2.4×10-4	4.8×10-3	1.5×10-2
Long-term storage	2.1×10-5	2.3×10-5	4.8×10-4
RAM storage	0.016	0.017	0.018
Downloading/streaming	0.00018	0.00180	0.01800
Total IT electricity	**0.017**	**0.024**	**0.051**

Source: Europe Economics analysis.

In addition to the IT electricity consumption, we include an allocation for non-IT electricity used to support data centre operations. This is calculated by applying a proportional uplift to the electricity consumption associated with each IT process using PUE factors.

The table below summarises both IT and non-IT electric components to give the total data centre electricity attributable to the user.

Table 5.37: Summary of data centre electricity consumption (kWh)

Component	Low	Medium	High

IT electricity	0.017	0.024	0.051
Non-IT electricity	0.002	0.007	0.026
Total data centre electricity	**0.018**	**0.031**	**0.077**

Source: Europe Economics analysis.

Internet transmission

Electricity is consumed in transmitting the training video content between the data centre and the end-user. This includes both the initial upload of content to the platform (e.g. by a training content producer) and the streaming of the video to the user. Upload transmission is considered a shared process, while download transmission is specific to the user.

Table 5.38: Transmission electricity attributable to each user (kWh)

Component	Low	Medium	High
Upload transmission	5.03×10^{-5}	3.14×10^{-4}	4.72×10^{-4}
Download transmission	0.755	0.943	1.132
Total transmission	**0.755**	**0.944**	**1.132**

Source: Europe Economics analysis.

End-user device

The end-user is assumed to watch the video using a desktop computer. Electricity consumption is estimated by multiplying the desktop computer's assumed power draw by the total hours of use.

Table 5.39: Electricity used in end-user devices (kWh)

Units	Low	Medium	High
Electricity used in device	2.4	3.0	3.6
Electricity used in set-top box (if used)	0	0	0
Total device electricity	**2.4**	**3.0**	**3.6**

Source: Europe Economics analysis.

Summary of electricity use

This section consolidates the estimated electricity consumption associated with watching a 12-hour pre-recorded training content once over a three-year period. The total electricity use accounts for data centre operations (both IT and non-IT), internet transmission, and electricity consumed by the end-user device.

As shown in the table below, internet transmission and end-user device usage each contribute a sizeable share of the overall electricity, across all scenarios.

Table 5.40: Total electricity consumption for representative scenario (kWh)

Component	Low	Medium	High
IT electricity (data centres)	0.017	0.024	0.051
Non-IT electricity (data centres)	0.002	0.007	0.026
Internet transmission	0.755	0.944	1.132
End-user device	2.400	3.000	3.600
Total	**3.173**	**3.974**	**4.809**

Source: Europe Economics analysis.

To illustrate the relative contributions of each component, the table below presents the percentage breakdown of total electricity consumption across the three scenarios.

Table 5.41: Breakdown of electricity consumption (%)

Component	Low	Medium	High
IT electricity (data centres)	1	1	1
Non-IT electricity (data centres)	0	0	1
Internet transmission	24	24	24
End-user device	76	75	75
Total	**100**	**100**	**100**

Source: Europe Economics analysis.

We treat electricity consumed within the data centre as centralised, while electricity used for transmission and end-user devices is decentralised. Overall, the share of electricity consumed in centralised locations ranges from 1 to 2 per cent across the three scenarios.

5.3.2 Counterfactual

This section sets out the estimated electricity use associated with delivering a short training course via in-person teaching. We assume that the teaching takes place for 6 hours per day. We still assume the same duration of the course – 12 hours of teaching in total. With an allowance for breaks, this results in two days of in-person teaching for our representative course.

Under the counterfactual we estimate electricity use for different components: (i) electricity consumed in the classroom and (ii) distance-related transportation electricity consumption.

Electricity consumed in the classroom

To estimate the electricity consumed in the classroom during the duration of the course we multiply the days spent in the classroom (two days) by the electricity per day per student and the electricity per day for a teacher (divided across all students) in an educational building. The table below shows the classroom electricity estimates across the three scenarios.

Table 5.42: Electricity consumption in the classroom (kWh)

Component	Low	Medium	High
Electricity consumed in the classroom by student	3.74	5.54	7.70
Electricity consumed in the classroom by teacher (divided by number of students)	0.25	0.55	1.54
Electricity consumed in the classroom	**3.988**	**6.099**	**9.240**

Source: Europe Economics analysis.

Distance-related energy consumption

Our primary metric for distance-related energy consumption is electricity consumption for transportation. Two further views are then considered: (i) a fully decarbonised scenario where travel is electrified, and (ii) a total energy scenario including liquid fuels such as diesel.

Under the counterfactual, we assume that both the teacher and the student travel to and from the educational venue on each day of attendance. We therefore assume two days of travel per individual, with a return journey made on each day.

Electricity-based view (primary presentation)

In the baseline case, distance-related energy use is not electricity-based, as travel is assumed to take place using a diesel vehicle. Electricity consumption from travel is zero in this scenario.

Decarbonised

To reflect a fully electrified system, we assume that travel is undertaken using an electric vehicle. Energy use is calculated based on assumed travel distances and representative EV electricity consumption. The table below presents the resulting electricity consumption from commuting.

Table 5.43: Distance-related electricity consumption – decarbonised (kWh)

Component	Low	Medium	High
Electricity from commuting for student	2.048	3.072	4.301
Electricity from commuting for teacher (divided by number of students)	0.137	0.307	0.860
Electricity from commuting	**2.185**	**3.379**	**5.161**

Source: Europe Economics analysis.

Total energy (all fuels)

For completeness, we also estimate total energy use including liquid fuels. In this case, energy consumption is calculated based on travel distance, diesel vehicle fuel efficiency, and the energy content of diesel. The table below presents our results.

Table 5.44: Distance-related energy consumption (kWh)

Component	Low	Medium	High

Energy from commuting for student	9.36	13.48	18.61
Energy from commuting for teacher (divided by number of students)	0.25	0.55	1.54
Energy from commuting	**9.986**	**14.831**	**22.328**

Source: Europe Economics analysis.

Summary of energy use

We first present results for electricity consumption only. As travel in the baseline case is assumed to rely on diesel vehicles, electricity consumption is driven entirely by classroom use.

Table 5.45: Summary of electricity consumption for a short in-person course (kWh)

Component	Low	Medium	High
Electricity consumed in the classroom	3.988	6.099	9.240
Electricity from commuting	0	0	0
Total	**3.988**	**6.099**	**9.240**

Source: Europe Economics analysis.

In a decarbonised scenario, in which travel is electrified, distance-related electricity consumption also takes the form of electricity consumption.

Table 5.46: Summary of electricity consumption for a short in-person course – decarbonised (kWh)

Component	Low	Medium	High
Electricity consumed in the classroom	3.988	6.099	9.240
Electricity from commuting	2.185	3.379	5.161
Total	**6.172**	**9.478**	**14.401**

Source: Europe Economics analysis.

Under the decarbonised scenario, total electricity consumption increases relative to the baseline view, as travel energy now takes the form of electricity. However, distance-related energy consumption remains relatively small, accounting for around 33 per cent of total energy use, with energy consumption in the educational building accounting for the remaining 66 per cent.

For completeness, we also present total energy use including liquid fuels such as diesel.

Table 5.47: Summary of energy consumption for a short in-person course – all fuels (kWh)

Component	Low	Medium	High
Electricity consumed in the classroom	3.988	6.099	9.240
Energy from commuting	9.986	14.831	22.328
Total	**13.973**	**20.931**	**31.568**

Source: Europe Economics analysis.

We note that the majority of energy consumption in this view comes from commuting to the educational building, which accounts for approximately 71 per cent of total energy consumption across our three scenarios.

We assume there is no centralised process under this counterfactual – all energy use occurs in decentralised locations.

5.3.3 Comparison of electricity consumption

Taking a short course online results in different energy consumption from attending an equivalent in-person course. Given that we have assessed the energy per student for taking an equivalent course in either format, we can consistently compare energy consumption between the two cases.

Table 5.48: Total energy consumption for representative scenario (kWh)

	Low	Medium	High
Online course	3.173	3.974	4.809
In-person course	3.988	6.099	9.240
In-person course (decarbonised)	6.172	9.478	14.401
In-person course (all fuels)	13.973	20.931	31.568

Source: Europe Economics analysis.

The figure below illustrates the same comparison graphically.

Figure 5.4: Comparison of electricity consumption – online education versus in-person education

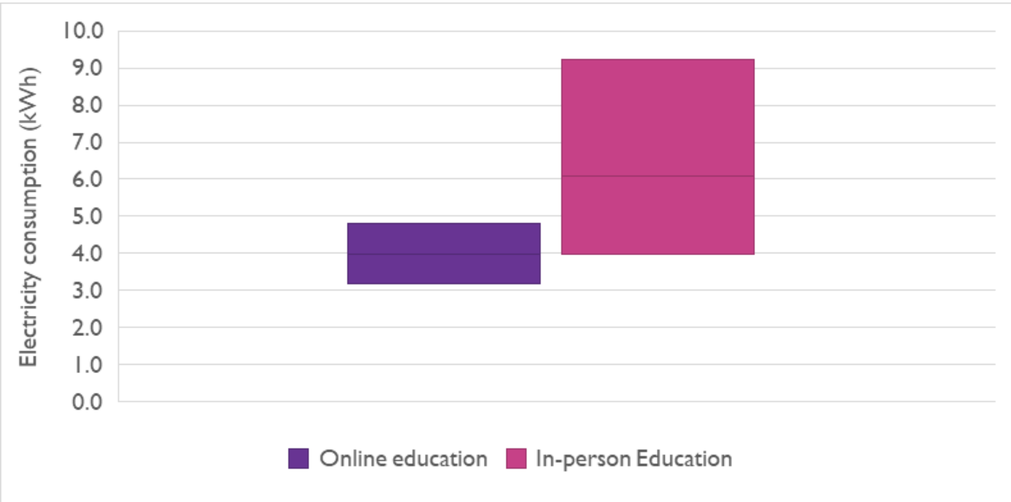

Source: Europe Economics analysis.

The in-person course shows higher energy consumption than the online alternative, reflecting the energy required within educational buildings, while travel does not contribute in our baseline view. When moving to a decarbonised scenario, total electricity use increases as travel is electrified, although the overall pattern remains unchanged, with and building-related electricity still the main driver of the difference. In contrast, when total energy use is considered, including liquid fuels, the in-person course becomes substantially more energy intensive, as commuting dominates overall consumption and is absent from the digital scenario.

A caveat to this comparison is that online and in-person courses may differ in the nature of the learning experience delivered. In-person teaching may provide benefits such as face-to-face interaction with instructors and peers, access to dedicated teaching facilities, and opportunities for practical or collaborative activities that may not be fully replicated in an online format. These differences mean that the two approaches may not always represent identical educational experiences. In some cases, replicating aspects of in-person learning in an online setting could require additional digital activity (such as interactive digital activities or supplementary materials), which could affect the associated electricity consumption. The comparison presented here therefore focuses on a representative scenario in which a student completes the same short course through either delivery mode, allowing for a broadly consistent assessment of the energy required for each approach.

Decarbonisation readiness

As with cloud gaming versus gaming on own device use case, we also assess the DRI for this use case. When students and teachers commute to the educational building by diesel car, the DRI ranges from 28.54 to 29.27 per cent across scenarios. These relatively low figures reflect the high share of energy use from commuting relative to building electricity consumption.

By contrast, the online education platform is entirely electricity-based, resulting in a DRI of 100 per cent.

5.4 Cloud company back-up versus server back-up onto physical media

This use case assesses the electricity consumption associated with backing up company data using cloud-based infrastructure (digital scenario) compared with backing up the same data using an on-premises server (counterfactual). The analysis is grounded in a representative scenario that defines the scale and storage duration of the back-up activity, ensuring that both approaches are assessed on a consistent basis.

The representative scenario assumes a small or medium-sized enterprise (SME) which generates and stores operational data, of which **100 GB is backed up for one year**. In the digital scenario, this back-up is created by copying data from a primary cloud data centre to a secondary cloud data centre located at a different site, where the back-up copy is stored for one year. In the counterfactual scenario, the same volume of data is backed up by copying data from a primary on-premises server to a dedicated back-up storage device located in another building, with the back-up copy likewise stored for one year.

Under both approaches, the back-up process is assumed to be automated and does not require any interaction with end-user devices. Electricity consumption is estimated for the processes required to create and store the back-up copy. In the digital approach, these include downloading data from the primary cloud data centre, internet transmission between data centres, and long-term storage in the secondary data centre. In the on-premises counterfactual, they include downloading data from the primary server, transmission over the internet, and storage on the back-up server in another building.

5.4.1 Digital scenario

Here we set out the estimated electricity use associated with a company cloud back-up. The analysis includes electricity use from two main sources: (i) data centre electricity consumption, including both IT and non-IT processes, and (ii) internet transmission.

Data centre

Cloud back-ups involve several IT processes within data centres: downloading/streaming, uploading to data centre and long-term storage.

The table below presents the activity volumes for each process under low, medium and high scenarios, along with their classification as shared or specific.

Table 5.49: Activity volumes and process classification for IT processes

Data Process	Units	Low	Medium	High	Shared/ Specific	Basis of Estimate
Uploading to data centres	GB Uploaded	100	100	100	Specific	*File size being uploaded to the back-up server*
Video processing	GB processed	0	0	0	-	-
Long-term storage	GB stored	100	100	100	Specific	*File size being stored for a year*
RAM storage	GB-hours	0	0	0	-	-
General processing	CPU-hours	0	0	0	-	-
GPU processing	GPU-hours	0	0	0	-	-
AI training	FLOP	0	0	0	-	-
AI interface	FLOP	0	0	0	-	-
Downloading/Streaming	GB downloaded/ streamed	100	100	100	Specific	*File size being downloaded from cloud server*

Source: Europe Economics analysis.

Using these activity volumes and corresponding electricity intensities, the total electricity consumption associated with each IT process is calculated for our representative scenario. Since all processes are specific, no electricity consumption is shared across a wider user base in this use case.

Table 5.50: Electricity used in each IT process – cloud back-up (kWh)

Data Process	Low	Medium	High
Uploading to data centres	0.0005	0.0050	0.0500
Long-term storage	0.294	0.625	1.067
Downloading	0.0005	0.0050	0.0500
Sub-total	**0.295**	**0.635**	**1.167**

Source: Europe Economics analysis.

Accounting for non-IT processes within the data centre using PUE values, we obtain total estimates for the data centre electricity consumption, presented in the table below.

Table 5.51: Summary of data centre electricity consumption (kWh)

Component	Low	Medium	High
IT electricity	0.295	0.635	1.167
Non-IT electricity	0.030	0.191	0.583
Total data centre electricity	**0.325**	**0.826**	**1.750**

Source: Europe Economics analysis.

Internet transmission

Electricity is consumed in transmitting the cloud back-up content between data centres over a fixed-line network. The back-up process is automated and does not involve any interaction with end-user devices. The table below shows the electricity consumed in internet transmission.

Table 5.52: Transmission electricity (kWh)

Component	Low	Medium	High
Download transmission	0.096	0.120	0.144

Source: Europe Economics analysis.

Summary of electricity used

This section consolidates the estimated electricity consumption associated with **creating and storing a cloud back-up copy** under the representative scenario. Total electricity use accounts for data centre operations (including both IT and non-IT electricity use) and internet transmission between data centres.

As shown in the table below, internet transmission and IT energy usage each contribute a sizeable share of the overall electricity.

Table 5.53: Total electricity consumption for representative scenario (kWh)

Component	Low	Medium	High
IT electricity (data centres)	0.295	0.635	1.167
Non-IT electricity (data centres)	0.030	0.191	0.583
Internet transmission	0.096	0.120	0.144
Total	**0.421**	**0.946**	**1.894**

Source: Europe Economics analysis.

To show the relative contributions of each component, the table below presents the percentage breakdown of total electricity consumption across the three scenarios.

Table 5.54 : Breakdown of electricity consumption (%)

Component	Low	Medium	High
IT electricity (data centres)	70.16	67.16	61.60
Non-IT electricity (data centres)	7.02	20.15	30.80
Internet transmission	22.82	12.69	7.60
Total	**100**	**100**	**100**

Source: Europe Economics analysis.

Electricity consumption in this use case is primarily associated with **centralised infrastructure**, with **around 77-92 per cent** of total electricity use supplied to centralised locations across the low, medium, and high scenarios, reflecting the dominant role of data centre operations in the cloud back-up process.

5.4.2 Counterfactual

This section sets out the counterfactual electricity use associated with backing up data in a local server room.

Electricity consumed in the local server rooms

The storage of back-up data involves several IT processes within local server rooms: downloading from a server room, uploading to the back-up server room and long-term storage. The total electricity consumption associated with each IT process is shown below for our representative scenario.

Table 5.55: Electricity used in each IT process – on-premise back-up (kWh)

Data Process	Low	Medium	High
Uploading from local server	0.001	0.030	0.500
Long-term storage	0.823	2.659	7.147
Download by back-up local server	0.001	0.030	0.500
Total	**0.825**	**2.719**	**8.147**

Source: Europe Economics analysis.

Accounting for non-IT processes within local server rooms (using PUE values specific to local server rooms), we obtain total estimates for local server room electricity consumption, presented in the table below.

Table 5.56: Summary of local server electricity consumption (kWh)

Component	Low	Medium	High
IT electricity	0.825	2.719	8.147
Non-IT electricity	0.330	2.039	8.961
Total local server electricity	**1.155**	**4.758**	**17.108**

Source: Europe Economics analysis.

Internet transmission

Electricity is consumed in transmitting the cloud back-up content over the internet between local servers. The back-up process is automated and does not involve any interaction with end-user devices. The table below shows the electricity consumption in internet transmission.

Table 5.57: Transmission electricity (kWh)

Component	Low	Medium	High
Download transmission	0.096	0.120	0.144

Source: Europe Economics analysis.

Summary of electricity use

This section consolidates the estimated electricity consumption associated with creating and storing a back-up copy in a local server room under the representative scenario. Total electricity use accounts for local server operations (including both IT and non-IT electricity use) and internet transmission between servers.

As shown in the table below, internet transmission and IT electricity usage each contribute a sizeable share of the overall electricity.

Table 5.58 : Total electricity for representative scenario (kWh)

Component	Low	Medium	High
IT electricity (local servers)	0.8248	2.7189	8.1467
Non-IT electricity (local servers)	0.3299	2.0392	8.9613
Internet transmission	0.0960	0.1200	0.1440
Total	**1.2507**	**4.8780**	**17.2520**

Source: Europe Economics analysis.

To illustrate the relative contributions of each component, the table below presents the percentage breakdown of total electricity consumption across the three scenarios.

Table 5.59: Breakdown of electricity consumption (%)

Component	Low	Medium	High
IT electricity (local servers)	65.95	55.74	47.22
Non-IT electricity (local servers)	26.38	41.80	51.94
Internet transmission	7.68	2.46	0.83
Total	**100**	**100**	**100**

Source: Europe Economics analysis.

We classify electricity consumed in the local server as decentralised, given that an office building will typically have a much lower electricity intensity than a data centre. Thus, the share of electricity consumption occurring in decentralised locations is **100 per cent**.

5.4.3 Comparison of electricity consumption

Cloud back-up and local server back-up result in different electricity consumption outcomes. Both scenarios assess the electricity required to back up the same representative data volume of 100 GB, allowing electricity use to be compared on a consistent basis.

Total Electricity Use Across Scenarios

Across all scenarios, cloud back-up is associated with lower total electricity consumption than the local server back-up counterfactual. This difference is driven primarily by inefficiencies in long-term storage in local server rooms, where hardware may be less electricity-efficient and typically operates at lower utilisation levels, as well as by the higher PUE of server rooms compared with large-scale cloud data centres.

Table 5.60: Total electricity consumption – cloud back-up versus server back-up (kWh)

Scenario	Low	Medium	High
Cloud back-up	0.42	0.95	1.89
Local server back-up	1.25	4.88	17.25

Source: Europe Economics analysis.

The figure below illustrates the same comparison graphically.

Figure 5.5: Comparison of electricity consumption – cloud versus local server back-up

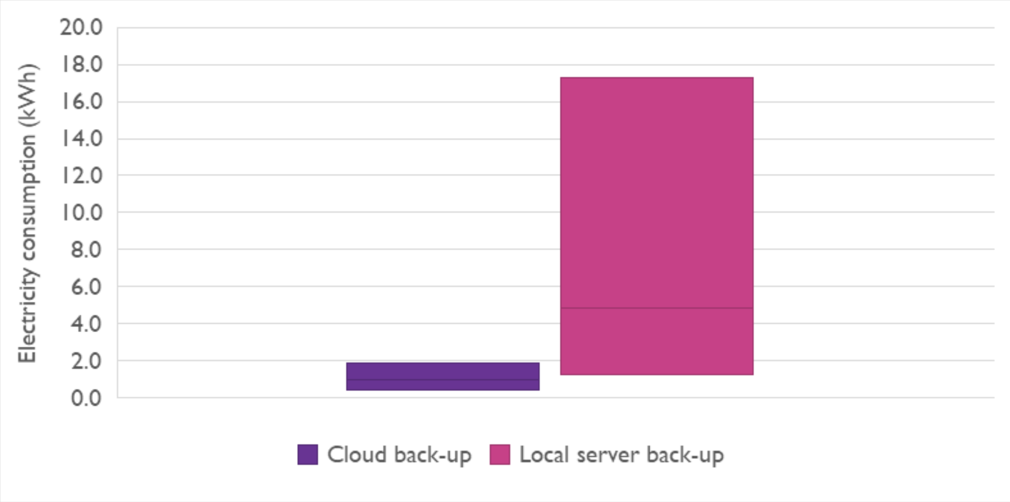

Source: Europe Economics analysis.

A caveat to this comparison is that the representative case reflects the behaviour of a small firm using a relatively small on-premises server set-up. Larger organisations may operate server infrastructure at higher utilisation levels, which could reduce the relative inefficiencies associated with local server rooms.

5.5 AI-assisted email drafting versus manual email drafting

This use case compares the electricity consumption when a worker drafts emails using an AI-based tool (digital scenario) with the electricity consumption when the worker drafts emails without AI assistance (counterfactual).

The representative scenario assumes that a knowledge worker drafts **50 emails per week**, in line with typical estimates of 8–12 emails per working day for office-based roles. Each email is assumed to contain **approximately 200 words** of original drafting text, excluding attachments and quoted material. This results in a total of around **10,000 words drafted per week**. The worker is assumed to be office-based. The large-scale AI model is assumed to have a useful life of six months before retraining is required.

5.5.1 Digital scenario

This section sets out the estimated electricity use associated with drafting emails using a digital AI-based general purpose tool. The analysis includes electricity use from three main sources: (i) data centre electricity consumption, including both IT and non-IT processes, (ii) internet transmission, and (iii) office electricity consumption.

Data centre

The AI-assisted process draws on several IT processes within data centres. These include uploading the prompt and downloading the response, long-term and RAM storage of the AI model, and compute-intensive operations such as AI training and inference. The electricity required for each process is determined by both the scale of activity (e.g. file size, duration, number of floating point operations or FLOP) and whether the process is shared across multiple users or is specific to the representative user.

The table below summarises the IT processes involved, the units used to quantify each activity, the assumed volumes under low, medium, and high scenarios, and whether the electricity for that process is treated as shared or specific to the user. A final column outlines the basis of each estimate.

Table 5.61: IT processes within data centre – AI-assisted email drafting

Data Process	Units	Low	Medium	High	Shared/ Specific	Basis of Estimate
Uploading to data centres	GB uploaded	0.0000625	0.0000625	0.0000625	Specific	*File size of prompts uploaded*
Video processing	GB processed	0	0	0	-	-
Long-term storage	GB stored per year	140	350	3,200	Shared	*Size of AI model*
RAM storage	GB-hours	613,620	1,534,050	4,733,640	Shared	*Storage of AI model for the period of representative scenario*
General processing	CPU-hours	0	0	0	-	-
AI training	FLOP	3.0×10^{24}	3.0×10^{25}	3.0×10^{26}	Shared	*Number of FLOPs for training AI model*
AI inference	FLOP	8.125×10^{14}	1.625×10^{15}	3.25×10^{15}	Specific	*Number of FLOPs per week of email drafted*
Downloading/ streaming	GB downloaded	0.00025	0.00025	0.00025	Specific	*File size of prompt response downloaded*

Source: Europe Economics analysis.

Using these activity volumes and corresponding electricity intensities (Table 3.1), the total electricity consumption associated with each IT process is calculated. These figures represent the total electricity usage, prior to allocating shared electricity across the user base.

Table 5.62: Electricity used in each IT process – AI-assisted email drafting (kWh)

Data Process	Low	Medium	High
Uploading to data centres	6.25×10^{-9}	1.13×10^{-8}	3.13×10^{-8}
Long-term storage	0.21	1.09	5.76
RAM storage	164.36	343.63	536.48
AI training	833,333	45,833,333	833,333,333
AI inference	0.0002	0.0025	0.0090
Downloading/streaming	1.0×10^{-8}	1.50×10^{-8}	3.00×10^{-8}
Sub-total	**833,564**	**45,833,936**	**833,333,270**

Source: Europe Economics analysis.

To estimate the electricity attributable to the representative user, we apply allocation factors for shared processes (e.g. RAM storage, long-term storage, and AI training), based on the assumed user base. For processes such as AI inference and transmission that are specific to each user, the full electricity is allocated to the user. The table below presents the electricity consumption attributed to the representative user after allocating shared electricity across the user base.

Table 5.63: IT electricity attributable to each user (kWh)

IT Process (attributable to user)	Low	Medium	High
Uploading to data centres	6.25×10^{-9}	1.13×10^{-8}	3.13×10^{-8}
Long-term storage	5.0×10^{-13}	1.6×10^{-11}	7.0×10^{-10}
RAM storage	5.6×10^{-10}	8.78×10^{-9}	2.4×10^{-7}
AI training	0.000002	0.0007	0.1014
AI inference	0.0002	0.0025	0.0090
Downloading/streaming	1.0×10^{-8}	1.50×10^{-8}	3.00×10^{-8}
Total IT electricity	**0.0002**	**0.0032**	**0.1104**

Source: Europe Economics analysis.

In addition to the IT electricity, we include an allocation for non-IT electricity used to support data centre operations (e.g. cooling, lighting, and power supply systems). This is calculated by applying a proportional uplift to each IT process (**Table 3.2**) using our assumed PUE values. The table below summarises both IT and non-IT electricity components to give the total data centre electricity attributable to the user.

Table 5.64: Summary of data centre electricity consumption (kWh)

Component	Low	Medium	High
IT electricity	0.0002	0.0032	0.1104
Non-IT electricity	0.0000	0.0009	0.0552
Total data centre electricity	**0.0003**	**0.0041**	**0.1656**

Source: Europe Economics analysis.

Internet transmission

Electricity is consumed in internet transmission when the user uploads the prompt to the AI tool and downloads the initial draft of the email using a fixed line network. As this transmission is user-specific, the entire electricity consumption is attributed directly to the representative user. This electricity is calculated using an assumed electricity intensity per GB of data transferred, applied to the file size of 0.00025 GB (as specified in the representative scenario).

Table 5.65: Internet transmission electricity attributable to each user (kWh)

Component	Low	Medium	High
Upload transmission	0.0000013	0.0000016	0.0000020
Download transmission	0.0000052	0.0000066	0.0000079
Total transmission	**0.0000066**	**0.0000082**	**0.0000098**

Source: Europe Economics analysis.

Office electricity consumption

The worker is assumed to be office based. The total electricity is estimated by multiplying the electricity per office worker per working day by the number of days required for emailing with AI-assisted tools.[89]

Table 5.66: Electricity used in office

Units	Unit	Low	Medium	High
Number of hours required for email per week with AI-assisted tools	Hours	5.00	12.00	28.33
Number of working days required for email with AI-assisted tools	Working days	0.67	1.60	3.78
Office electricity use with AI-assisted tools	kWh	**1.764**	**7.534**	**27.447**

Source: Europe Economics analysis.

Summary of electricity use

The table below combines the three components of electricity use. The office electricity consumption accounts for the largest share of total consumption in all scenarios, with the data centre electricity and internet transmission electricity being modest in comparison.

Table 5.67: Total electricity for representative scenario (kWh)

[89] This is calculated as (total number of words drafted per week/ speed of writing emails)/ Number of hours in a workday

Component	Low	Medium	High
Data centre IT electricity	0.00023	0.00315	0.11043
Data centre non-IT electricity	0.00002	0.00095	0.05521
Internet transmission	0.00001	0.00001	0.00001
Office electricity consumption	1.7637	7.5341	27.4476
Total electricity use	**1.7640**	**7.5382**	**27.6133**

Source: Europe Economics analysis.

Percentage breakdown of electricity use

The table below shows the percentage contribution of each component of total electricity use across the three scenarios. Office electricity consumption accounts for the largest share of total consumption in all scenarios.

Table 5.68: Breakdown of electricity consumption (%)

Component	Low	Medium	High
Data centre IT electricity	0.01	0.04	0.40
Data centre non-IT electricity	0.00	0.01	0.20
Internet transmission	0.00	0.00	0.00
Office electricity consumption	99.99	99.95	99.40
Total	**100**	**100**	**100**

Source: Europe Economics analysis.

For AI-assisted emails, we treat electricity consumed within the data centre as centralised, while electricity used for transmission and office electricity consumption is considered decentralised. Across the low, medium, and high scenarios, the share of electricity consumed in centralised locations is around **0.01-0.60 per cent**.

5.5.2 Counterfactual

This section sets out the estimated electricity use associated with drafting emails without any AI assistance. The analysis focuses on the electricity used in the worker's office. Range estimates are presented to reflect variations in task duration and office electricity intensity.

Task duration and office electricity consumption

The time required to complete email drafting is calculated by dividing the total word count (10,000 words, as defined in the representative scenario) by assumed email drafting speeds. This defines the total number of hours required to complete the task, which is then converted into working days based on a 7.5-hour workday.

Table 5.69: Time required for email drafting

	Units	Low	Medium	High
Number of hours required for email drafting	Hours	12.5	20.0	33.3
Number of working days required	Working days	1.67	2.67	4.44

Source: Europe Economics analysis.

Electricity consumption in the worker's office is calculated by multiplying the assumed EUI (kWh/m²/year) by office space per worker (m²) to derive electricity consumption per worker per year. This annual figure is then scaled to reflect the number of working days spent completing the task, assuming 253.25 working days per year.

Table 5.70: Electricity consumed in worker's office (kWh)

	Low	Medium	High
Electricity for worker for emailing	4.41	12.56	32.29

Source: Europe Economics analysis.

We classify electricity consumed in the office as decentralised, given that an office will typically have a much lower electricity intensity than a data centre. Thus, the share of electricity consumption occurring in decentralised locations is **100 per cent**.

5.5.3 Comparison of electricity consumption

The table below summarises the total electricity consumption for the task in our low, medium, and high scenarios. Our calculations show that manual email drafting requires more electricity than its AI-based counterpart. Under medium assumptions, the manual scenario results in over 12 kWh of electricity consumption - more than 50 per cent than the 7.5 kWh associated with the AI-assisted email drafting. The difference is primarily driven by office electricity consumption. In the manual scenario, more time is required to draft the email, meaning that office lighting, heating, cooling and computing equipment are used for longer for a given task. By contrast, AI-assisted drafting can reduce the time required to produce a draft email, increasing worker productivity and therefore reducing the office electricity consumed per task.

Table 5.71: Total electricity consumption – AI-assisted versus manual email drafting (kWh)

Scenario	Low	Medium	High
AI-assisted email drafting	1.76	7.54	27.61
Manual email drafting	4.41	12.56	32.29

Source: Europe Economics analysis.

The figure below illustrates the same comparison graphically.

Figure 5.6: Comparison of electricity consumption – AI-assisted versus manual email drafting

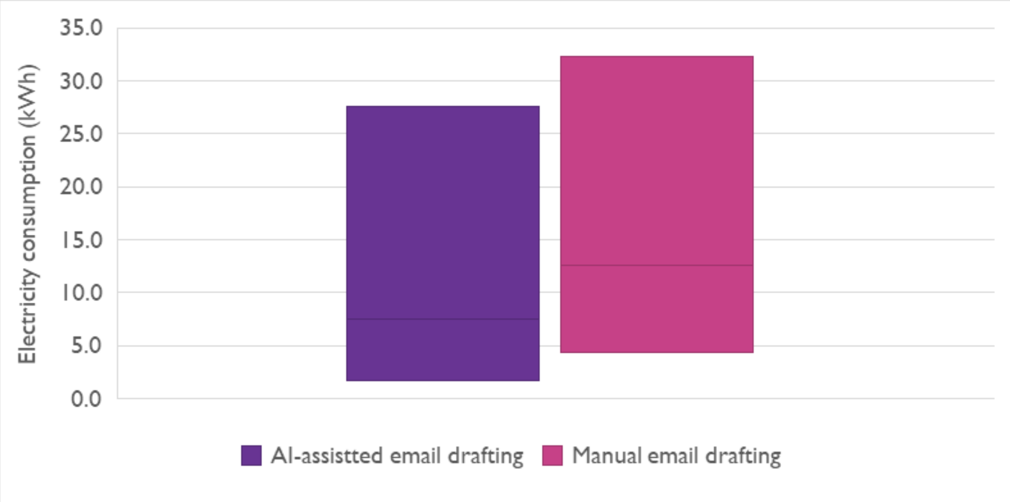

Source: Europe Economics analysis.

5.6 AI-assisted code development versus manual code development

This use case assesses the electricity consumption involved in a worker developing code using an AI-based tool (digital scenario) compared with manual code development without any AI-assistance (counterfactual).

The representative scenario assumes a small software development task comprising approximately 500 lines of code. This is consistent with published benchmarks for small software projects and discrete development tasks. Average line length is assumed to be 40 characters per line. This results in a total of around **20,000 character of code being developed**. The worker is assumed to be office-based. The coding-specialised model is assumed to have a useful life of one year before retraining is required.

5.6.1 Digital scenario

This section presents our estimates of the electricity use associated with code development using a dedicated AI tool. We cover: (i) data centre electricity consumption, including both IT and non-IT processes, (ii) electricity used in internet transmission, and (iii) office electricity consumption.

Data centre

The table below provides an overview of the IT processes that take place within data centres to support the AI-assisted workflow. It identifies the units used to measure each activity, the assumed volume under low, medium and high scenarios, and whether the associated electricity use is treated as shared across users or specific to the representative user. The final column explains the basis for the estimates used.

Table 5.72: IT processes within data centre – AI-assisted code development

Data Process	Units	Low	Medium	High	Shared/ Specific	Basis of Estimate
Uploading to data centres	GB uploaded	0.000002	0.000004	0.000010	Specific	*File size of prompts uploaded*
Video processing	GB processed	0	0	0	-	-
Long-term storage	GB stored per year	70	80	90	Shared	*Size of AI model*
RAM storage	GB-hours	613,620	701,280	788,940	Shared	*Storage of AI model for the period of representative scenario*
General processing	CPU-hours	0	0	0	-	-
AI training	FLOP	1.0×10^{21}	1.0×10^{22}	3.0×10^{23}	Shared	*Number of FLOPs for training AI model*
AI inference	FLOP	5.25×10^{14}	2.1×10^{15}	1.05×10^{16}	Specific	*Number of FLOPs per coding task*
Downloading/ streaming	GB downloaded	0.00004	0.00008	0.00020	Specific	*File size of prompt response downloaded*

Source: Europe Economics analysis.

Combining these activity volumes with our corresponding electricity intensities (**Table 3.1**), we calculate the total electricity consumption associated with each IT process. These figures represent the total electricity usage, prior to the allocation of shared electricity across the user base.

Table 5.73: Electricity used in each IT process – AI-assisted code development (kWh)

Data process	Low	Medium	High
Uploading to data centres	2.0×10^{-10}	7.2×10^{-10}	5.0×10^{-9}
Long-term storage	0.21	0.50	0.96
RAM storage	230.11	274.90	321.89
AI training	277.78	15,277.78	833,333.33
AI inference	0.00015	0.00321	0.02917
Downloading/streaming	1.6×10^{-9}	4.8×10^{-9}	2.4×10^{-8}
Sub-total	**508.09**	**15,553.18**	**833,656.21**

Source: Europe Economics analysis.

The table below shows the electricity attributed to the representative user for each data centre process, after the allocation of shared electricity across the user base.

Table 5.74: IT electricity attributable to each user (kWh)

IT Process (attributable to user)	Low	Medium	High
Uploading to data centres	2.0×10^{-10}	7.2×10^{-10}	5.0×10^{-9}
Long-term storage	5.1×10^{-9}	2.5×10^{-8}	9.6×10^{-8}
RAM storage	5.8×10^{-6}	1.4×10^{-5}	3.2×10^{-5}
AI training	6.9×10^{-6}	7.6×10^{-4}	8.33×10^{-2}
AI inference	0.00015	0.00321	0.02917
Downloading/streaming	1.6×10^{-9}	4.8×10^{-9}	2.4×10^{-8}
Total IT electricity	**0.0002**	**0.0040**	**0.1125**

Source: Europe Economics analysis.

To account for non-IT electricity used in supporting data centre operations, we apply a proportional uplift to the electricity used in each IT process (**Table 3.2**) based on the assumed PUE values. The table below summarises the total data centre electricity attributable to the user, including both IT and non-IT electricity components.

Table 5.75: Summary of data centre electricity consumption (kWh)

Component	Low	Medium	High
IT electricity	0.00016	0.00399	0.11254
Non-IT electricity	0.00002	0.00120	0.05627
Total data centre electricity	**0.00017**	**0.00518**	**0.16880**

Source: Europe Economics analysis.

Internet transmission

Electricity consumed in internet transmission during upload and download is calculated using an assumed electricity intensity per GB of data transferred, applied to the file size of 0.00002 GB (as specified in the representative scenario).

Table 5.5.76: Transmission electricity attributable to each user (kWh)

Component	Low	Medium	High
Upload transmission	0.0000000	0.0000001	0.0000003
Download transmission	0.0000008	0.0000021	0.0000063
Total transmission	**0.0000009**	**0.0000022**	**0.0000066**

Source: Europe Economics analysis.

Office electricity consumption

Office electricity consumption is estimated by multiplying the electricity per office worker per working day by the number of days required for coding with AI-assisted tools.[90]

[90] This is calculated as (total number of words drafted per week/ speed of writing emails)/ Number of hours in a workday

Table 5.77: Electricity used in office electricity consumption

Units	Unit	Low	Medium	High
Number of hours required for coding per project with AI-assisted tools	Hours	14	30	105
Number of working days required for coding with AI-assisted tools	Working days	1.80	4.00	14.00
Office electricity use with AI-assisted tools	**kWh**	**4.762**	**18.835**	**101.718**

Source: Europe Economics analysis.

Summary of electricity use

The table below brings together the three components of electricity use, showing that office electricity consumption represents the largest share across all scenarios, while data centre electricity and internet transmission contribute comparatively small amounts.

Table 5.78: Total electricity for representative scenario (kWh)

Component	Low	Medium	High
Data centre IT electricity	0.0002	0.0040	0.1125
Data centre non-IT electricity	0.0000	0.0012	0.0563
Internet transmission	0.0000	0.0000	0.0000
Office electricity consumption	4.7621	18.8351	101.7177
Total electricity use	**4.7623**	**18.8403**	**101.8865**

Source: Europe Economics analysis.

Percentage breakdown of electricity use

The table below shows the percentage contribution of each component of total electricity use across the three scenarios.

Table 5.79: Breakdown of electricity consumption (%)

Component	Low	Medium	High
Data centre IT electricity	0.003	0.021	0.110
Data centre non-IT electricity	0.000	0.006	0.055
Internet transmission	0.000	0.000	0.000
Office electricity consumption	99.996	99.972	99.834
Total	**100**	**100**	**100**

Source: Europe Economics analysis.

Electricity used within the data centre is treated as centralised, while electricity associated with transmission and office use is classified as decentralised. Across the low, medium and high scenarios, the share of electricity consumed in centralised locations is around **0–0.20 per cent**.

5.6.2 Counterfactual

This section sets out the estimated electricity use associated with manual code development without any AI assistance. Range estimates for the electricity used in the worker's office are presented to reflect variations in task duration and office electricity intensity.

Task duration and office electricity consumption

The time required to complete manual code development is calculated by dividing the total character count (20,000 characters, as defined in the representative scenario) by assumed speed of code development. This defines the total number of hours required to complete the task, which is then converted into working days based on a 7.5-hour workday.

Table 5.80: Time required for coding

	Units	Low	Medium	High
Number of hours required for coding	Hours	30	50	150
Number of working days required for coding	Working days	4.00	6.67	20.00

Source: Europe Economics analysis.

Electricity consumption in the worker's office is estimated by applying the assumed EUI (kWh/m²/year) to the office space per worker (m²) to obtain annual electricity use per worker, and then scaling this figure to the number of working days required to complete the task, assuming 253.25 working days per year.

Table 5.81: Electricity consumed in worker's office (kWh)

	Low	Medium	High
Electricity for worker during coding	10.58	31.39	145.31

Source: Europe Economics analysis.

As we classify electricity consumed in the office as decentralised, the share of electricity consumption occurring in decentralised locations is **100 per cent**.

5.6.3 Comparison of electricity consumption

The table below reports total electricity use for the task under low, medium and high scenarios. The results indicate that manual coding is more electricity-intensive than AI-assisted coding. Under the medium case, the manual approach leads to electricity consumption of just over 31 kWh, which is more than 50 per cent higher than the electricity consumption of 19 kWh when AI tools are used. This gap is primarily explained by differences in the time required to complete the task. Writing code manually takes longer, meaning that office lighting, heating, cooling and computing equipment remain in use for a greater period. By shortening the time needed to produce code, AI assistance increases worker productivity and correspondingly reduces the office electricity required per task.

Table 5.82: Total electricity consumption – AI-assisted versus manual code development (kWh)

Scenario	Low	Medium	High
AI-assisted code development	4.76	18.84	101.89
Manual code development	10.58	31.39	145.31

Source: Europe Economics analysis.

The figure below illustrates the same comparison graphically.

Figure 5.7: Comparison of electricity consumption – AI-assisted versus manual code development

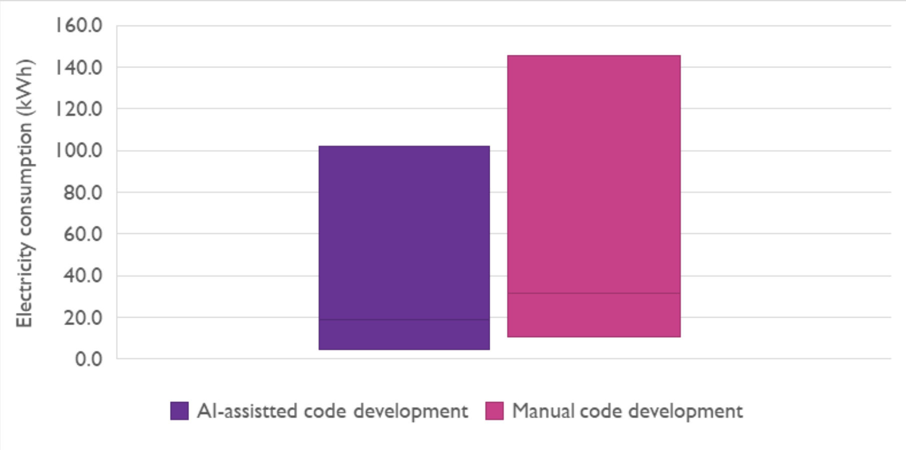

Source: Europe Economics analysis.

5.7 AI-assisted copy-editing versus manual copy-editing

This use case compares the electricity consumption when a copy-editor uses an AI-based tool to assist them in copy-editing professional text (digital scenario) with the electricity consumption of manual copy-editing without AI assistance (counterfactual).

The representative scenario assumes that a knowledge worker copy-edits a single piece of professional text containing **approximately 10,000 words**. The task involves reviewing and refining existing text rather than producing new content, with the AI tool providing suggested edits for the worker to review and apply as appropriate. The worker is assumed to be office-based. As with the other AI use cases, the large-scale language model used for copy-editing is assumed to have a useful life of six months before retraining is required.

5.7.1 Digital scenario

Below, we set out our estimates of the electricity use associated with AI-assisted copy-editing, considering in turn: (i) data centre electricity consumption, including both IT and non-IT processes, (ii) internet transmission electricity consumption, and (iii) office electricity consumption.

Data centre

The AI-assisted process draws on several IT processes within data centres. These include uploading and downloading the prompt and prompt response, long-term and RAM storage of the AI model, and compute-intensive operations such as AI training and inference. The electricity required for each process is determined by both the scale of activity (e.g. file size, duration, number of floating point operations or FLOP) and whether the process is shared across multiple users or is specific to the representative user.

The table below summarises the IT processes involved in AI-assisted copy-editing. For each process, it sets out the units used to quantify the process, assumed volumes under low, medium, and high scenarios, whether the electricity for that process is treated as shared or specific to the user, and the basis for our estimate.

Table 5.83: IT processes within data centre – AI-assisted copy-editing

Data Process	Units	Low	Medium	High	Shared/ Specific	Basis of Estimate
Uploading to data centres	GB uploaded	0.0003	0.0003	0.0003	Specific	*File size of prompts uploaded*
Video processing	GB processed	0	0	0	-	-
Long-term storage	GB stored per year	140	350	3,200	Shared	*Size of AI model*
RAM storage	GB-hours	613,620	1,534,050	4,733,640	Shared	*Storage of AI model for the period of representative scenario*
General processing	CPU-hours	0	0	0	-	-
AI training	FLOP	3.0×10^{24}	3.0×10^{25}	3.0×10^{26}	Shared	*Number of FLOPs for training AI model*
AI inference	FLOP	1.3×10^{15}	2.6×10^{15}	5.2×10^{15}	Specific	*Number of FLOPs for text copy-edited*
Downloading/ streaming	GB downloaded	0.0003	0.0003	0.0003	Specific	*File size of prompt response downloaded*

Source: Europe Economics analysis.

The table below presents figures for total electricity usage, prior to allocating the electricity for share processes across the user base. These figures have been calculated on the basis of the volumes in the previous table and our assumed electricity intensities (**Table 3.1**).

Table 5.84: Electricity used in each IT process – AI-assisted copy-editing (kWh)

Data Process	Low	Medium	High
Uploading to data centres	3.0×10^{-8}	5.4×10^{-8}	1.5×10^{-7}
Long-term storage	0.21	1.09	5.76
RAM storage	230.11	601.35	1,93136.33
AI training	833,333	45,833,333	833,333,333
AI inference	0.00037	0.00407	0.01481
Downloading/streaming	1.20×10^{-8}	1.80×10^{-6}	3.6×10^{-8}
Sub-total	**833,564**	**45,833,936**	**833,335,270**

Source: Europe Economics analysis.

The table below presents the electricity attributed to the user in our representative scenario, after the electricity involved in shared processes has been allocated across the user base.

Table 5.85: IT electricity attributable to each user (kWh)

IT Process (attributable to user)	Low	Medium	High
Uploading to data centres	3.0×10^{-8}	5.4×10^{-8}	7.5×10^{-8}
Long-term storage	4.0×10^{-13}	1.3×10^{-11}	5.6×10^{-10}
RAM storage	4.5×10^{-10}	7.0×10^{-9}	1.9×10^{-7}
AI training	1.6×10^{-6}	5.4×10^{-4}	0.08112
AI inference	0.00037	0.00407	0.01481
Downloading/streaming	1.2×10^{-8}	1.8×10^{-8}	3.6×10^{-8}
Total IT electricity	**0.0004**	**0.0046**	**0.0959**

Source: Europe Economics analysis.

An uplift for non-IT electricity used to support data centre operations is calculated for each IT process (Table 3.2) using our assumed PUE values. The table below summarises both IT and non-IT electricity consumption to give total data centre electricity consumption.

Table 5.86: Summary of data centre electricity consumption (kWh)

Component	Low	Medium	High
IT electricity	0.0004	0.0046	0.0959
Non-IT electricity	0.0000	0.0014	0.0480
Total data centre electricity	**0.0004**	**0.0060**	**0.1439**

Source: Europe Economics analysis.

Internet transmission

The table below sets out our estimates of the electricity consumed in internet transmission when the user uploads the raw text and downloads suggestions for edits to the text using a fixed line network.

Table 5.87: Transmission electricity attributable to each user (kWh)

Component	Low	Medium	High
Upload transmission	0.000006	0.000008	0.000009
Download transmission	0.000006	0.000008	0.000009
Total transmission	**0.000013**	**0.000016**	**0.000019**

Source: Europe Economics analysis.

Office electricity consumption

Our estimates of office electricity consumption are shown below. They were calculated by multiplying the electricity per office worker per working day by the number of days required for AI-assisted copy-editing .[91]

Table 5.88: Electricity used in office

Units	Unit	Low	Medium	High
Number of hours required for editing text	Hours	2.77	4.90	6.23
Number of working days required for editing text	Working days	0.37	0.65	0.83
Office electricity use with AI-assisted tools	kWh	**0.977**	**3.079**	**6.037**

Source: Europe Economics analysis.

Summary of Electricity Use

As shown in the table below, office electricity consumption accounts for the largest share of total consumption in all scenarios.

Table 5.89: Total electricity for representative scenario (kWh)

Component	Low	Medium	High
Data centre IT electricity	0.0004	0.0046	0.0959
Data centre non-IT electricity	0.0000	0.0014	0.0480
Internet transmission	0.0000	0.0000	0.0000
Office electricity consumption	0.9770	3.0794	6.0372
Total electricity use	**0.9775**	**3.0854**	**6.1811**

Source: Europe Economics analysis.

Percentage Breakdown of Electricity Use

The table below shows the percentage contribution of each component of electricity consumption in our low, medium and high scenarios.

Table 5.90: Breakdown of electricity consumption (%)

[91] This is calculated as (total number of words copy-edited/ speed of copy-editing)/ Number of hours in a workday

Component	Low	Medium	High
Data centre IT electricity	0.04	0.15	1.55
Data centre non-IT electricity	0.00	0.04	0.77
Internet transmission	0.00	0.00	0.00
Office electricity consumption	99.96	99.81	97.67
Total	**100**	**100**	**100**

Source: Europe Economics analysis.

The only element of this electricity consumption which we regard as centralised is data centre electricity consumption. Across the low, medium, and high scenarios, the share of electricity consumed in centralised locations ranges from **0 to 2.5 per cent**.

5.7.2 Counterfactual

This section sets out estimated electricity use in the worker's office associated with copy-editing without any AI assistance. We calculate a range estimate to reflect variations in task duration and office electricity intensity.

Task duration and office electricity consumption

We calculate the time required for copy editing by dividing the total word count (10,000 words, as defined in the representative scenario) by assumed copy-editing speeds. We convert the total number of hours required to complete the task into working days using an assumption that a workday comprises 7.5 hours.

Table 5.91: Time required for copy-editing

	Units	Low	Medium	High
Number of hours required for editing	Hours	4.62	5.77	6.92
Number of working days required for editing	Working days	0.62	0.77	0.92

Source: Europe Economics analysis.

Below, we present our estimates of electricity consumption in the worker's office, based on the number of working days required for copy-editing.

Table 5.92: Electricity consumed in worker's office (kWh)

	Low	Medium	High
Electricity for worker for editing	1.628	3.623	6.708

Source: Europe Economics analysis.

The share of electricity consumption occurring in decentralised locations is **100 per cent**, given that we classify electricity consumed in offices as decentralised.

5.7.3 Comparison of electricity consumption

As shown in the table below, our calculations suggest that manual copy-editing requires more electricity to complete the same task than AI-assisted copy-editing. Under medium assumptions, the manual scenario results in over 3.6 kWh of electricity consumption – 15 per cent more than the 3 kWh associated with the AI-assisted copy-editing. The magnitude of the difference is largely driven by time savings from using AI tools, which as well as increasing worker productivity also reduces the office electricity consumed per task.

Table 5.93: Total electricity consumption – AI-assisted versus manual copy-editing (kWh)

Scenario	Low	Medium	High
AI-assisted copy-editing	0.9775	3.0853	6.1806
Manual copy-editing	1.628	3.623	6.708

Source: Europe Economics analysis.

The figure below illustrates the same comparison graphically.

Figure 5.8: Comparison of electricity consumption – AI-assisted versus manual copy-editing

Source: Europe Economics analysis.

5.8 AI-assisted data extraction versus manual data extraction

This use case compares the electricity consumption when a worker extracts structured data from unstructured documents using an AI-based tool (digital scenario) with the electricity consumption associated with manual data extraction without AI assistance (counterfactual).

The representative scenario assumes that a knowledge worker extracts predefined data elements from a single unstructured document, assumed to be **a dense, full-text clinical study** (approximately 15,000 words). The task involves identifying and recording structured information from the document rather than generating new text. The worker is assumed to be office-based.

As with the other AI use cases, the large-scale language model used for data extraction is assumed to have a useful life of six months before retraining is required.

5.8.1 Digital scenario

This section presents our estimates of electricity use associated with data extraction using a digital AI-based general purpose tool. It covers three main sources: (i) data centre electricity consumption, including both IT and non-IT processes, (ii) electricity used in internet transmission, and (iii) office electricity consumption.

Data centre

The table below shows the IT processes involved in AI-assisted data extraction, along with the units used to measure each process, assumed volumes under low, medium, and high scenarios, whether the electricity for that process is shared or specific to the user, and the basis of each estimate.

Table 5.94: IT processes within data centre – AI-assisted data extraction

Data Process	Units	Low	Medium	High	Shared/ Specific	Basis of Estimate
Uploading to data centres	GB uploaded	0.00045	0.00045	0.00045	Specific	*File size of prompts uploaded*
Video processing	GB processed	0	0	0	-	-
Long-term storage	GB stored per year	140	350	3,200	Shared	*Size of AI model*
RAM storage	GB-hours	613,620	1,534,050	4,733,640	Shared	*Storage of AI model for the period of representative scenario*
General processing	CPU-hours	0	0	0	-	-
AI training	FLOP	3.0×10^{24}	3.0×10^{25}	3.0×10^{26}	Shared	*Number of FLOPs for training AI model*
AI inference	FLOP	9.96×10^{14}	2.1×10^{15}	4.27×10^{15}	Specific	*Number of FLOPs for data extracted*
Downloading/ streaming	GB downloaded	$1{,}44 \times 10^{-6}$	3.6×10^{-6}	7.2×10^{-6}	Specific	*File size of prompt response downloaded*

Source: Europe Economics analysis.

Below, we present our figures for total electricity used in each process, prior to allocating shared electricity across the user base. These figures were calculated using the above activity volumes and corresponding electricity intensities (**Table 3.1**).

Table 5.95: Electricity used in each IT process – AI-assisted data extraction (kWh)

Data Process	Low	Medium	High
Uploading to data centres	4.5×10^{-8}	8.1×10^{-8}	2.3×10^{-7}
Long-term storage	0.21	1.09	5.76
RAM storage	230.11	601.35	1,931.33
AI training	833,333	45,833,333	833,333,333
AI inference	0.00028	0.00312	0.01186
Downloading/streaming	5.8×10^{-11}	2.2×10^{-10}	8.6×10^{-10}
Sub-total	**833,564**	**45,833,936**	**833,335,270**

Source: Europe Economics analysis.

The table below presents the electricity attributable to the representative user after allocating shared electricity across the user base.

Table 5.96: IT electricity attributable to each user (kWh)

IT Process (attributable to user)	Low	Medium	High
Uploading to data centres	4.5×10^{-8}	8.1×10^{-8}	2.3×10^{-7}
Long-term storage	0.21	1.09	5.76
RAM storage	6.0×10^{-13}	2.0×10^{-11}	9.0×10^{-10}
AI training	0.000002	0.000821	0.129896
AI inference	0.0003	0.0031	0.0119
Downloading/streaming	5.8×10^{-11}	2.2×10^{-10}	8.6×10^{-10}
Total IT electricity	**0.0003**	**0.0039**	**0.1418**

Source: Europe Economics analysis.

Using our assumed PUE values, we also calculate an uplift to account for non-IT electricity used in data centre operations (e.g. for cooling, lighting, and power supply systems). The table below shows total data centre electricity attributable to the user, calculated by summing IT and non-IT electricity.

Table 5.97: Summary of data centre electricity consumption (kWh)

Component	Low	Medium	High
IT electricity	0.00029	0.00395	0.14176
Non-IT electricity	0.00003	0.00118	0.07088
Total data centre electricity	**0.00031**	**0.00513**	**0.21264**

Source: Europe Economics analysis.

Internet Transmission

Below, we present our estimates of the electricity consumed in internet transmission when the user uploads the prompt and the document to the AI tool and downloads the data using a fixed line network.

Table 5.98: Transmission electricity attributable to each user (kWh)

Component	Low	Medium	High
Upload transmission	0.000009	0.000012	0.000014
Download transmission	0.000000	0.000000	0.000000
Total transmission	**0.000009**	**0.000012**	**0.000014**

Source: Europe Economics analysis.

Office electricity consumption

We have estimated office electricity consumption by multiplying the electricity per office worker per working day by the number of days required for data extraction with AI-assisted tools.[92]

Table 5.99: Electricity used in office

Units	Unit	Low	Medium	High
Number of hours required for data extraction	Hours	0.80	1.40	2.48
Number of working days required for data extraction	Working days	0.11	0.19	0.33
Office electricity use with AI-assisted tools	kWh	**0.28**	**0.88**	**2.40**

Source: Europe Economics analysis.

Summary of electricity use

The table below presents the three components of electricity use. Office electricity consumption accounts for most of the total electricity consumption in all scenarios, with electricity used in data centres and in internet transmission being small in comparison.

Table 5.100: Total electricity for representative scenario (kWh)

Component	Low	Medium	High
Data centre IT electricity	0.0003	0.0040	0.1418
Data centre non-IT electricity	0.0000	0.0012	0.0709
Internet transmission	0.0000	0.0000	0.0000
Office electricity consumption	0.2811	0.8790	2.4004
Total electricity use	**0.2814**	**0.8841**	**2.6131**

Source: Europe Economics analysis.

Percentage breakdown of electricity use

The table below shows the percentage breakdown of electricity use across the three scenarios.

Table 5.101: Breakdown of electricity consumption (%)

[92] This is calculated as (total number of documents reviewed/ speed of data extraction)/ Number of hours in a workday

Component	Low	Medium	High
Data centre IT electricity	0.10	0.45	5.43
Data centre non-IT electricity	0.01	0.13	2.71
Internet transmission	0.00	0.00	0.00
Office electricity consumption	99.88	99.42	91.86
Total	**100**	**100**	**100**

Source: Europe Economics analysis.

We treat electricity consumed within the data centre as centralised, while electricity used for transmission and office electricity consumption is considered decentralised. The share of electricity consumed in centralised locations varies across the low, medium and high scenarios in the range **0.1 to 8.14 per cent**.

5.8.2 Counterfactual

This section sets out the estimated office electricity consumption associated with data extraction without any AI assistance. We derive a range estimate given potential variations in task duration and office electricity intensity.

Task duration and office electricity consumption

The time required to extract data from a full-text clinical study (as defined in the representative scenario) is estimated using assumed data-extraction speeds. The resulting number of hours required to complete the task is then converted into working days based on an assumed 7.5-hour workday.

Table 5.102: Time required for data extraction

	Units	Low	Medium	High
Number of hours required for data extraction	Hours	1.38	2.08	3.07
Number of working days required for data extraction	Working days	0.18	0.28	0.41

Source: Europe Economics analysis.

Our estimate of office electricity consumption is calculated by multiplying electricity consumption per worker per working day by the number of working days spent completing the task.

Table 5.103: Electricity consumed in worker's office (kWh)

	Low	Medium	High
Electricity for worker for data extraction	0.488	1.308	2.971

Source: Europe Economics analysis.

We treat electricity consumed in the office as decentralised, and hence the share of electricity consumption occurring in decentralised locations is **100 per cent**.

5.8.3 Comparison of electricity consumption

Our estimates, show in the table below, suggest that manual data extraction requires more electricity than its AI-based counterpart. In the medium scenario, the manual approach results in over 1.3 kWh of electricity consumption – approximately 50 per cent more than the 0.88 kWh associated with the AI-assisted data extraction. The difference is largely driven by the time savings from using AI to assist the worker, which as well as raising worker productivity also reduces office electricity consumption attributable to the task.

Table 5.104: Total electricity consumption – AI-assisted versus manual data extraction (kWh)

Scenario	Low	Medium	High
AI-assisted data extraction	0.28	0.88	2.61
Manual data extraction	0.49	1.31	2.97

Source: Europe Economics analysis.

The figure below illustrates the same comparison graphically.

Figure 5.9: Comparison of electricity consumption – AI-assisted versus manual data extraction

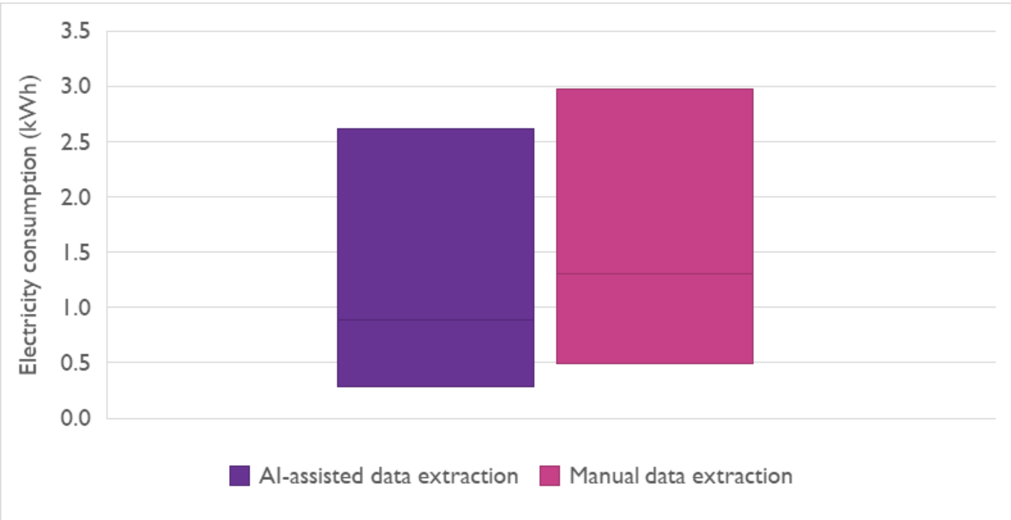

Source: Europe Economics analysis.

5.9 AI-assisted transcription versus manual transcription

This use case assesses the electricity consumption involved in a worker transcribing recorded audio into text using an AI-based transcription tool and then checking the output (digital scenario). It compares this with the electricity consumption associated with manual transcription carried out entirely by a human (counterfactual).

The representative scenario assumes that a knowledge worker transcribes a single audio recording of **one hour in duration**. The task involves converting spoken audio into written text rather than generating new content. The worker is assumed to be based in an office. The speech-to-text model used for transcription is assumed to have a useful life of one year before retraining or replacement is required.

5.9.1 Digital scenario

This section sets out the estimated electricity use associated with data extraction using a dedicated AI transcription tool following by human checking. It covers: (i) data centre electricity consumption, including both IT and non-IT processes, (ii) internet transmission electricity consumption, and (iii) office electricity consumption.

Data centre

The table below summarises the IT processes involved in using an AI transcription tool. It sets out the units used for each process, the assumed volumes under low, medium, and high scenarios, whether the electricity for that process is shared or specific to the user, and the basis of each estimate.

Table 5.105: IT processes within data centre – AI-assisted transcription

Data Process	Units	Low	Medium	High	Shared/ Specific	Basis of Estimate
Uploading to data centres	GB uploaded	0.10	0.10	0.10	Specific	*File size of prompts uploaded*
Video processing	GB processed	0	0	0	-	-
Long-term storage	GB stored per year	1.5	3	6	Shared	*Size of AI model*
RAM storage	GB-hours	13,149	26,298	52,596	Shared	*Storage of AI model for the period of representative scenario*
General processing	CPU-hours	0	0	0	-	-
AI training	FLOP	3.11×10^{17}	3.80×10^{21}	$7.6 \times 10^{2`}$	Shared	*Number of FLOPs for training AI model*
AI inference	FLOP	1.93×10^{15}	1.12×10^{16}	5.62×10^{16}	Specific	*Number of FLOPs for transcription*
Downloading/ streaming	GB downloaded	0.00002	0.00002	0.00002	Specific	*File size of prompt response downloaded*

Source: Europe Economics analysis.

Using these activity volumes and corresponding electricity intensities (**Table 3.1**), we calculate the total electricity consumption associated with each IT process prior to allocating shared electricity across the user base.

Table 5.106: Electricity used in each IT process – AI-assisted transcription (kWh)

Data Process	Low	Medium	High
Uploading to data centres	0.0015	0.002	0.003
Long-term storage	0.004	0.019	0.064
RAM storage	4.931	10.309	21.459
AI training	0.09	5,805.79	21,111.11
AI inference	0.0005	0.0686	0.4680
Downloading/streaming	1.1×10^{-10}	1.05×10^{-9}	1.1×10^{-8}
Sub-total	**5.02**	**5,816.19**	**21,133.10**

Source: Europe Economics analysis.

To estimate the electricity attributable to the representative user, we then allocate shared electricity across the user base. The table below presents the electricity attributable to the user in the representative scenario.

Table 5.107: IT electricity attributable to each user (kWh)

IT Process (attributable to user)	Low	Medium	High
Uploading to data centres	0.0015	0.002	0.003
Long-term storage	3.2×10^{-12}	2.7×10^{-11}	4.6×10^{-10}
RAM storage	3.6×10^{-9}	1.5×10^{-8}	1.5×10^{-7}
AI training	6.2×10^{-11}	8.4×10^{-6}	1.52×10^{-4}
AI inference	0.0005	0.0686	0.468
Downloading/streaming	1.1×10^{-10}	1.05×10^{-9}	1.1×10^{-8}
Total IT electricity	**0.0005**	**0.0686**	**0.4682**

Source: Europe Economics analysis.

An uplift to account for non-IT electricity within data centres is calculated using our assumed PUE values. The table below sums IT and non-IT electricity to give the total data centre electricity attributable to the user.

Table 5.108: Summary of data centre electricity consumption (kWh)

Component	Low	Medium	High
IT electricity	0.0005	0.0687	0.4682
Non-IT electricity	0.0001	0.0206	0.2341
Total data centre electricity	**0.0006**	**0.0892**	**0.7023**

Source: Europe Economics analysis.

Internet transmission

Using an assumed electricity intensity per GB of data transferred, we calculate the electricity consumed in internet transmission when the user uploads the audio recording and downloads the AI-generated transcript using a fixed line network.

Table 5.109: Transmission electricity attributable to each user (kWh)

Component	Low	Medium	High
Upload transmission	0.002096	0.002620	0.003144
Download transmission	0.000000	0.000001	0.000001
Total transmission	**0.002096**	**0.002621**	**0.003145**

Source: Europe Economics analysis.

Office electricity consumption

We estimate office electricity consumption by multiplying the electricity per office worker per working day by the number of days required for a human to carry out the transcription tasks with the assistance of AI tools.[93]

Table 5.110: Electricity used in office

Units	Unit	Low	Medium	High
Number of hours required for human transcription	Hours	0.75	1.50	2.50
Number of working days required for transcription	Working days	0.10	0.20	0.33
Office electricity use with AI-assisted tools	kWh	**0.26**	**0.94**	**2.42**

Source: Europe Economics analysis.

Summary of electricity use

The table below shows that office electricity consumption represents the majority of electricity consumption in all scenarios. Data centre electricity and internet transmission electricity are modest by comparison.

Table 5.111: Total electricity for representative scenario (kWh)

Component	Low	Medium	High
Data centre IT electricity	0.0005	0.0687	0.4682
Data centre non-IT electricity	0.0001	0.0206	0.2341
Internet transmission	0.0021	0.0026	0.0031
Office electricity consumption	0.2646	0.9418	2.4218
Total electricity use	**0.2672**	**1.0336**	**3.1273**

Source: Europe Economics analysis.

Percentage breakdown of electricity use

The table below shows the percentage breakdown total electricity use across the three scenarios.

Table 5.112: Breakdown of electricity consumption (%)

[93] This is calculated as (Number of audio recordings* average length of audio recordings * time taken to transcribe 1 hour of audio)/ Number of hours in a workday

Component	Low	Medium	High
Data centre IT electricity	0.20	6.64	14.97
Data centre non-IT electricity	0.02	1.99	7.49
Internet transmission	0.78	0.25	0.10
Office electricity consumption	99.00	91.11	77.44
Total	**100**	**100**	**100**

Source: Europe Economics analysis.

We treat electricity consumed within the data centre as centralised, while electricity used for transmission and office electricity consumption is considered decentralised. Across the low, medium, and high scenarios, the share of electricity consumed in centralised locations ranges from **0.8 to 9 per cent**.

5.9.2 Counterfactual

This section sets out the estimated office electricity consumption associated with transcription without any AI assistance, with range estimates used to reflect uncertainty over task duration and office electricity intensity.

Task duration and office electricity consumption

We calculate the time required to transcribe a one-hour audio recording (as defined in the representative scenario) using the assumed data extraction speeds. The hours required to complete the task are converted into working days assuming a 7.5-hour workday.

Table 5.5.113: Time required for transcription

	Units	Low	Medium	High
Number of hours required for transcription	Hours	3.00	4.00	5.00
Number of working days required for transcription	Working days	0.40	0.53	0.67

Source: Europe Economics analysis.

The table below shows estimated office electricity consumption, calculated by multiply electricity consumption per worker per working day by the number of days required for the transcription task.

Table 5.114: Electricity consumed in worker’s office (kWh)

	Low	Medium	High
Electricity for worker for transcription	1.058	2.511	4.844

Source: Europe Economics analysis.

The share of electricity consumption occurring in decentralised locations is **100 per cent**, given that we classify offices as a decentralised.

5.9.3 Comparison of electricity consumption

The table below shows that manual transcription requires more electricity than transcription carried out with AI assistance. The manual approach results in over 1.3 kWh of electricity consumption in the medium scenario – approximately 50 per cent more than the 0.88 kWh associated with the AI-assisted transcription. Our findings are largely driven by time savings from using AI, which increase worker productivity and reduce the office electricity consumed per task.

Table 5.115: Total electricity consumption – AI-assisted versus manual transcription (kWh)

Scenario	Low	Medium	High
AI-assisted transcription	0.267	1.034	3.127
Manual transcription	1.058	2.511	4.844

Source: Europe Economics analysis.

The figure below illustrates the same comparison graphically.

Figure 5.10: Comparison of electricity consumption – AI-assisted versus manual transcription

Source: Europe Economics analysis.

5.10 AI-assisted demand forecasting model

This use case examines the development of a short-term electricity demand forecasting model using an automated machine learning (AutoML) system (digital scenario), compared with a conventional manual modelling approach carried out by analysts (counterfactual).

The representative scenario assumes that in the counterfactual two analysts develop the model over 10 working days. The task involves developing and testing multiple candidate model specifications, using a common underlying dataset and standard statistical tools, to produce a short-term electricity demand forecasting model. The analysts are assumed to be based in an office environment.

In the digital approach, the AutoML system is assumed to have a useful life of two years, consistent with its development period, after which retraining or replacement is required.

5.10.1 Digital scenario

This section sets out the estimated electricity use associated with the development of a demand forecasting model using an AutoML system guided by human analysts. It covers: (i) data centre electricity consumption, including both IT and non-IT processes, (ii) internet transmission electricity consumption, and (iii) office electricity consumption.

Data centre

The table below summarises the IT processes involved in using an AutoML system to build a statistical model. Unlike other AI-assisted use cases, where training electricity is estimated using FLOPs, this use case employs a direct estimate of the electricity consumed in developing the AutoML system (in kWh). This is because an AutoML platform is not trained in the same way as an AI model as it is a system for developing models.

Table 5.116: IT processes within data centre – AI-assisted model development

Data Process	Units	Low	Medium	High	Shared/ Specific	Basis of Estimate
Uploading to data centres	GB uploaded	2	2	2	Specific	*File size of data uploaded*
Video processing	GB processed	0	0	0	-	-
Long-term storage	GB stored per year	0.5	1.55	2.6	Shared	*Size of AutoML system*
RAM storage	GB-hours	8,766	27,175	45,583	Shared	*Storage of AutoML system for the period of representative scenario*
General processing	CPU-hours	0	0	0	-	-
AI training	FLOP	0	0	0	-	*Replaced with figure for electricity used in development of AutoML system*
AI inference	FLOP	1.68×10^{17}	8.99×10^{17}	3.37×10^{18}	Specific	*Number of FLOPs for model development*
Downloading/ streaming	GB downloaded	0.02	0.02	0.02	Specific	*File size of data downloaded*

Source: Europe Economics analysis.

Using these activity volumes and corresponding electricity intensities (Table 3.1), we calculate the total electricity consumption associated with each IT process prior to allocating shared electricity across the user base.

Table 5.117: Electricity used in each IT process – AI-assisted model development (kWh)

Data Process	Low	Medium	High
Uploading to data centres	1.0×10^{-5}	1.0×10^{-4}	1.0×10^{-3}
Long-term storage	0.0029	0.0194	0.0555
RAM storage	3.287	10.652	18.598
Development of AutoML system	194,900	213,400	222,700
AI inference	0.004	0.749	9.360
Downloading/streaming	1.0×10^{-7}	1.0×10^{-6}	1.0×10^{-5}
Sub-total	**194,903**	**213,411**	**222,728**

Source: Europe Economics analysis.

To estimate the electricity attributable to the representative user, we then allocate shared electricity across the user base. The table below presents the electricity attributable to the user in the representative scenario.

Table 5.118: IT electricity attributable to each user (kWh)

IT Process (attributable to user)	Low	Medium	High
Uploading to data centres	1.0×10^{-5}	1.0×10^{-4}	1.0×10^{-3}
Long-term storage	1.8×10^{-10}	6.0×10^{-8}	1.3×10^{-5}
RAM storage	2.0×10^{-7}	3.3×10^{-5}	4.3×10^{-3}
Development of AutoML system	0.004	0.165	10.310
AI inference	0.004	0.749	9.360
Downloading/streaming	1.0×10^{-7}	1.0×10^{-6}	1.0×10^{-5}
Total IT electricity	**0.008**	**0.914**	**19.672**

Source: Europe Economics analysis.

An uplift to account for non-IT electricity within data centres is calculated using our assumed PUE values. The table below sums IT and non-IT electricity to give the total data centre electricity attributable to the user.

Table 5.119: Summary of data centre electricity consumption (kWh)

Component	Low	Medium	High
IT electricity	0.008	0.914	19.672
Non-IT electricity	0.001	0.274	9.836
Total data centre electricity	**0.009**	**1.188**	**29.508**

Source: Europe Economics analysis.

Internet transmission

Here we calculate the electricity consumed in internet transmission when data is exchanged between the user and the AutoML platform during model development, including data upload and retrieval of outputs, over a fixed line network.

Table 5.120: Transmission electricity attributable to each user (kWh)

Component	Low	Medium	High
Upload transmission	0.04192	0.05240	0.06288
Download transmission	0.00042	0.00052	0.00063
Total transmission	**0.04234**	**0.05292**	**0.06351**

Source: Europe Economics analysis.

Office electricity consumption

The office electricity consumption reflects the electricity use associated with two analysts working to develop, review and validate the statistical model using the AutoML workflow.

Table 5.121: Electricity used in office

Units	Unit	Low	Medium	High
Number of person hours required to build model	Hours	34	60	93.8
Number of person working days required to build model	Working days	4.5	8.0	12.5
Office electricity use with AI-assisted tools	kWh	**11.91**	**37.67**	**90.82**

Source: Europe Economics analysis.

Summary of electricity use

The table below shows that office electricity consumption represents the majority of electricity consumption in all scenarios. Data centre electricity and internet transmission electricity are modest by comparison.

Table 5.122: Total electricity for representative scenario (kWh)

Component	Low	Medium	High
Data centre IT electricity	0.008	0.914	19.672
Data centre non-IT electricity	0.001	0.274	9.836
Internet transmission	0.042	0.053	0.064
Office electricity consumption	11.905	37.670	90.819
Total electricity use	**11.956**	**38.912**	**120.391**

Source: Europe Economics analysis.

Percentage breakdown of electricity use

The table below shows the percentage breakdown total electricity use across the three scenarios.

Table 5.123: Breakdown of electricity consumption (%)

Component	Low	Medium	High
Data centre IT electricity	0.07	2.35	16.34
Data centre non-IT electricity	0.01	0.70	8.17
Internet transmission	0.35	0.14	0.05
Office electricity consumption	99.57	96.81	75.44
Total	**100**	**100**	**100**

Source: Europe Economics analysis.

Electricity consumed within data centres is treated as centralised, while electricity used for internet transmission and office-based activity is treated as decentralised. The share of electricity supplied to centralised locations ranges from **0.07 per cent** in the low case to **24.51 per cent** in the high case, with a medium estimate of **3.05 per cent**. This reflects the fact that office electricity dominates in the low and medium scenarios, while data centre electricity becomes more significant in the high scenario due to higher electricity per user for the development of the AutoML system.

5.10.2 Counterfactual

This section sets out the estimated office electricity consumption associated with model development without any AI assistance, with range estimates used to reflect uncertainty over task duration and office electricity intensity.

Task duration and office electricity consumption

The time required to develop the statistical model is estimated based on a manual workflow in which two analysts build and test multiple model specifications using standard statistical tools.

Table 5.5.124: Time required for model development

	Units	Low	Medium	High
Number of hours required for developing model	Hours	113	150	188
Number of working days required for model development	Working days	15	20	25

Source: Europe Economics analysis.

The table below presents the associated office electricity consumption for model development in the absence of AutoML assistance.

Table 5.125: Electricity consumed in worker's office (kWh)

	Low	Medium	High
Electricity for worker for model development	39.68	94.18	181.64

Source: Europe Economics analysis.

The share of electricity consumption occurring in decentralised locations is **100 per cent**, given that we classify offices as a decentralised.

5.10.3 Comparison of electricity consumption

The table below compares total electricity consumption for AutoML-assisted and manual model development. AutoML-assisted development results in lower electricity consumption across the low, medium and high scenarios, primarily due to reduced analyst time and therefore lower office electricity use. In the medium scenario, total electricity consumption is **38.9 kWh** for the AutoML-assisted approach, compared with **94.2 kWh** for manual development.

The results should be interpreted with a degree of caution, as they are sensitive to assumptions on both the allocation of electricity consumed developing the AutoML system and the time required to develop the statistical model.

Table 5.126: Total electricity consumption – AutoML-assisted versus manual model development (kWh)

Scenario	Low	Medium	High
AutoML-assisted model development	11.96	38.91	120.39
Manual model development	39.68	94.18	181.64

Source: Europe Economics analysis.

The figure below illustrates the same comparison graphically.

Figure 5.11: Comparison of electricity consumption – AutoML-assisted versus manual model development

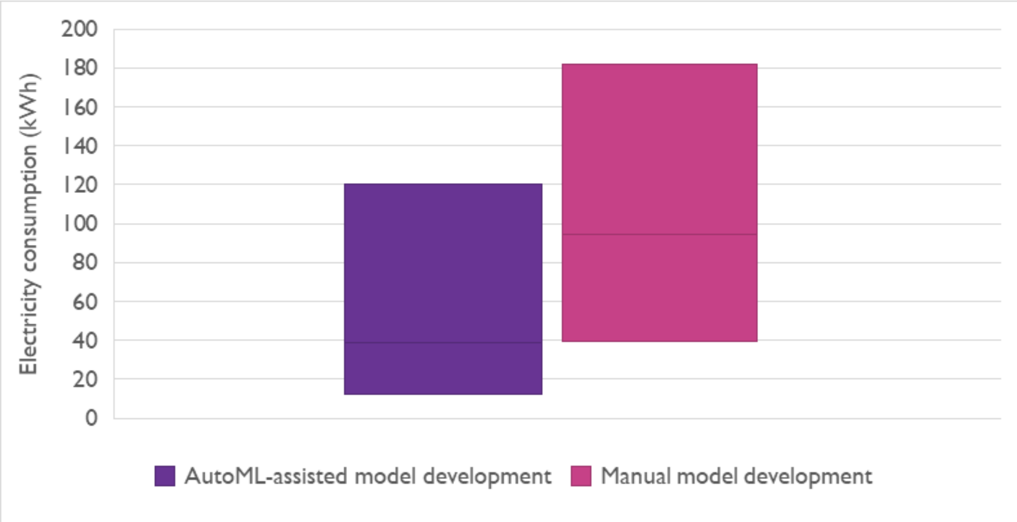

Source: Europe Economics analysis.

6 Conclusions

This study builds on Europe Economics' earlier work by further developing and applying a consistent framework for assessing the electricity consumption associated with digital services, compared with plausible counterfactuals. As in the earlier study, the analysis abstracts from any change in the scale of economic activity, and instead analyses how the electricity consumption of a specific activity depends on the way in which that activity is delivered.

Whereas the previous study examined three illustrative use cases, this report extends the framework to a further ten case studies, covering a wider range of digital services supported by data centres and their corresponding counterfactuals.

This chapter summarises the findings across the ten case studies, drawing on the results of the modelling. We present results for electricity consumption, the share of electricity consumed in centralised locations, and decarbonisation readiness. We also discuss the policy implications of our analysis.

6.1 Electricity Consumption

The table below summarises electricity consumption in the UK across the ten use cases under low, medium, and high assumptions.

Table 6.1: Total electricity consumption (kWh) in the UK

	Scenario	Low	Medium	High
Video streaming versus TV streaming	Digital	0.47	0.77	1.05
	Counterfactual	0.19	0.41	0.62
Cloud gaming versus gaming on computer	Digital	5.23	8.49	20.77
	Counterfactual (UK manufactured)	3.53	7.22	10.94
	Counterfactual (Overseas manufactured)	3.16	6.75	10.37
	Counterfactual (UK, decarbonised)	3.60	7.32	11.07
	Counterfactual (Overseas decarbonised)	3.18	6.78	10.41
Online education versus in-person education	Digital	3.17	3.97	4.81
	Counterfactual	3.99	6.10	9.24
	Counterfactual (decarbonised)	6.17	9.48	14.40
Cloud back-up versus local server back-up	Digital	0.42	0.95	1.89
	Counterfactual	1.25	4.88	17.25
AI-assisted email drafting versus manual email drafting	Digital	1.76	7.54	27.61
	Counterfactual	4.41	12.56	32.29
AI-assisted code development versus manual code development	Digital	4.76	18.84	101.89
	Counterfactual	10.58	31.39	145.31
AI-assisted copy-editing versus manual copy-editing	Digital	0.98	3.09	6.18
	Counterfactual	1.63	3.62	6.71
AI-assisted data extraction versus manual data extraction	Digital	0.28	0.88	2.61
	Counterfactual	0.49	1.31	2.97
AI-assisted transcription versus manual transcription	Digital	0.27	1.03	3.13
	Counterfactual	1.06	2.51	4.84
AI-assisted demand forecasting model development	Digital	11.96	38.91	120.39
	Counterfactual	39.68	94.18	181.64

Source: Europe Economics analysis.

As can be seen from the above table, the pattern of results differs between non-AI and AI use cases.

For non-AI use cases, there is a significant variation between the four use cases we studied in how electricity consumption under the digital approach compares with electricity consumption under the counterfactual. In the video streaming and cloud gaming, the digital approaches consume more electricity than their counterfactuals, reflecting the significant electricity requirements associated with data transmission in these use cases. By contrast, for online education and cloud back-up, the digital pathway results in lower electricity consumption than the counterfactual. For online education, this reflects the building electricity use and travel associated with in-person teaching, while for cloud back-up the difference is driven primarily by inefficiencies in long-term storage in local server environments. These cases highlight that digitalisation may not universally reduce electricity consumption and that results depend on the nature of the service.

We conducted further analysis on the potential UK electricity impact of shifting all current users of broadcast television (DTT) to video streaming, and found that this could increase UK electricity demand by around 1.8 TWh.

By contrast, our results suggest that when workers use AI to assist them in their jobs it leads to lower electricity consumption per task than in the counterfactual for all of the AI use cases that we considered, across low, medium and high scenarios. For each of these use cases, reductions in task duration caused by AI productivity gains translate into substantially lower office electricity use per task, more than offsetting the additional electricity consumption associated with AI training and inference and internet transmission.

As in the earlier study, the analysis distinguishes between total energy consumption and electricity use. Some of the counterfactuals rely partly on non-electric fuels, whereas digital services are almost entirely electricity-based. For the two use cases where this distinction materially affects the comparison (cloud gaming, and online education), we also present results for decarbonised counterfactuals in which all energy use is assumed to be electricity, reflecting what the comparison may look like as the energy system decarbonises.

To provide a concise comparison across use cases, the figure below presents results in terms of the ratio of electricity consumption in the digital scenario relative to electricity consumption in the counterfactual, using the medium scenario in each case. Expressing results as ratios allows comparisons to be made across services with very different absolute levels of electricity consumption and avoids the visual dominance of use cases with particularly large kWh values. Ratios above 100 per cent indicate that the digital approach consumes more electricity than the counterfactual approach, while ratios below 100 per cent indicate lower electricity consumption in the digital scenario.

Figure 6.1: Ratio of digital to physical electricity consumption (medium scenario)

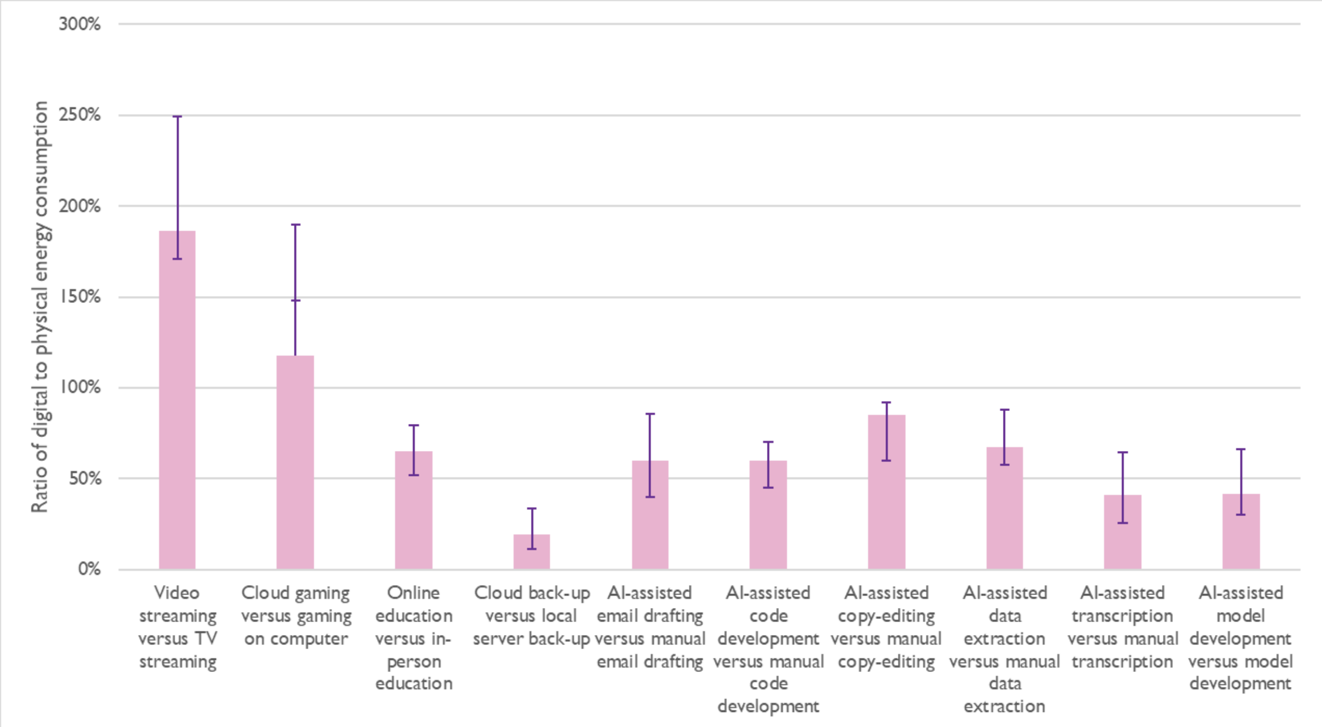

Note: The pink bars show ratios calculated using the medium scenario for both digital and counterfactual pathways. Error bars reflect the range obtained when comparing low digital with low counterfactual estimates and high digital with high counterfactual estimates. In some cases, the top end of the error bar may come from the comparison of the low estimates, while the low end of the error bar may come from comparison of the high estimates. In one case, cloud gaming versus gaming on computer, the ratio obtained using medium estimates is the lowest value, and so the uncertainty range for this use case extends only above the central estimate.

Source: Europe Economics analysis.

Overall, the results suggest that when workers use AI to increase their productivity it typically reduces electricity consumption per task, whereas for non-AI use cases digitalisation can either increase or reduce consumption depending on the nature of the service.

6.2 Centralised Electricity Use

A distinction can be made between centralised electricity use (e.g. data centres, warehousing, retail outlets) and decentralised use (e.g. offices, end-user devices, road freight, internet transmission). This distinction is important for understanding how the digital and counterfactual approaches to different use cases may place demands on the electricity system.

The table below presents the estimated share of electricity consumption that is centralised versus decentralised across use cases.

These results show that digital use cases generally involve a higher share of centralised electricity consumption than their physical counterfactuals, reflecting the role of data centres. This is particularly evident in cloud gaming and cloud back-up, where centralised electricity accounts for a substantial share of total consumption in the digital scenario. By contrast, counterfactuals for office-based services exhibit no centralised electricity use, as electricity consumption arises entirely from office buildings.

Among the digital use cases, AI-assisted tasks display the widest variation in centralised electricity shares. Some AI-assisted activities such as email drafting, copy-editing and code development have very low centralised shares, while data extraction, transcription and forecasting model development show higher ranges due to greater reliance on data centre processing.

Video streaming and online education sit between these extremes, with a small but material share of electricity consumed centrally.

Table 6.2: Centralised electricity consumption by use-case

Use-case	Digital	Counterfactual
Video streaming versus TV streaming	0.26–1.17	2.5–5.38
Cloud gaming versus gaming on computer	18.60–61.33	6–13
Online education versus in-person education	0.58–1.59	0
Cloud back-up versus local server back-up	77.18–92.40	0
AI-assisted email drafting versus manual email drafting	0.01–0.60	0
AI-assisted code development versus manual code development	0.01–0.17	0
AI-assisted copy-editing versus manual copy-editing	0.04–2.33	0
AI-assisted data extraction versus manual data extraction	0.11–8.14	0
AI-assisted transcription versus manual transcription	0.22–22.46	0
AI-assisted demand forecasting model development	0.07-24.51	0

Note: For cloud gaming, the counterfactual shown is the UK-manufactured case. Alternative counterfactual estimates for overseas manufactured and decarbonised cases are reported elsewhere in the report. Source: Europe Economics analysis.

6.3 Decarbonisation Readiness

Decarbonisation readiness assesses the extent to which current energy use in each scenario could feasibly be supplied by low-carbon electricity. This assessment considers only direct energy use within each scenario and does not include emissions associated with the production of capital equipment, including digital infrastructure. Digital use cases are, by construction, already fully electrified, while some counterfactuals may rely partly on fossil fuels (such as diesel and gas), including for manufacturing and transportation, whether these occur in the UK or overseas.

Among non-AI use cases, gaming on a physical device and in-person courses are the use cases in which material use is made of other fuels. The digital gaming scenario and online education scenarios are fully electrified and therefore scores 100 per cent on decarbonisation readiness. The counterfactuals to the digital gaming scenario also exhibit a very high degree of electrification, with decarbonisation readiness of around 98 per cent for UK-manufactured devices and around 99 per cent for overseas-manufactured devices across all assumptions. The small residual difference reflects limited non-electric energy use in manufacturing and transport for physical devices. In contrast, the in-person course has a materially lower decarbonisation readiness of around 29 per cent. This is driven by commuting, which in the representative scenario is assumed to take place by diesel car, leading to a substantial share of total energy use being derived from liquid fuels rather than electricity. However, we acknowledge that in practice it would also be possible for commuting to take place using an electric vehicle, in which case commuting would already be decarbonised.

The figure below compares decarbonisation readiness for the digital and counterfactual scenarios in cloud gaming and online education. For cloud gaming, the counterfactual shown reflects the UK-manufactured scenario.

Figure 6.2: Decarbonisation readiness in relevant use cases

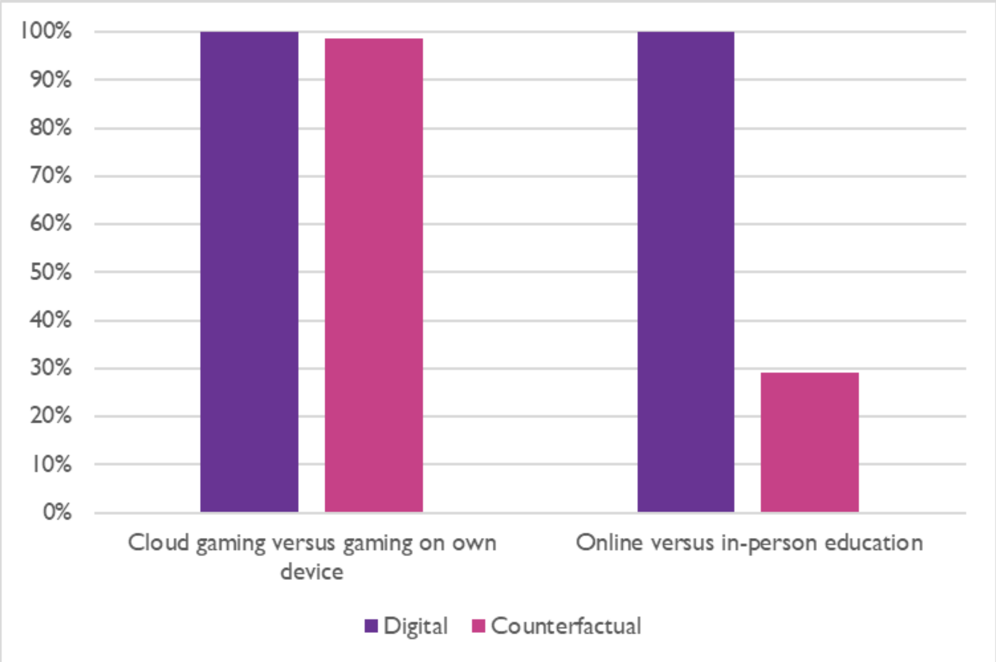

Source: Europe Economics analysis.

As a result, while the decarbonisation readiness metric confirms that both digital and physical approaches are largely compatible with a decarbonised energy system, it provides clearer differentiation in the education use case, where travel-related fuel use remains significant.

For AI-assisted use cases, the low value for office Electrical Energy Use Intensity (EUI) that we use in our modelling is based on offices which use another fuel (most probably gas for office heating), whereas the high value is based on all-electric offices.[94] This implies that, for these

[94] The same thing implicitly applies to the EUI for educational buildings that we use in the counterfactual to an online education course.

use cases, decarbonisation could move the electricity impacts of digitalization from the lower end of the range towards the upper end, as offices switch from gas heating to electric heating. It should be remembered, however, that other input assumptions also change between the low and high scenarios, and hence this does not mean the electricity impact will necessarily reach the upper bound of the range.

6.4 Policy implications

The results of this study suggest that the implications of digitalisation for UK electricity consumption are more nuanced than a simple increase in data centre demand. Digital services increase electricity use in data centres and internet transmission networks, but they can also reduce electricity consumption elsewhere in the economy by replacing electricity-intensive activities. The overall effect on UK electricity consumption therefore depends on the service in question and the nature of the counterfactual it replaces.

A consistent pattern emerges across the use cases considered in this report which involve workers using AI to assist them. In all of these cases, the AI-assisted approach consumes less electricity than the manual counterfactual under the low, medium and high assumptions. This reflects the fact that the additional electricity required for AI inference, training and other associated data centre activity is more than offset by reductions in task duration and hence lower office-based electricity consumption per task. Hence, the examples modelled in our analysis suggest that when workers use AI to increase their productivity, it can reduce the electricity required to deliver certain office-based tasks.

By contrast, the results for the non-AI use cases are mixed. Video streaming and cloud gaming consume more electricity than their counterfactuals, driven primarily by transmission requirements in the digital pathway. Online education and cloud back-up, however, consume less electricity than the counterfactual, reflecting the electricity associated with buildings, travel and less efficient local storage environments in the counterfactual case. The results for use cases that do not involve AI therefore illustrate that the electricity implications of digitalisation depend heavily on the specific service being examined, rather than pointing to a single general conclusion for non-AI digital activities as a whole.

The results also have implications for how growth in digital services may affect electricity demand at the system level. For example, the report’s additional analysis of a transition from DTT to OTT delivery suggests that replacing current terrestrial television viewing with streaming could materially increase UK electricity demand for content delivery, because streaming is more electricity-intensive per hour of viewing. Under the stylised assumption that all existing DTT viewing transfers to streaming (with no change in viewing behaviour), the report estimates a net increase in electricity demand of around 1.8 TWh per year.

It is also important to note that the analysis is conducted on a constant-activity basis. The modelling compares the electricity consumption of digital and counterfactual pathways for delivering the same level of activity, and does not consider whether digitalisation may increase total demand by making services cheaper, faster or more convenient. The results should therefore be interpreted as showing the direct electricity implications of substituting one mode of delivery for another, rather than the economy-wide effect of digitalisation if digital services are adopted more widely or used more intensively over time.

Appendix A: Detailed assumptions and results for the gaming on own computer

This appendix sets out the detailed assumptions and results for the gaming on own computer counterfactual. It supplements the summary description in the main report by providing the underlying modelling detail.

Input Assumptions

Overview of approach

This counterfactual involves a user playing a video game, using a video game disc, on the end-user's device, assumed to be a gaming console or a gaming PC. Energy consumption metrics are assumed across three components: manufacturing and storage, transportation, and end-user device usage. We account for the type of fuel used at each stage, distinguishing between electricity, gas, diesel and marine fuel, and consider how energy consumption varies depending on whether the product is manufactured domestically or overseas.

Fuel type

In order to differentiate between the fuel types associated with each step of energy consumption and to isolate electricity consumption, we multiply the energy consumption of each step by a percentage allocation to each fuel type.

Manufacture of raw materials

It is assumed that 80 per cent of the energy used in the production of raw materials such as plastic is from electricity. This reflects the use of large-scale electrically powered machinery for polymerisation, chemical processing, and refining. The remaining 20 per cent is attributed to gas, primarily for heating and other thermal processes.

Manufacture of video gaming discs

The manufacturing process – including injection moulding, stamping, and cooling – is assumed to be 90 per cent electricity-based, due to its reliance on high-precision automated machinery. The remaining 10 per cent is attributed to gas use for heating applications in specific manufacturing steps.

Burning of video onto disc

This stage is assumed to be fully powered by electricity (100 per cent), as laser burning and encoding rely entirely on electrically-operated optical drives.

Packaging

Packaging is assumed to use a mix of energy sources: 70 per cent electricity, supporting automated assembly lines; 20 per cent gas, for processes such as heat sealing; and 10 per cent diesel, reflecting the use of fuel-powered forklifts or material handling equipment within packaging facilities.

Warehouse operations and retail store

Warehouse energy consumption is assumed to be 80 per cent electricity and 20 per cent gas. Electricity supports lighting, conveyors, and cooling systems, while gas is used for space heating, particularly in colder climates.

Transportation fuel

All road-based transport, covering inbound logistics, ferry-port transfers, and retail distribution, is assumed to be powered entirely by diesel. For overseas manufacturing, ferry segments are assumed to use marine diesel exclusively, consistent with typical fuel use by roll-on/roll-off freight vessels.

Table A.1: Share of energy consumption by fuel type for various activities (%)

Activity	Electricity	Gas	Diesel	Marine Diesel
Manufacture of raw materials	80	20	0	0
Manufacture of disc	90	10	0	0
Burning of video onto disc	100	0	0	0
Packaging	70	20	10	0
Warehouse operations	80	20	0	0
Retail store energy	80	20	0	0
Transportation – road	0	0	100	0
Transportation – ferry	0	0	0	100
Retail distribution	0	0	100	0

Source: Europe Economics analysis.

Energy used in manufacturing and storage

The manufacture and storage of video gaming discs involves multiple stages, including raw material processing, disc manufacturing, video burning, packaging, and warehousing. Each step contributes to the overall energy footprint.

Energy manufacturing raw material – case and sleeve

The energy consumption values for manufacturing the gaming disc case and sleeve have been derived from Shehabi et al. (2014). Specific data for gaming disc cases were unavailable;

however, due to the similarity in materials used for both DVD and gaming discs cases (primarily polycarbonate and polypropylene), it is reasonable to assume comparable energy requirements. According to Shehabi et al. (2014), the energy consumption for manufacturing a DVD case is approximately 0.05004 kWh per unit. Similarly, the energy required for the sleeve is approximately 0.03614 kWh per unit. To account for variations in manufacturing conditions and material processing efficiency, a ±20 per cent adjustment has been applied to derive the low and high values.

Manufacturing energy per disc

The energy required to manufacture a video gaming disc itself has also been sourced from Shehabi et al. (2014), using DVD manufacturing data as a proxy. Both DVDs and video gaming discs share similar production processes, including injection moulding and metallisation. The study estimates that the energy consumption for manufacturing a DVD is approximately 0.3753 kWh per disc. A ±20 per cent range has been assumed to reflect potential differences in efficiency, material wastage, and process optimisations across different manufacturers.

Burning of video game onto disc

The electricity consumption for burning the video game onto a disc has been estimated based on the power consumption of SATA Blu-ray drives, which operate at 25 to 30 W during the burning process. Given that Blu-ray burning typically takes around 0.5 hours per disc, the total electricity consumption ranges from 0.0125 kWh (low, at 25 W) to 0.015 kWh (high, at 30 W). The medium value is calculated as 0.01375 kWh, derived from the average of the low and high estimates.

Packaging – shrink wrapping, cardboard boxes and pallet shrink wrapping

For the energy associated with packaging materials and processes, no directly applicable references were found. Instead, values have been estimated based on assumptions about material energy intensity and the typical energy cost of shrink-wrapping, corrugated cardboard production, and pallet wrapping. The medium values are intended to reflect typical industry practices, while the low and high values are intended to account for variations in material thickness, production efficiency, and automation levels.

Shrink wrapping is assumed to require 0.03 kWh per gaming CD (medium value), with a low estimate of 0.015 kWh and a high estimate of 0.045 kWh. Cardboard boxes, used for bulk transport and retail display, have an estimated energy requirement of 0.025 kWh per gaming CD (medium value), ranging from 0.018 to 0.032 kWh. We assume that pallet shrink wrapping, which is required for bulk shipping, has a medium energy estimate of 0.01 kWh per gaming CD, with a range of 0.007 to 0.014 kWh.

Warehouse storage energy per disc

The storage energy consumption per gaming CD has been derived from Seetharam et al. (2010), which analyses energy usage in warehouse operations. The study estimates that the energy consumption for storing a DVD in a warehouse is approximately 0.026 kWh per disc. A ±20 per cent variation has been applied to derive the low and high values, accounting for differences in warehouse efficiency, climate control energy usage, and storage density.

Retail store energy (lighting, heating, ventilation and air conditioning (HVAC), shelving)

To estimate the energy consumed in retail environments attributable to gaming disc sales, we draw on industry data on the energy intensity of commercial buildings and make assumptions about typical store operations and product distribution. The calculation begins with figures from the U.S. Energy Information Administration (EIA), which reports an average annual electricity consumption of approximately 13.5 kWh per square foot (equivalent to around 145 kWh per square metre) for retail stores.[95]

To estimate the retail energy use associated with gaming disc sales, we assume that a dedicated media section of 10 square metres is used to display and sell gaming discs. If this area supports the sale of 10,000 discs per year, this implies 0.001 m² of floor space per unit sold. We apply the above energy intensity of 145kWh per square per year. This yields an estimated retail energy use of 0.145 kWh per disc sold. To provide a conservative estimate, we round this figure down to 0.1 kWh per disc. To accommodate uncertainty in parameters such as store size, layout, lighting efficiency, HVAC operation, and product turnover, a ±20 per cent range is applied around this medium estimate.

Table A.2: Manufacturing and storage energy (kWh/CD)

Category	Low	Medium	High	Location
Manufacture of raw materials	0.069	0.086	0.103	UK/Overseas
Case	*0.04*	*0.05*	*0.06*	
Sleeve	*0.029*	*0.036*	*0.043*	
Manufacture of CD	0.3	0.375	0.45	UK/Overseas
Burning of game onto CD	0.011	0.0138	0.0165	UK/Overseas
Packaging	0.04	0.065	0.091	UK/Overseas
Shrink wrapping	*0.015*	*0.03*	*0.045*	
Cardboard boxes	*0.018*	*0.025*	*0.032*	
Pallet shrink wrapping	*0.007*	*0.01*	*0.014*	
Warehouse operations	0.022	0.026	0.031	UK
Retail store energy	0.08	0.1	0.12	UK

Note: The Location column shows where each stage is assumed to occur in the representative scenario. UK/Overseas means the stage is attributed to the country of manufacture, depending on whether manufacturing is assumed to take place in the UK or overseas. UK means the stage is always assumed to occur in the UK.

[95] EIA (2016), Electricity consumption totals and conditional intensities by building activity subcategories, 2012 [online]

Source: Europe Economics analysis.

Transportation

This section outlines the key assumptions used to estimate transportation energy consumption. The calculation framework accounts for multiple transportation stages, including inbound logistics (from the manufacturing facility to a distribution warehouse) and retail distribution (from the warehouse to retail stores). Key parameters include distance (km), fuel consumption intensity (L/km), energy content per litre (kWh/L), and the number of units transported per vehicle (units/truck and trucks/ferry).

To determine the transportation energy per unit, the following formula is applied:

$$kWh\ per\ unit = \left(Distance\ (km) \times Fuel\ Consumption\ \left(\frac{l}{km}\right) \times Energy\ content\ per\ litre\ \left(\frac{kWh}{l}\right)\right) \div\ Number\ of\ units\ per\ vehicle$$

Distance

Inbound logistics (manufacture to warehouse)

Inbound logistics can be classified into two scenarios: (1) when manufacturing occurs within the UK, and (2) when manufacturing takes place outside the UK (e.g., in Germany). Each scenario involves different transportation routes and distances.

UK manufacturing - road to warehouse

For gaming CDs produced in the UK, we assume they are transported from the manufacturing site to a warehouse for national retail distribution via road. The distance depends on the location of the factory and the position of the warehouse within the UK. A low estimate of 100 km is applied where both the manufacturing site and warehouse are within the same region, such as in the Midlands. A medium estimate of 150 km is applied for a typical route, such as from Enfield (North London) to a warehouse in the Midlands (e.g. Daventry or Milton Keynes). A high estimate of 200 km accounts for longer domestic transport, such as when manufacturing occurs in northern England or Scotland, requiring greater distances to reach distribution centres in the Midlands.

Overseas manufacturing (factory to warehouse via ferry)

For gaming CDs manufactured overseas, the supply chain involves multiple transport stages: (A) overland transport from the factory to a ferry port, (B) a ferry crossing to the UK, and (D) final road transport from the UK port to a distribution warehouse.

Overseas road (factory to ferry port)

Before shipping to the UK, gaming CDs must first be transported from the factory to the nearest ferry port in Europe. The distance depends on the factory location and its proximity to a major port. A low estimate of 300 km applies when the factory is near a major port, such as in western Germany with access to Rotterdam. A medium estimate of 500 km accounts for typical transport distances from central Germany to a ferry port. A high estimate of 700 km is used when manufacturing is farther east, such as in Poland or the Czech Republic, requiring a longer overland journey before reaching the ferry port.

Ferry distance (cross-channel)

The ferry distance depends on the port of departure in Europe and the UK port of entry. The shortest and most common ferry route, Dover-Calais, is 35 km, which applies as the low estimate. A medium estimate of 150 km accounts for longer ferry crossings such as Rotterdam to Harwich or Zeebrugge to Hull. The high estimate of 300 km represents significantly longer sea routes, such as Hamburg to Immingham or Liverpool, used when shipments take a North Sea route instead of the English Channel. For overseas-manufactured products, only a share of ferry fuel use is attributed to the UK; we assume 50 per cent in the central case, with a range of 40 to 60 per cent in the low and high cases.

UK road (ferry port to warehouse)

Once the shipment arrives in the UK, it must be transported by truck from the ferry port to a distribution warehouse. The distance depends on which UK port is used and the warehouse location. A low estimate of 100 km applies when the warehouse is near the port, such as in South East England. The medium estimate of 200 km accounts for a typical transport route to a Midlands warehouse, such as in Daventry or Milton Keynes, which are central logistics hubs. A high estimate of 300 km represents longer distances, such as when the shipment is transported from Dover to a warehouse in Manchester or Leeds.

Retail distribution (warehouse to store)

Once the gaming discs have arrived at the warehouse, they are distributed to retail stores within the UK. The distance varies depending on store locations, warehouse networks, and regional demand. A low estimate of 20 km represents a case in which the retail store is near the warehouse, typically in urban areas. The medium estimate of 50 km assumes that stores are located in regional retail hubs. The high estimate of 100 km accounts for longer distribution routes, such as deliveries to rural locations or smaller towns.

Table A.3: Assumed transport distances for gaming discs

Category	Unit	Low	Medium	High
Inbound logistics (manufacture to warehouse)				
UK manufacturing				
UK road (factory to warehouse)	km	100	150	200

Overseas manufacturing				
Overseas road (factory to ferry port)	km	300	500	700
Ferry distance (cross-channel)	km	35	150	300
UK road (ferry port to warehouse)	km	100	200	300
Retail distribution (warehouse to store)				
UK road (warehouse to store)	km	20	50	100

Source: Europe Economics analysis.

Fuel consumption intensity

To estimate the energy use associated with transportation, the model draws on fuel consumption intensity (FCI) values for both road and maritime modes. These represent the amount of fuel consumed per kilometre travelled and are a critical input for calculating transport energy per gaming disc.

For road freight, the medium estimate is taken from the International Council on Clean Transportation (ICCT), which reports that a typical long-haul tractor-trailer in Europe consumes approximately 32.6 litres of diesel per 100 kilometres (L/100 km).[96] Converted to a per-kilometre basis, this equates to 0.326 L/km. To reflect variability in vehicle efficiency, road conditions, and driving behaviour, a ±20 per cent adjustment is applied. This yields a lower bound of 0.2608 L/km for more efficient operations and an upper bound of 0.3912 L/km for less efficient scenarios.

For ferry transport, the model uses data from ARUP (2014), which provides typical fuel consumption values for roll-on/roll-off (Ro-Ro) freight vessels operating on cross-channel routes.[97] The medium case assumes a fuel consumption intensity of 5.54 L/km, based on reported averages for these vessels. Again, a ±20 per cent range is applied to account for variations in vessel type, sea conditions, and operational practices. This results in a low estimate of 4.432 L/km for high-efficiency or optimally loaded ferries, and a high estimate of 6.648 L/km for older vessels or less efficient operations.

Table A.4: Energy content per litre (kWh/L)

Category	**Low**	**Medium**	**High**
Truck energy fuel consumption intensity	0.2608	0.326	0.3912
Ferry energy fuel consumption intensity	4.432	5.54	6.648

Source: Europe Economics analysis.

Energy content per litre

[96] ICCT(2018), Comparison of fuel consumption and emissions for representative heavy-duty vehicles in Europe [online]
[97] ARUP (2014), Cost of Emissions for NSW Ferry Networks [online]

To convert fuel consumption into energy use, the model draws on the lower calorific value (LCV) of diesel and marine diesel fuels. These values represent the amount of usable energy released during combustion and are expressed in kilowatt-hours per litre (kWh/L).

For road diesel, the Oak Ridge National Laboratory (ORNL, 2022) reports an LCV of 38.6 megajoules per litre (MJ/L).[98] Using the standard conversion factor of 1 MJ = 0.2778 kWh, this equates to an energy content of 10.72 kWh/L.

For marine diesel, the International Maritime Organization (IMO) provides a lower calorific value of 42,700 kilojoules per kilogram (kJ/kg), alongside a typical fuel density of 0.8744 kilograms per litre (kg/L).[99] Multiplying these gives a volumetric energy content of 37.35 MJ/L, which converts to 10.38 kWh/L using the same MJ-to-kWh factor.

These energy content values are used to convert fuel consumption intensity (FCI, in litres per kilometre) into energy intensity (EI, in kilowatt-hours per kilometre):

$$EI\left(\frac{kWh}{km}\right) = FCI\left(\frac{L}{km}\right) \times Energy\ content\left(\frac{kWh}{L}\right)$$

Table A.5: Energy efficiency for decarbonised scenario (kWh/L)

	Energy content per litre
Diesel energy content per litre	10.72
Marine diesel energy content per litre	10.68

Source: ORNL (2022), IMO (2016)

Decarbonised scenario

We also modelled a decarbonised scenario in which fossil fuels are replaced with electricity. To estimate the energy consumption in the decarbonised scenario, the model adopts the following assumptions for road and marine transport, reflecting the expected performance of electric HGVs and hydrogen-powered vessels based on emerging real-world data.

For electric HGVs, we use data from the UK’s Battery Electric Truck Trial (BETT, 2024) to model energy use for electric HGVs.[100] The trial, based on 287,000 km of real-world fleet data, reported an average efficiency of 1.08 km per kWh (0.93 kWh/km), which we adopt as our medium estimate. Urban stop-start driving reduced efficiency to 0.90 km per kWh (1.11 kWh/km, high case), while rural routes achieved 1.20 km per kWh (0.83 kWh/km, low case). While higher efficiency figures have been reported in test-cycle and idealised conditions, we consider that the BETT trial offers the most robust dataset for UK fleet use. We therefore consider that these figures provide a representative basis for modelling energy consumption under a decarbonised transport scenario.

98 ORNL (2022), Transportation Energy Data Book [online]
99 IMO's MEPC.281(70) Report [online]
100 Cenex (2024) “BETT: Battery Electric Truck Trial Final Report” [online]

For maritime transport, we assume the use of a hydrogen-powered vessel. Based on published specifications for commercially piloted models,[101] we assume that the hydrogen-powered vessel has two 240 kW fuel cells and a 380 km range per refill. Operating at full power, this yields a medium estimate of 12.6 kWh/km. We assume 9.5 kWh/km for the low case (slower cruising or higher efficiency) and 16.8 kWh/km for the high case (less efficient operation or longer trip duration).

Number of video game CDs per truck

To determine transport energy per video gaming disc, it is necessary to estimate the number of units that can be transported within a standard freight vehicle. In the case of ferry transport, it is also necessary to estimate the number of freight vehicles that can be transported on a ferry. A 40-foot lorry typically offers an internal cargo volume of 67.7 cubic metres. Given that a standard video game disc case occupies approximately 0.00035 cubic metres, this implies a theoretical maximum capacity of around 193,000 units per lorry.

However, practical constraints such as palletisation, stacking inefficiencies, and space reserved for manoeuvring and packaging reduce the usable volume significantly. Assuming that only 50 per cent of the internal space is effectively utilised, the capacity falls to 96,500 gaming discs per lorry. After applying a further downward adjustment to account for additional packing inefficiencies, a rounded medium estimate of 90,000 discs per lorry is adopted.

For ferry transport, the model assumes a large roll-on/roll-off (Ro-Ro) ferry capable (in the medium scenario) of accommodating 120 lorries. This is consistent with specifications from major ferry operators. A typical large Ro-Ro ferry includes three to five decks and offers a total lane length of between 2,000 and 3,500 metres. Given that a 40-foot lorry occupies approximately 17 metres of lane length, the assumed capacity of 120 lorries per vessel aligns with real-world configurations after accounting for manoeuvring space and deck utilisation.

To reflect uncertainties in loading efficiency, pallet dimensions, and real-world constraints, a sensitivity range is introduced: the number of discs per lorry is assumed to vary between 70,000 and 110,000, while the number of lorries per ferry is assumed to vary between 100 and 140. This variation is important for energy calculations, as a lower number of units per lorry and lorries per ferry results in higher transport energy per video gaming disc, and vice versa.

Table A.6: Number of gaming CDs

Category	Low	Medium	High
Discs per truck	110,000	90,000	70,000
Trucks per ferry	140	120	100

Source: Europe Economics analysis.

[101] Global Times (2024) “China launches hydrogen-powered container ship, capable of sailing 380 kilometers” [online]

End-user device

Users may access the game through a range of devices with different power requirements. To reflect this, we apply a range of electricity consumption values: the low estimate corresponds to the low-intensity gaming console assumption, while the high estimate reflects gaming on a gaming PC. The medium figure is based on the medium-intensity gaming console assumption. No additional hardware is assumed. No additional hardware is assumed. Additionally, we account for the electricity consumed during installation, which we estimate using the assumption that 1.6 GB of data is transferred per minute during the installation process.

Results

This section presents the detailed assumptions, calculation steps and results underpinning the gaming on own computer counterfactual.

UK manufacturing

For products manufactured within the UK, all energy-intensive processes, including raw material processing, manufacturing, packaging, warehousing, and retail operations, contribute to the UK's energy consumption footprint. Transportation energy is also fully accounted for, covering road transport from the manufacturing facility to warehouses and from warehouses to retail stores.

Manufacturing and storage

The energy consumption for manufacturing and storage is determined using figures for the energy use for each stage of production, allocated across fuel types using an assumed percentage breakdown by type of fuel. In this way, the results distinguish between electricity use and other fuel types (gas and diesel) when assessing the energy footprint of gaming disc production.

The table below sets out electricity consumption for manufacturing and storage when gaming discs are produced within the UK. The majority of electricity consumption comes from disc manufacturing, followed by raw material processing and retail store electricity use. Warehousing and packaging contribute a smaller share.

Table A.7: Energy Consumption for Manufacturing and Storage (Manufactured within the UK) – Electricity (kWh/CD)

	Low	Medium	High
Manufacture of raw materials	0.055	0.069	0.083
Manufacture of gaming CD	0.27	0.0.338	0.405
Burning video onto gaming CD	0.011	0.014	0.017
Packaging	0.028	0.046	0.064

Warehouse operations	0.017	0.021	0.025
Retail store electricity	0.064	0.08	0.096
Total	0.445	0.567	0.689

Source: Europe Economics analysis.

The table below presents the energy consumption from other fuel sources, including gas and diesel. Gas is primarily used in raw material processing, manufacturing, and packaging, particularly for heating and chemical processes, while diesel consumption is minimal, appearing only in packaging operations. In total, consumption of other fuels is very small compared with consumption of electricity.

Table A.8: Energy Consumption for Manufacturing and Storage (Manufactured within the UK) – other fuel (kWh/CD)

	Low	Medium	High
Gas			
Manufacture of raw materials	0.0138	0.0172	0.0207
Manufacture of disc	0.03	0.0375	0.045
Burning of video onto disc	-	-	-
Packaging	0.008	0.013	0.0182
Warehouse operations	0.0041	0.0052	0.0062
Retail store energy	0.016	0.02	0.024
Diesel			
Packaging	0.004	0.0065	0.0091
Total	0.076	0.099	0.123

Source: Europe Economics analysis.

Transportation

The table below presents the energy consumption for transportation when gaming CDs are manufactured within the UK. The transportation energy is calculated by multiplying the distance travelled (km) by the fuel consumption intensity per kilometre (L/km) and the energy content per litre of fuel (kWh/L), then dividing by the number of gaming CDs transported per vehicle to obtain the per-unit energy consumption.

For UK-manufactured CDs, all transportation occurs by road, with no marine transportation involved. The total transport energy consists of inbound logistics (from the factory to the warehouse) and retail distribution (from the warehouse to stores). We have assumed that all road transport relies on diesel, meaning that all transport energy in this scenario is attributed to diesel consumption.

The results show that transportation energy per gaming disc is minimal, particularly when compared to the energy required for manufacturing and storage. This is because the high number of gaming discs transported per vehicle significantly reduces per-unit transport energy. Inbound logistics from the factory to the warehouse contributes the largest share of transportation energy, while retail distribution has a smaller impact due to shorter distances.

Table A.9: Energy Consumption for Transportation (Manufactured within the UK) – other fuel (kWh/CD)

	Low	Medium	High
Inbound logistics			
Factory to warehouse	0.0025	0.0058	0.0119
Retail distribution			
Warehouse to store	0.0005	0.0019	0.006
Total			
Total road transportation inside UK (diesel)	0.003	0.0078	0.018
Marine transportation inside UK (marine diesel)	0	0	0

Source: Europe Economics analysis.

In the decarbonised scenario, all road transport is assumed to be carried out by electric trucks. The transportation electricity is calculated by multiplying the distance travelled (km) by the electricity efficiency of electric trucks (kWh/km), then dividing by the number of gaming CDs transported per vehicle to obtain the per-unit electricity consumption.

Table A.10: Electricity Consumption for Transportation (Manufactured within the UK) – Decarbonised (kWh/CD)

	Low	Medium	High
Inbound logistics (factory to warehouse)	0.00075	0.0016	0.0032
Retail distribution (warehouse to store)	0.00015	0.0005	0.0016
Total - road	0.0009	0.0021	0.0048

Source: Europe Economics analysis.

End-user device

Similarly to the digital case, the electricity consumption for end-user devices is calculated by multiplying the hours of use by the power draw of the end-user device. We also account for electricity used during the installation process.

End-user device electricity consumption remains the same regardless of whether the gaming CD was manufactured in the UK or overseas and is assumed to be entirely electricity-based.

Table A.11: Electricity consumption for end-user device – Electricity (kWh)

Units	Low	Medium	High
Electricity used by device	3.04	6.52	10.00

Installation electricity	0.04	0.18	0.42
Total device electricity	**3.08**	**6.70**	**10.42**

Source: Europe Economics analysis.

Summary of energy use

Electricity

We only consider electricity consumption within the UK, and consider a scenario in which the gaming disc is manufactured domestically and a scenario in which it is manufactured overseas.

Since all transportation energy is assumed to be diesel or marine diesel, the electricity consumption in this scenario includes only the manufacturing and storage electricity consumption (for UK manufacturing with the relevant electricity apportioning) and the end-user device electricity consumption.

Electricity consumption for the representative scenario ranges from 3.5 kWh to 11.1 kWh, with manufacturing, packaging, and storage making up a significant share in the low and medium estimates. However, in the high scenario, end-user device consumption becomes the dominant factor, reflecting the variability in device electricity efficiency.

Table A.12: Total energy for representative scenario (kWh) - (Manufactured within the UK) – Electricity

Component	Low	Medium	High
Electricity for manufacturing, packaging, storage, and retailing	0.445	0.567	0.689
Electricity for end-user device	3.081	6.702	10.417
Total electricity for representative scenario	3.526	7.269	11.106

Source: Europe Economics analysis.

When distributed over 20 hours of assumed gaming, electricity use per hour of gaming ranges from 0.18 kWh to 0.55 kWh. End-user device electricity contributes a significant share, increasing with higher power draw assumptions.

Table A.13: Electricity per hour of video gaming - (Manufactured within the UK) – Electricity

Component	Low	Medium	High
Electricity for manufacturing, packaging, storage, and retailing	0.022	0.028	0.034
Electricity for end-user device	0.154	0.335	0.521
Total electricity per hour of video watching	**0.176**	**0.363**	**0.555**

Source: Europe Economics analysis.

Table below presents the percentage breakdown of electricity consumption.

Table A.14: Breakdown of electricity consumption (%) - (Manufactured within the UK) – Electricity

Component	Low	Medium	High
Electricity for manufacturing, packaging, storage, and retailing	13	9	6
Electricity for end-user device	87	92	94

Source: Europe Economics analysis.

In the UK manufacturing scenario, we classify electricity used in production, warehousing, and retail as centralised and electricity for end-user devices as decentralised. Thus, the share of electricity consumption occurring in centralised locations ranges from **6 to 13 per cent**.

All fuel

The *"all fuel"* scenario accounts for all fuel types (electricity, gas, diesel, and marine diesel) consumed across manufacturing, storage, transportation, and end-user device usage. All transportation energy is assumed to be diesel or marine diesel, the manufacturing and storage energy includes electricity, gas and diesel consumption, while end-user device energy remains entirely electricity-based.

Table below presents the total energy consumption, with energy from end-user devices being the dominant energy source. Gas contributes to manufacturing and storage processes, but its share is relatively small. Diesel is mainly used for road transportation, making its impact minor due to the high number of gaming CDs transported per vehicle. Marine diesel consumption is zero, as UK manufacturing does not involve maritime shipping.

Table A.15: Total energy for representative scenario (kWh) - (Manufactured within the UK) – All Fuel

Component	Low	Medium	High
Energy for manufacturing, packaging, storage, and retailing - electricity	0.445	0.5667	0.6892
Energy for manufacturing, packaging, storage, and retailing - other fuel	0.076	0.0994	0.1233
Energy for transportation - road - other fuel	0.003	0.0078	0.018
Energy for transportation - marine - other fuel	0	0	0
Energy for end-user device - electricity	3.081	6.702	10.417
Total energy for representative scenario	**3.605**	**7.376**	**11.247**

Source: Europe Economics analysis.

Table below shows energy consumption per hour of video watching, derived by dividing the total energy by the assumed 20-hour gaming period. Energy from end-user devices remains the largest contributor, while other energy contributions remain relatively minimal. End-user device energy consumption is significantly higher in the high-energy scenario, making it a key driver of total energy use.

Table A.16: Energy per hour of video watching (kWh)- (Manufactured within the UK) – All Fuel

Component	Low	Medium	High
Energy for manufacturing, packaging, storage, and retailing - electricity	0.1112	0.1417	0.1723
Energy for manufacturing, packaging, storage, and retailing - other fuel	0.0189	0.0248	0.0308
Energy for transportation - road - other fuel	0.0008	0.0019	0.0046
Energy for end-user device- electricity	0.154	0.335	0.521
Total energy per hour of video gaming	**0.180**	**0.369**	**0.562**

Source: Europe Economics analysis

The table below provides a breakdown of energy consumption across different components. Energy form end-use devices remains the dominant energy source, accounting for over 90 per cent of total energy consumption. Electricity from manufacturing, packaging and storage is the second component when it comes it its importance, remaining low in comparison to end-user device energy – between 6 and 7 per cent of the total. Other fuels (mainly gas) used in manufacturing and storage account for only a small fraction, highlighting their limited role. Diesel for road transport remains negligible (0 per cent), reinforcing that logistics have a minor impact on the total energy footprint.

Table A.17: Breakdown of energy consumption (%) - (Manufactured within the UK) – All Fuel

Component	Low	Medium	High
Energy for manufacturing, packaging, storage, and retailing - electricity	12	8	6
Energy for manufacturing, packaging, storage, and retailing - other fuel	2	1	1
Energy for transportation - road - other fuel	0	0	0
Energy for end-user device- electricity	85	91	93

Source: Europe Economics analysis.

Under the UK manufacturing scenario (all fuel), we classify energy used in production, warehousing, and retailing as centralised. Transport fuel and end-user device electricity are treated as decentralised. The proportion of total energy consumption that occurs in centralised locations ranges from 7 to 14 per cent.

Decarbonised scenario

In the decarbonised case, all processes, including transport and heating, are assumed to be powered by low-carbon electricity. As a result, all energy consumption is reported in kWh of electricity, regardless of whether it would previously have been diesel or gas. The figures below therefore reflect the total grid-relevant electricity demand under a fully electrified supply chain.

Table below presents total electricity consumption under the decarbonised scenario. End-user device electricity remains the largest contributor.

Table A.18: Total electricity for representative scenario (kWh) - UK Manufactured – Decarbonisation

	Low	Medium	High
Electricity for manufacturing, packaging and storage	0.521	0.666	0.812
Electricity for transportation - road	0.001	0.002	0.005
Electricity for end-user device	3.081	6.702	10.417
Total electricity for representative scenario	**3.603**	**7.370**	**11.234**

Source: Europe Economics analysis.

Table below shows electricity consumption per hour of video watching, derived by dividing total electricity by the assumed 20-hour gaming period.

Table A.19: Electricity per hour of video gaming (kWh) - (UK Manufactured) – Decarbonisation

	Low	Medium	High
Electricity for manufacturing, packaging and storage	0.026	0.033	0.0.41
Electricity for transportation - road	0.0000	0.0001	0.0002
Electricity for end-user device	0.154	0.335	0.521
Total electricity for representative scenario	**0.180**	**0.369**	**0.562**

Scenario: Europe Economics analysis.

To illustrate the relative contributions of each component, the table below presents the percentage breakdown of total electricity consumption across the three scenarios.

Table A.20: Breakdown of electricity consumption (%) - (UK Manufactured) – Decarbonisation

	Low	Medium	High
Electricity for manufacturing, packaging and storage	14	9	7

Electricity for transportation - road	0	0	0
Electricity for end-user device	86	91	93
Total electricity for representative scenario	100	100	100

Source: Europe Economics analysis.

Under the UK manufacturing decarbonised scenario, we classify electricity used in production, warehousing, and retailing as centralised. Transport fuel and end-user device electricity are treated as decentralised. The proportion of total electricity consumption that occurs in centralised locations ranges from **7 to 14 per cent**.

Overseas Manufacturing

For products manufactured overseas, we only consider the energy consumption directly associated with UK-based activities in the analysis. This consists of warehousing, retail operations, and the domestic portion of transportation, which includes road freight from the UK port to warehouses and final retail distribution. Additionally, the energy from maritime transport is allocated between the UK and overseas. The energy used in raw material processing, manufacturing, and packaging are excluded as it remains outside the UK's energy footprint, as these processes occur abroad.

Manufacturing and Storage

For gaming discs manufactured overseas, the UK energy consumption for manufacturing and storage is limited to warehouse operations and retail store energy use, as all other production processes occur outside the UK.

Electricity consumption, is primarily from retail store operations, driven by lighting, HVAC, and shelving systems, while warehouse operations consume less electricity.

Table A.21: Energy Consumption for Storage and Retailing (Manufactured Overseas) – Electricity (kWh/CD)

	Low	Medium	High
Warehouse operations	0.01664	0.0208	0.02496
Retail store electricity	0.064	0.08	0.096
Total	0.0806	0.1008	0.121

Source: Europe Economics analysis.

Gas consumption follows a similar pattern, with retail store heating accounting for the majority and warehouse operations using a smaller share for climate control.

Table A.22: Energy Consumption for Storage and Retailing (Manufactured Overseas) – gas (kWh/CD)

	Low	Medium	High

Warehouse operations	0.00416	0.0052	0.00624
Retail store energy	0.016	0.02	0.024
Total	0.02016	0.0252	0.03024

Source: Europe Economics analysis.

Transportation

The transportation energy consumption is calculated using the same methodology across all transport stages, with energy derived from the distance travelled, fuel consumption intensity, energy content per litre of fuel and the number of gaming discs transported per vehicle. For ferry transport, energy consumption is calculated using marine diesel, factoring in ferry fuel consumption intensity, energy content per litre and the number of lorries per ferry.

- The UK allocation of **marine transportation energy (marine diesel)** represents the share of ferry fuel attributed to the UK. It is calculated as the total ferry energy consumption multiplied by the proportion allocated to the UK. Due to the high cargo capacity of ferries, the per-unit energy consumption is low.
- Road transportation inside the UK (diesel) includes transport from the UK ferry port to the warehouse and final retail distribution to stores. Since UK distances are generally shorter than overseas transport, this component contributes less to the total footprint.

The results show that UK road transport contributes the most to the total UK transport energy consumption, as it includes both inbound logistics from the UK port to the warehouse and final retail distribution to stores. Marine transport energy allocated to the UK has a minimal per-unit energy impact, as the ferry energy consumption is allocated across large cargo shipments.

Table A.23: Energy Consumption for Transportation (Manufactured Overseas) – other fuel (kWh/CD)

	Low	Medium	High
Inbound logistics			
Ferry (cross-channel)	0.0001	0.0008	0.0029
UK road (ferry port to warehouse)	0.0025	0.0078	0.018
Retail distribution			
Warehouse to store	0.0005	0.002	0.006
Total			
Marine transportation allocated to the UK (marine diesel)	0.0000418	0.0004	0.0018
Road transportation inside UK (diesel)	0.003	0.0097	0.024

Source: Europe Economics analysis.

In the decarbonised scenario, all road transport and marine transport are assumed to be carried out by electric trucks and hydrogen powered vessels. The transportation energy is calculated by multiplying the distance travelled (km) by the electricity efficiency of electric trucks or hydrogen powered vessels (kWh/km), then dividing by the number of gaming discs transported per vehicle to obtain the per-unit electricity consumption.

Table A.24: Electricity Consumption for Transportation (Manufactured Overseas) – Decarbonised (kWh/CD)

	Low	Medium	High
Inbound logistics			
Ferry (cross-channel)	$2.15*10^{-5}$	0.00018	0.0007
UK road (ferry port to warehouse)	0.00075	0.0021	0.0048
Retail distribution			
Warehouse to store	0.00015	0.0005	0.0016
Total			
Marine transportation allocated to the UK	8.63×10^{-6}	0.00009	0.00043
Road transportation inside UK	0.00009	0.0026	0.0063

Source: Europe Economics analysis.

End-User Device

End-user device electricity consumption is the same for the UK manufacturing case.

Appendix B: Impact of excluding AI training energy

This appendix presents results excluding electricity consumption associated with large-scale AI model training (or equivalently, electricity associated with development of the AutoML system in the case of AI-assisted demand forecasting model development).

As noted in the section 3.2.1, large-scale AI training is typically undertaken outside the UK due to relatively high domestic electricity costs. In the core analysis presented in the main report, AI training energy is included and allocated across the user base. This results in a relatively small contribution to electricity consumption for the representative user in most use cases.

The table below compares results with and without AI training energy across all use cases and scenarios. In almost all cases, excluding AI training energy makes little or no difference to the results at the number of decimal places reported. The main noticeable difference arises in the high scenario for AI-assisted demand forecasting model development, where the per-user allocation of energy developing the AutoML system is larger.

Table B.1: Electricity consumption for AI-assisted use cases, with and without AI training energy (kWh)

	Scenario	Low	Medium	High
AI-assisted email drafting versus manual email drafting	Without AI training	1.76	7.54	27.46
	With AI training	1.76	7.54	27.61
AI-assisted code development versus manual code development	Without AI training	4.76	18.84	101.76
	With AI training	4.76	18.84	101.89
AI-assisted copy-editing versus manual copy-editing	Without AI training	0.98	3.08	6.06
	With AI training	0.98	3.09	6.18
AI-assisted data extraction versus manual data extraction	Without AI training	0.28	0.88	2.42
	With AI training	0.28	0.88	2.61
AI-assisted transcription versus manual transcription	Without AI training	0.27	1.03	3.13
	With AI training	0.27	1.03	3.13
AI-assisted demand forecasting model development	Without AutoML system development	11.95	38.70	104.93
	With AutoML system development	11.96	38.91	120.39

Note: in the case of AI-assisted demand forecasting model development, the data centre electricity consumed in developing the AutoML system is treated conceptually as the equivalent of AI training energy.
Source: Europe Economics analysis.

This publication is available from: https://www.gov.uk/government/publications/impact-of-growth-of-data-centres-on-energy-consumption