

THE BANKING SECTOR'S OWN AI FOOTPRINT

The energy and carbon impacts of artificial intelligence applications in banking operations

*Péter Bagó*¹

ABSTRACT

Over the past decade, the banking sector has become a key actor in the sustainability transition through the measurement of financed emissions (Scope 3) and the decarbonisation of portfolios. At the same time, however, artificial intelligence (AI) solutions have been proliferating rapidly in the operations of credit institutions – from customer-service chatbots and credit-risk models to agentic systems that automate multiple steps of business workflows. The energy and resource requirements of these applications constitute a dynamically growing component of the sector's own environmental footprint that has so far received little attention. This study has three objectives. First, it reviews the international literature and measurement methodologies concerning the lifecycle environmental impacts of AI systems, including training, inference, embodied emissions, and water and resource consumption. Second, it examines how current trends in the adoption of AI in banking – particularly the growing use of generative and agentic AI – affect credit institutions' Scope 2 and Scope 3 emissions categories and, through an exploratory document analysis of the 2025 disclosures of nine credit institutions, provides empirical evidence that AI-related energy consumption cannot currently be identified in prevailing reporting practices. Third, it proposes a methodological framework for quantifying the “banking AI footprint” and integrating it into sustainability reporting (CSRD/ESRS). The study argues that, alongside the narrative of “greening by AI”, the sector must also institutionalise a “greening of AI” perspective: without energy-related criteria being incorporated into AI procurement decisions, model sizing, and cloud-service agreements, digital transformation may undermine the credibility of banks' own climate targets.

JEL codes: G21, O33, Q54, Q56, M14

Keywords: artificial intelligence, banking sector, carbon footprint, data centres, ESG, sustainability reporting, CSRD, EU AI Act

¹ Péter Bagó Assistant Professor, Corvinus University of Budapest, Institute of Entrepreneurship and Innovation, Department of Innovation and Business Incubation. E-mail: peter.bago@uni-corvinus.hu.

1 INTRODUCTION

Over the past decade, the debate on the role of the financial sector in climate policy has focused almost exclusively on the *indirect* environmental impact of banks arising from their financing activities. This is understandable: the greenhouse gas (GHG) emissions profile of credit institutions is dominated by so-called financed emissions associated with lending and investment portfolios (Scope 3, Category 15), whereas direct emissions from their own operations (Scope 1) and indirect emissions associated with energy consumption (Scope 2) generally account for a negligible share (PCAF, 2022). Previous research analysing the ESG disclosure practices of major Hungarian banks has likewise focused primarily on portfolio-level risks and the substantive maturity of disclosures (Vőneki–Lamanda, 2020), while the literature on green financial innovation has typically treated technology as an *instrument* for achieving sustainability objectives rather than as a *source* of environmental impacts (Horváth, 2022).


This approach, however, is becoming increasingly incomplete. The digital transformation of the banking sector – the strategic framework for which had already been outlined in Hungarian academic literature prior to the emergence of generative AI (Vives, 2019) – entered a new phase with the generative artificial intelligence wave that began in late 2022, making the financial sector one of the most intensive corporate adopters of the technology (ECB, 2024). The international institutional literature identified the opportunities and risks associated with the use of AI in banking at an early stage (EBA, 2020; FSB, 2017; Königstorfer–Thalmann, 2020; OECD, 2021), while recent systematic reviews document a broad spectrum of applications ranging from bankruptcy prediction and risk management to the optimisation of customer relationships (Aldasoro et al., 2024; Cao, 2022; Fares et al., 2023; Lülök–Sebestyén, 2025). According to data from the Evident AI Index, which measures the AI maturity of the world’s largest banks, the number of publicly announced AI applications in banking has been increasing quarter by quarter, and in the first quarter of 2026, nearly one-third of new use cases were so-called agentic applications that automate multiple steps of a workflow rather than a single task (Evident, 2026).

Meanwhile, the energy requirements of the computing infrastructure underpinning AI applications have become one of the fastest-growing components of global electricity demand. The long-term growth trajectory of energy consumption by information and communication technologies (ICT) had already been identified in the literature in the mid-2010s (Andrae–Edler, 2015; Belkhir–Elmeligi, 2018; Jones, 2018; Malmodin–Lundén, 2018), while several authors had also drawn attention to the depletion of efficiency gains and to methodological debates surrounding such estimates (Freitag et al., 2021; Koomey et al., 2011; Masanet et al.,

2020; Mytton–Ashtine, 2022). Generative AI has altered this trajectory: according to estimates by the International Energy Agency (IEA), global electricity consumption by data centres could increase from approximately 415 TWh in 2024 to around 945 TWh by 2030, with AI workloads accounting for the largest share of this growth (IEA, 2025); in the United States, data centres' share of national electricity consumption could rise from 4.4% in 2023 to as much as 12% by 2028 (Shehabi et al., 2024). The projected trajectory is illustrated in *Figure 1*.

Figure 1

Global electricity consumption of data centres and projected trajectory, 2020–2030 (TWh/year)



Source: Author's compilation based on IEA (2025); the 2024 and 2030 values are IEA estimates, while the intervening years represent an interpolated trajectory.

At the intersection of these two trends lies the subject of this study: the banking sector's *own* AI footprint, i.e. the environmental impacts – primarily in terms of energy, carbon, water and resource use – arising over the full lifecycle of AI systems developed, procured and operated by credit institutions. The issue is timely for several reasons. First, although banks' own operational emissions are small in absolute terms compared with financed emissions, the pace of AI adoption may alter the structure and dynamics of both Scope 2 emissions and Scope 3 emissions associated with purchased services, with implications for the achievability of net-zero commitments concerning own operations (*net zero own operations*). Second, the European Union's sustainability reporting (CSRD/ESRS) and AI regulatory frameworks (Regulation (EU) 2024/1689, hereinafter as: the AI Act) both contain – in some respects still evolving – energy-related transparency requirements that

also apply to banks as AI deployers. Third, from the perspective of the sector's reputational risk, it is not sustainable in the long term for an asymmetry to persist whereby banks require their clients to provide detailed climate data and transition plans while reporting little or no information on the environmental impacts of their own digital operations.

Against this background, the study's central research question is: how can the environmental footprint of a bank's AI portfolio be defined and integrated into corporate GHG accounting and sustainability reporting? This question is explored through three sub-questions: (i) where should the accounting boundaries be drawn, i.e. which lifecycle elements and impact pathways should be included in the banking footprint (*Sections 2 and 5.1*); (ii) what data sources are actually available, and what methodological limitations do they entail (*Sections 2.3 and 2.4*); and (iii) how does the accounting approach differ depending on how access to AI is provided – through proprietary infrastructure, cloud services, or API-based model services (*Section 5.1*)?

It is important to clarify the empirical status of the study. *Section 3.3* presents an exploratory document analysis of the disclosures of nine credit institutions for the 2025 financial year, examined according to a common set of criteria. The sample is small, non-random and not representative of the Hungarian banking sector; its purpose is to determine whether current disclosure practices make the AI footprint visible *in principle*, rather than to measure its magnitude. The study does not include measurements based on banks' energy consumption or emissions data, nor has the proposed framework been empirically validated. The numerical example for a banking group with 5,000 employees presented in *Section 5.2* is purely a stylised illustration intended to demonstrate the applicability of the framework and the order of magnitude of parameter uncertainty; its values should not be interpreted either as an estimate for any specific institution or as sector-level data.

Accordingly, the study's primary contribution is conceptual and methodological and consists of two elements. First, it synthesises dispersed findings from three strands of the literature. Second – and this constitutes the backbone of the study – it develops the three-tier accounting framework set out in *Section 5*, differentiated according to the mode of access (proprietary infrastructure, cloud services, API-based model services), with an associated Scope classification, explicit system boundaries, and a range-based, scenario-based estimation procedure, together with a proposal for integrating the framework into sustainability reporting. The empirical review in *Section 3.3* does not validate this proposal; rather, it provides a rationale for its necessity.

The study is structured as follows. *Section 1.1* documents the methodology used for source collection. *Section 2* reviews the conceptual frameworks and measure-

ment methodologies for assessing the environmental footprint of AI systems, with a separate section evaluating the customer-level carbon allocation practices of cloud service providers. *Section 3* analyses international and domestic trends in the adoption of AI in banking along dimensions relevant to environmental impacts and empirically examines the visibility of the AI footprint through a review of the 2025 disclosures of nine credit institutions. *Section 4* sets out the relevant provisions of the regulatory framework, including the AI Act, the CSRD/ESRS and related data-centre regulations. *Section 5* proposes a methodological framework for estimating the banking AI footprint, illustrated by a numerical example. *Section 6* examines the question of net impact – the balance between AI applications that support sustainability and AI's own environmental footprint. *Section 7* discusses expected development trajectories and the associated management and regulatory recommendations, while *Section 8* concludes the study.

1.1 Methodology of the study and the source collection process

The study adopts the methodology of a targeted narrative literature review, aimed at developing a conceptual framework and proposing a methodology; it is not a systematic review and does not seek to provide a comprehensive, protocol-driven survey of the relevant literature. This choice of approach is justified by the interdisciplinary nature of the subject: the research question lies at the intersection of three strands of literature that have hitherto had little interaction with one another – the computer science and energy research addressing the environmental impacts of AI, the financial literature on AI adoption in banking, and the literature on ESG reporting and climate regulation – where a systematic search based on a uniform set of keywords would not achieve comprehensive coverage.

Sources were collected from the Scopus, Web of Science and Google Scholar databases, as well as from the official publication platforms of the institutions concerned (IEA, OECD, BIS, ECB, EBA, European Commission, National Bank of Hungary). Searches were conducted using English- and Hungarian-language keywords, both individually and in combination. The main search terms were: *AI energy consumption, carbon footprint of machine learning, green AI, data centre electricity demand, inference energy, embodied emissions computing, AI in banking, generative AI financial services, Scope 3 financial institutions, CSRD ESRS climate disclosure, EU AI Act energy*, and, in Hungarian, *mesterséges intelligencia bankszektor, karbonlábnyom, ESG-közzététel*. The review of Hungarian-language literature was supplemented by a thematic survey of the archives of *Hitelintézet* *Szemle, Gazdaság és Pénzügy* and *Vezetéstudomány*.

The time frame was differentiated by source category. For technological and energy-related sources, we limited the review to the period 2015–2026, as measurements of the energy requirements of deep learning from earlier periods are not relevant given current model sizes; no time limit was applied to foundational conceptual and accounting works (such as the GHG Protocol); while for sources concerning AI adoption in banking and the regulatory environment, given the rapid pace of change in these fields, we primarily considered publications from after 2020, with particular emphasis on the most recent publications available.

The selection criterion was *relevance to the intersection* of the fields: we included all sources that (i) address the measurement of the energy or emissions impacts of AI systems, (ii) address the application of AI in the financial sector, or (iii) address the regulatory and reporting interface between the two, and whose findings can be interpreted at the level of credit institution operations. We excluded works dealing solely with model architecture, algorithmic performance or the methodology for financed emissions where they had no environmental or own-operations dimension.

The study distinguishes three types of sources and assigns them different evidentiary weight. (a) *Peer-reviewed academic sources* – journal articles and conference papers – which provide the empirical and conceptual basis for the argument. (b) *Institutional and regulatory documents* (IEA, OECD, BIS, ECB, MNB, EU legal sources and standards), which serve as primary sources for describing the regulatory environment and energy-related aggregates. (c) *Industry and consultancy analyses* (Evident, McKinsey), which are not subject to academic peer review, whose data-collection methodologies are only partially publicly available, and whose full databases are accessible only by subscription; these are used solely to illustrate the scale and direction of market adoption and are not assigned independent evidentiary weight. This distinction is also reflected in the bibliography through square-bracketed labels.

The limitations of the review follow from the chosen methodology: source selection is not protocol-driven, and therefore selection bias on the part of the authors cannot be ruled out; the status of rapidly evolving regulatory documents (in particular, the 2026 Commission consultation and the related standardisation process) reflects the state of affairs prevailing at the time the manuscript was finalised; and the literature data on specific energy values are themselves subject to significant methodological uncertainty, as detailed in Section 2.3.

2. THE ENVIRONMENTAL FOOTPRINT OF AI SYSTEMS: CONCEPTUAL FRAMEWORKS AND MEASUREMENT CHALLENGES


2.1 Energy requirements of the AI lifecycle: training and inference

An analysis of the environmental impact of AI systems should appropriately begin from a lifecycle perspective. From an energy perspective, two principal operational phases can be distinguished in the lifecycle of machine learning models, and particularly large language models (LLMs): training and inference.

The energy requirements of the training phase have increased markedly with the growing size of models. Strubell et al. (2019) were among the first to draw widespread attention to the carbon cost of training natural language processing models, demonstrating that training a single large model in combination with hyperparameter optimisation could result in GHG emissions comparable to the full lifecycle emissions of several passenger cars. This phenomenon is underpinned by the exponential growth in the computational requirements of machine learning – which has doubled every few months since the 2010s (Sevilla et al., 2022) – while the marginal cost of improvements in accuracy has been rising increasingly steeply in computational and energy terms (Thompson et al., 2021). The environmental and social costs of the “bigger is always better” development paradigm were also addressed early in the critical literature (Bender et al., 2021; Crawford, 2021), giving rise to the “Greening of AI” research agenda, which seeks to elevate computational efficiency alongside accuracy as a primary criterion for scientific evaluation (Schwartz et al., 2020). Quantification has also been supported by methodological tools: calculators have been developed for estimating the emissions associated with machine learning experiments (Lacoste et al., 2019), alongside systematic reporting frameworks (Henderson et al., 2020). Later estimates suggest that training the GPT-3 model required approximately 1,300 MWh of electricity, corresponding to several hundred tonnes of CO₂ emissions depending on the carbon intensity of the electricity mix used (Patterson et al., 2021); a detailed lifecycle analysis of the open-source BLOOM model further demonstrated that, in addition to operational energy consumption, embodied emissions and idle power consumption also constitute material components (Luccioni et al., 2023). Training, however, is a one-off cost – or a recurring cost for each new model version – that is incurred primarily by foundation model providers; the vast majority of banks do not train their own foundation models, but instead fine-tune pre-trained models or access them through APIs or cloud services.

From the perspective of the banking AI footprint, the inference phase is therefore the more significant one: every query, document analysis, code-generation task or chatbot interaction consumes computational capacity and energy, and this consumption may scale faster than linearly as the number of uses increases, for example in the case of longer context windows or multi-step agentic workflows. According to measurements by Luccioni et al. (2024), individual inference operations performed by general-purpose generative models can be orders of magnitude more energy-intensive than those performed by conventional task-specific models, while research estimates indicated already at an early stage of the technology’s development that, with mass adoption, aggregate energy consumption from inference could exceed that associated with training (de Vries, 2023; Desislavov et al., 2023; Wu et al., 2022). According to the widely cited comparison by the IEA (2024), the electricity requirement per generative AI query may be nearly an order of magnitude higher than that of a conventional internet search; although the methodologies and precise numerical values underlying individual estimates are subject to debate (Dhar, 2020; Freitag et al., 2021), the overall direction – the dominance of inference in cumulative consumption – is widely accepted in the literature. The differences in specific energy consumption across task types, spanning several orders of magnitude, are illustrated in Figure 2; this also indicates the energy-related implications of banking model portfolio management (see Section 5.3).

Figure 2
Specific energy requirements of inference by task type (kwh/1,000 inferences, logarithmic scale)



Source: Author’s compilation based on the measurement results of Luccioni et al. (2024).

2.2 Beyond direct energy consumption: embodied emissions, water and resource use

The AI footprint cannot be reduced to operational electricity consumption. At least three additional impact pathways can be identified.

First, *embodied emissions*, which arise from the manufacture, transportation and end-of-life treatment of specialised hardware (graphics processing units /GPUs/, accelerators, servers, network and cooling infrastructure). Industry analyses indicate that the share of the manufacturing phase within the total lifecycle emissions of IT equipment is continuously increasing, and that the relative importance of embodied emissions in hyperscale AI infrastructure will continue to rise in parallel with the decarbonisation of electricity grids (Gupta et al., 2021; Wu et al., 2022), while accelerating hardware cycles (frequent generational replacement of AI accelerators) further increase the significance of this component, and manufacturers' lifecycle data remain only partially available (Green Software Foundation, 2026). The broader resource implications of AI's physical infrastructure, extending from rare-earth extraction to electronic waste, are documented in detail in the critical literature (Crawford, 2021).

Second, *water consumption*. Depending on the cooling technology and geographical location, cooling data centres may involve significant direct water consumption, supplemented by the indirect water requirements of electricity generation (Mytton, 2021; Siddik et al., 2021). According to estimates by Li et al. (2023), the training and operation of large language models could generate a global water footprint on the order of several billion cubic metres by the middle of the decade; in the case of capacity located in water-stressed regions, this may also become a source of local sustainability and social conflict.

Third, *network and system-level impacts*. The geographical concentration of data centres has become a structural factor affecting electricity-grid demand in certain regions; in Ireland, for example, data centres account for a share of national electricity consumption reaching as much as one fifth, while in the United States, data-centre demand has become an independent planning factor in federal energy policy (Shehabi et al., 2024). At the system level, according to the IEA's projections, a significant share of the additional demand from data centres over the period to 2030 will still be met by fossil-fuel sources, meaning that AI expansion – without being balanced by clean energy – may make it more difficult to achieve climate targets (IEA, 2025).

2.3 Measurement methodologies and accounting challenges

There is currently no uniformly accepted mandatory methodology for measuring the AI footprint at the corporate level; several partially overlapping frameworks are competing with one another, while measurement and reporting gaps are also identified by the international policy literature as a major constraint in the field (OECD, 2022).

The general framework for corporate carbon accounting is provided by the GHG Protocol, which classifies emissions into direct emissions (Scope 1), emissions associated with purchased energy (Scope 2), and other indirect emissions (Scope 3) (WRI–WBCSD, 2004). The banking AI footprint can in principle be situated within this framework, but practical allocation is not straightforward: AI workloads operated on an institution's own server infrastructure constitute a Scope 2 item, whereas AI services accessed through the cloud or via APIs fall within Scope 3 Category 1 (purchased goods and services) – precisely the category that service-intensive companies typically cover using the most rudimentary, cost-based estimates. AI adoption is therefore partly “invisible” from the perspective of banks' emissions inventories: under current disclosure practices, increasing consumption of cloud-based AI typically does not appear as a separate item based on quantitative data.

Pointing towards software-specific measurement is the Software Carbon Intensity (SCI) methodology developed by the Green Software Foundation and based on the ISO/IEC 21031:2024 standard. Its AI extension (SCI for AI) converts energy consumption into emissions using location-specific grid carbon-intensity factors, takes into account embodied emissions over the hardware lifecycle, and defines standardised functional units to facilitate the comparability of AI systems (Green Software Foundation, 2026). On the research side, systematic reviews of the “Greening of AI” literature document a rapidly expanding yet fragmented ecosystem of measurement tools, efficiency techniques and reporting practices (Verdecchia et al., 2023), while consistent energy reporting has been proposed for institutionalisation as a quality standard in machine learning research (Henderson et al., 2020). Methodological adaptation of lifecycle assessment (LCA) to AI solutions has likewise been undertaken, specifically to uncover the hidden impacts of AI applications designed for environmental purposes (Ligozat et al., 2022).

Among the practical obstacles to measurement, *transparency asymmetry* is particularly significant: the vast majority of banks' AI consumption is realised through a small number of major cloud and model service providers, which provide customer-level energy or emissions data only to a limited extent, using different methodologies (market-based versus location-based Scope 2) and at different levels of granularity. As a result, banks as AI deployers typically have to rely on

estimates whose input parameters (specific energy values based on tokens, queries or GPU-hours) are themselves subject to substantial uncertainty (Luccioni et al., 2024; OECD, 2022).

2.4 Cloud service provider carbon allocation: theoretically measurable and practically available data

Since the vast majority of the banking AI footprint is realised through cloud or model services, the practical question of measurability ultimately comes down to this: what data does the service provider make available to the customer, and according to what methodology are those data generated? A comparison of the customer-level carbon allocation tools of the three major cloud service providers shows that there is a significant gap between theoretical measurability and practical data accessibility, and that this gap is not neutral from the perspective of banks.

The fundamental allocation problem. A cloud service provider's own emissions are aggregated: they consist of energy consumption measured at datacentre, cluster or regional level and purchased electricity. These emissions must be allocated to individual customers according to some allocation key, typically in proportion to the resource hours used. Allocation decisions, however, are not neutral: AWS's Version 2.0 methodology, for example, also allocates the estimated emissions associated with *unused capacity* proportionally among all customers (AWS, 2025) – an approach that is methodologically defensible, but means that a customer's footprint partly depends on factors over which it has no influence.

Market-based accounting as a key issue. The most significant methodological issue is the choice between the market-based and location-based approaches to accounting for Scope 2 emissions. The market-based approach takes the provider's renewable energy purchases (PPAs, guarantees of origin) into account, meaning that in regions covered by renewable energy contracts, the Scope 2 emissions allocated to a customer may even fall to zero – while the workload physically consumes electricity from the same grid, with its carbon intensity unchanged. From a bank's perspective, this means that its reported AI footprint can be reduced drastically simply by changing service provider or region, without any reduction in actual physical emissions. This is precisely why the CSRD requires disclosure of Scope 2 emissions using both the market-based and location-based methods; according to the proposal advanced in this study, both values should likewise be reported when estimating the banking AI footprint, with the difference between them treated as separate information requiring interpretation.

Asymmetries in data availability. The practices of the three providers differ materially. Google Cloud Carbon Footprint reports both Scope 2 approaches in parallel, uses hourly grid emissions factors and machine-level consumption data for location-based calculations, and provides breakdowns by project, product and region, as well as database export functionality – while the provider itself states that, unlike the company’s audited, top-down footprint, these customer-level reports have not undergone third-party verification (Google Cloud, 2026). For a long time, AWS’s tool reported only market-based Scope 1–2 data and did not include embodied or other upstream emissions; the Version 3.0 methodology, which was third-party-verified, extended the reporting in 2025 to material Scope 3 categories – upstream emissions associated with energy procurement, cradle-to-gate embodied emissions of IT hardware, and emissions associated with buildings and equipment – with embodied items amortised over the useful life of the assets (six years for IT hardware and fifty years for buildings) (AWS, 2025). There is, however, a material time lag in the data: Scope 2 data are updated monthly, with a delay of approximately three months. Microsoft’s Emissions Impact Dashboard covers all three Scopes and both Scope 2 approaches, but access is dependent on the type of contract and requires a separate licence, which may constitute a practical constraint for smaller institutions.


Implications for banking practice. Three conclusions follow from this. First, customer-level data obtained from different service providers cannot be aggregated without methodological harmonisation, as they are based on different allocation keys, different Scope 2 definitions, different scopes of coverage and different verification statuses; accordingly, the data source and methodological choice must be documented in detail in the bank’s emissions inventory. Second, the data are historical and subject to delays and are therefore unsuitable on their own for operational decision-making – for example, selecting the region in which a model is run; the carbon-aware scheduling discussed in Section 7.2, together with real-time data on grid carbon intensity, provides a basis for such decisions. Third (and this is the most important limitation), these provider tools are designed for infrastructure-level (IaaS/PaaS) usage and currently do not extend to model services accessed via APIs, identified as the third mode of access in Section 5.1. Yet this is precisely the layer in which banks’ generative AI use is concentrated. The gap between theoretically measurable and practically available data is therefore greatest precisely where the banking AI footprint is growing most rapidly – this is the gap that Section 7.4 proposes to address through contractual data-provision clauses and collective action at sector level.

3 AI ADOPTION IN THE BANKING SECTOR: THE DEMAND SIDE OF THE FOOTPRINT

3.1 International trends

The banking sector is one of the leading adopters of generative AI in corporate applications. International institutions had already analysed the transformative impact of AI on the financial system before the current wave (FSB, 2017; OECD, 2021), while the European Central Bank identified the rapid adoption of generative AI by banks as a development relevant to financial stability as well (ECB, 2024). Systematic reviews of the academic literature document rapidly expanding fields of research on commercial banking applications (Königstorfer–Thalmann, 2020), customer relationship solutions (Hentzen et al., 2022) and financial AI techniques (Cao, 2022; Milana–Ashta, 2021), while the intersection of banking operations research and AI has been identified as a distinct discipline (Doumpos et al., 2023). From the perspective of central-bank research, Aldasoro et al. (2024) argue that AI is transforming the entire value chain of financial intermediation – from lending and payment systems to regulatory compliance. The 2025 edition of the Evident AI Index, which assesses the AI maturity of the world's fifty largest banks across four pillars (talent, innovation, leadership and transparency), documents accelerating adoption across the sector (Evident, 2025). It should be noted that the Evident and McKinsey materials are industry and consultancy analyses, respectively, and are not subject to academic peer review; this study uses them to illustrate the scale of market adoption rather than as methodological evidence, and their detailed databases are available only by subscription. *Three developments are particularly noteworthy from the perspective of the environmental footprint.*

First, *the growing adoption of agentic AI*. Nearly one-third of the banking AI applications publicly announced in the first quarter of 2026 were agentic, compared with 15% in the previous quarter (Evident, 2026). From an energy perspective, agentic systems represent a qualitative shift: a single user request may trigger dozens of model calls, tool invocations, and iterations, causing energy consumption per task to increase several-fold compared with conventional single-turn queries. When forecasting the banking AI footprint, the number of use cases alone is therefore a potentially misleading indicator; weighting by compute intensity is essential. The development of the share of agentic applications is shown in *Figure 3*.

Figure 3**Share of agentic applications among newly announced banking AI use cases**

Source: Author's compilation based on Evident (2026).

Second, *the expanding user base*. Generative AI has evolved from early, data-scientist-centric applications towards general-purpose assistants (copilots) available to the entire workforce and, according to the latest trends, from enterprise-wide copilots towards function-specific solutions embedded in workflows (Evident, 2026). At a bank employing tens of thousands of people, aggregate inference consumption resulting from widespread daily AI use may already become comparable in order of magnitude to the energy consumption of conventional office IT infrastructure.

Third, *the growing importance of risk management and compliance applications*. Banks make extensive use of AI for regulatory compliance, combating financial crime and fraud detection – the cost-efficiency benefits of which are also supported by domestic empirical literature (Benedek–Nagy, 2023) – as well as for credit risk analysis and, importantly, for ESG data processing and the estimation of Scope 3 emissions (BIS, 2024; McKinsey, 2024). The latter creates a paradoxical situation: the automation of sustainability reporting itself relies on energy-intensive AI processes, as researchers at the Bank for International Settlements (BIS) Innovation Hub also noted in connection with the Symbiosis project, observing that any application of AI entails material emissions while the decarbonisation of electricity grids is progressing only slowly (BIS, 2024). The concept of “AI washing”, which has become established in Hungarian academic literature

as well, treats the credibility of communications concerning value creation from banks' AI investments as a risk in its own right (Prisznyák, 2023); according to the argument advanced in this study, this phenomenon also has an environmental dimension where banks communicate the sustainability benefits of AI without accounting for their own AI footprint.

3.2 The domestic and regional context

AI adoption in the Central and Eastern European banking sector, including Hungary, is taking place with a time lag relative to the international frontrunners, but with a similar structure: customer-service automation, fraud detection, document processing and internal efficiency applications (development support and reporting) predominate (Bagó, 2023a; MNB, 2025). In addition to the opportunities presented by AI applications in finance (Bagó, 2023a), the domestic literature has addressed cybersecurity implications (Bagó, 2023b) and the transformation of the compliance function (Harkácsi-Szegfű, 2021); the systemic-risk implications of dependence on technology giants have also been examined in the domestic literature in the context of the "Too Big(Tech) to Fail" problem (Bódi et al., 2023). The environmental footprint dimension, however, is at most addressed only tangentially in these works – the present study seeks to narrow this gap, including through the empirical review in Section 3.3.

Three domestic characteristics are particularly noteworthy from the perspective of the footprint. First, the AI consumption of Hungarian credit institutions relies almost exclusively on *imported* computing capacity: inference of large language models typically takes place in data centres located abroad – in Western Europe or the United States – so the associated emissions appear in the inventories of Hungarian subsidiaries as a Scope 3 item, reflecting the carbon intensity of the energy mix in the host countries, if they appear at all. This dependence also represents the environmental dimension of the big-tech concentration risk analysed by Bódi et al. (2023): the measurability of the banking AI footprint depends on the willingness of a small number of global providers to supply data. Second, the carbon intensity of regional electricity mixes varies considerably, giving data-residency and localisation decisions – for example, choosing processing in low-carbon-intensity regions within the EU – measurable emissions implications. Third, the ESG frameworks of the parent companies of Hungarian banking groups – in line with the expectations of the green finance programme of the National Bank of Hungary (MNB, 2024) – increasingly also include group-level digital carbon-footprint targets, which may influence AI procurement decisions at subsidiary banks.

3.3 Empirical review of sustainability disclosures by Hungarian banks

Limitations of the sample and methodology

The assertion made in *Section 3.2* – that energy consumption attributable to AI remains invisible in reporting – requires empirical verification. To this end, we conducted a documentary analysis of publicly available documents relating to the 2025 financial year from nine credit institutions. Three limitations of the sample should be noted at the outset, as these are themselves part of the findings of the analysis.

First, not all credit institutions publish sustainability reports. Under the applicable regulatory framework, institutions operating in the form of branches are exempt from the obligation to make entity-level ESG disclosures – the Hungarian branch of ING explicitly states this, noting that ING Group applies the ESRS on a voluntary basis (ING Bank N.V., Hungarian Branch, 2026) – while smaller institutions become subject to the obligation only at a later stage: according to MagNet Bank’s own disclosure, as a small and non-complex institution, it will first be required to publish an integrated, group-level sustainability report for the 2028 financial year (MagNet Bank Zrt., 2026).

Second, even where a report does exist, it is not necessarily a standalone document. In the sample, the typical format is a sustainability statement incorporated into the annual report or consolidated management report (MBH Bank Nyrt., 2026; GRÁNIT Bank Nyrt., 2026); the only standalone sustainability report is that of CIB Bank, which itself states that it does not constitute an ESG report under the CSRD, as this is prepared at group level by its parent company, Intesa Sanpaolo, while the local document is voluntary and serves a transparency purpose for Hungarian stakeholders (CIB Bank Zrt., 2026).

Third, not every institution has a downloadable document for 2025. Group-level reports are often available only as part of the parent company’s disclosures; downloading and processing annual reports running to several hundred pages presents technical constraints; and for some institutions, only documents of a different type were available, such as an investor presentation (Oberbank AG, 2026) or a sustainability financing framework (Deutsche Bank, 2026). The sample is therefore not representative of the Hungarian banking sector; its purpose is illustrative and exploratory: to examine whether the AI footprint can, *in principle*, be visible in current disclosure practices.

We analysed the documents according to five uniform criteria: (i) whether digital, IT or data-centre energy consumption or emissions are separately identified; (ii) which Scope 3 categories are disclosed and according to what methodology; (iii) whether data from cloud service providers are reported; (iv) whether AI is men-

tioned in connection with energy requirements; and (v) whether there are targets relating to the digital footprint.

Results

Table 1

Analysis of the sustainability disclosures of the credit institutions examined (financial year 2025)

Institution	Document type	Separate digital/ IT energy	Scope 3 disclosure	AI's energy requirements
MBH Bank Nyrt.	consolidated sustainability report embedded in the annual report	none	Category 15 only (PCAF)	none
GRÁNIT Bank Nyrt.	sustainability statement embedded in the annual report	none (E1-5 'not relevant')	not disclosed (phased introduction)	none
CIB Bank Zrt.	independent, voluntary sustainability report	none	not disclosed	none
UniCredit Bank Hungary Zrt.	annual report	none	not included in the document	none
MagNet Bank Zrt.	disclosure document as per CRR	none	voluntary, unaudited (PCAF, from 2025)	none
ING Bank N.V. HU Branch	IFRS standalone financial statements	exempt from disclosure	–	none
Bank of China HU Branch	IFRS separate financial statements	no ESG disclosures	–	none
Oberbank AG	investor presentation	none	–	none
Deutsche Bank	sustainable financing framework	asset-side PUE threshold	–	none

Source: author's compilation based on the documents cited.

The first and most important finding of the review is that none of the documents in the sample separately identifies digital, IT or data-centre energy consumption, or the emissions attributable to it. The absence of such information does not result from a scarcity of data: the energy data disclosed by the institutions providing more detailed reporting are in fact very comprehensive but are broken down along different dimensions.

This is best illustrated by the most comprehensive publication. The MBH Group's total energy consumption in 2025 was 76,467.19 MWh (2024: 82,054.23 MWh), which the report breaks down by fossil, nuclear and renewable sources, but not by purpose of use. The text itself states that most energy consumption was attributable to the utility consumption of buildings, followed by fuel consumption by the vehicle fleet; IT does not even appear as a consumption category. The Group also identifies group-wide energy data consolidation as a challenge in its own right, promising a unified IT-supported system by 2026 (MBH Bank Nyrt., 2026). A similar pattern can be observed at CIB Bank: electricity consumption in 2025 amounted to 4,921,960 kWh, while the measures resulting in savings of 383,836 kWh are, without exception, building management-related (heat pumps, LED lighting, a building management system and window-film installation); but no IT-related item is reported (CIB Bank Zrt., 2026).

The second finding concerns the structure of Scope 3 and confirms the allocation analysis presented in *Section 5.1*: where Scope 3 is disclosed at all, it relates exclusively to Category 15. The MBH Scope 3 table contains a single item, investments (4,615,674.55 tCO₂e); Category 1 (purchased goods and services), which would include cloud-based AI consumption, is not reported. The consequence in terms of scale is that, within the Group's total footprint of 4,630,535.72 tCO₂e, its own operations (Scopes 1–2) account for only 0.32%, meaning that precisely the category in which AI consumption would be expected to increase remains methodologically undeveloped. GRÁNIT Bank, making use of the option for the phased introduction of the ESRs, does not disclose any Scope 3 emissions at all, deferring this to its 2026 report, while it classifies the E1-5 energy consumption disclosure as “not material” (GRÁNIT Bank Nyrt., 2026). MagNet Bank has applied the PCAF methodology since 2025, but explicitly presents the resulting emissions data as voluntary and unaudited (MagNet Bank Zrt., 2026).

The third – and, from the perspective of the study's argument, most revealing – finding is that AI adoption is documented in detail in these same documents, while its energy requirements are nowhere reported. The sharpest contrast is seen at GRÁNIT Bank: it was the first bank in Hungary to introduce AI-based selfie account opening; in 2024, it launched its generative-AI-powered mortgage assistant; it began developing generative-AI-based solutions for its lending processes; it uses AI solutions to process Otthon Start loans; and in 2025 it received the

“AI-Based Innovation of the Year” award – while the same institution discloses neither energy consumption nor Scope 3 data (GRÁNIT Bank Nyrt., 2026). CIB Bank reports the participation of 492 employees in its basic AI training and 196 in its advanced training (CIB Bank Zrt., 2026), while UniCredit introduced AI-based voicebots and digital assistants in customer service as part of the Genesys project (UniCredit Bank Hungary Zrt., 2026); in both cases, these are reported in the digitalisation and human-resources sections, respectively, separately from the environmental section. In Oberbank’s investor presentation, this phenomenon is also evident visually: the slide presenting AI as an efficiency tool is followed by the sustainability slide, with no connection between their content – AI forms part of the efficiency narrative, whilst sustainability forms part of the portfolio narrative (Oberbank AG, 2026).

Finally, the fourth finding is a documented instance of the dual position discussed in *Section 4.3*. Deutsche Bank’s sustainable finance framework lists the construction, modernisation and operation of data centres as financeable sustainable activities in the ICT section, subject to compliance with a specified PUE threshold (Deutsche Bank, 2026). The energy-efficiency metric therefore does appear, but on the *asset side*, as a customer-assessment criterion rather than as a measure of the bank’s own operations. In its concluding statement, the same framework acknowledges that there is no uniform market standard, that the availability of reliable data is limited, and that different but equally acceptable measurement techniques may produce substantially different results; this constitutes a corporate acknowledgement of the measurement uncertainty discussed in *Sections 2.3 and 2.4*.

Overall, the review provides empirical support for the study’s central claim: several of the institutions examined are demonstrably intensive users of AI, yet none identifies AI-related energy consumption or emissions in its sustainability disclosures. The absence itself is the finding. This is not a matter of omission – the current ESRS datapoints and established banking practice do not require such information – but rather a structural blind spot whose significance is increasing in parallel with the pace of adoption. The framework proposed in *Section 5* is intended precisely to provide a procedure for eliminating this blind spot.

4 REGULATORY ENVIRONMENT: TRANSPARENCY REQUIREMENTS FROM MULTIPLE DIRECTIONS

4.1 Energy-related provisions of the EU's AI Act

The AI Act – the world's first comprehensive, horizontal AI regulatory framework – contains several provisions relating to energy in the context of the twin (digital and green) transitions, but these primarily impose obligations on model providers. Providers of general-purpose AI (GPAI) models are required, in the technical documentation referred to in Annex XI to the Regulation, to record the model's known or estimated energy consumption; where actual consumption is not known, the estimate may be based on the computational resources used (European Parliament and Council, 2024). Article 40 of the Regulation empowers the Commission to initiate the development of harmonised standards to improve the resource efficiency – including energy efficiency – of AI systems, while Article 95 provides for the development of voluntary codes of conduct covering environmental sustainability and energy-efficient programming.

The literature nevertheless points to significant shortcomings in the Regulation's environmental provisions: disclosure is made to authorities rather than publicly; the Regulation measures energy consumption rather than actual emissions and therefore disregards potentially order-of-magnitude differences in the carbon intensity of electricity mixes; embodied emissions remain outside its scope; and, in the absence of a uniform measurement methodology, the data are not comparable (Green Software Foundation, 2026). Of particular relevance from the perspective of the banking AI footprint is the fact that the GPAI Code of Practice – which operationalises the requirements of Annex XI in the form of a model documentation form and covers consumption in both the training and inference phases – allows for an exception: a provider may be exempted from reporting energy data if it does not have the necessary information from the partner providing the computing capacity or hardware (Green Software Foundation, 2026). This exception effectively institutionalises precisely the transparency asymmetry identified in Section 2.3 as a major obstacle to estimating the banking footprint: the lack of data can be passed along the supply chain, leaving the bank at the end of the chain, as the user, reliant on estimates. In the spring of 2026, the Commission launched a targeted consultation on a measurement framework for the energy consumption and emissions of AI models and systems, which also serves to prepare a possible EU AI energy and emissions label (European Commission, 2026). For banks, this means that standardised energy data may become available on the *supplier side* over the medium term, potentially improving the quality of their own footprint

estimates significantly; obligations on the deployer side, however, remain indirect for the time being.

4.2 CSRD, ESRS and double materiality

The Corporate Sustainability Reporting Directive (EU) 2022/2464 (CSRD) and the European Sustainability Reporting Standards (ESRS) require large credit institutions, in accordance with the principle of double materiality, to disclose material environmental impacts, risks and opportunities, including energy consumption and the breakdown of GHG emissions by Scope 1–3 (European Parliament and Council, 2022). From the perspective of the AI footprint, the key issue is *materiality*: where a bank's AI-related energy consumption and emissions from purchased services account for a material and growing proportion of its own operational footprint, this is difficult to disregard in the double materiality assessment from an impact-materiality perspective. In current disclosure practice, however, digital consumption, including AI-specific consumption, is generally not reported as a separate item – which, consistently with earlier content-analysis findings on disclosures by major Hungarian banks (Vőneki–Lamanda, 2020), indicates that reporting practice follows the regulatory minimum requirements and incorporates emerging areas of impact only with a time lag. The framework proposed in *Section 5* of the study provides a practical solution for making this visible.

4.3 Data centre and taxonomy links

The regulatory environment surrounding the banking AI footprint is shaped, from a third direction, by regulation governing data-centre energy efficiency. Following the revision of the EU Energy Efficiency Directive, the reporting of key energy-performance indicators for larger data centres – such as the PUE metric measuring power usage effectiveness – has become mandatory, while some Member States, Germany for example, have also introduced requirements concerning the share of renewable energy and the utilisation of waste heat through their national energy-efficiency legislation. The EU Taxonomy Climate Delegated Regulation also establishes the framework for assessing the taxonomy compliance of data-centre activities. Banks occupy a dual position: they are simultaneously *consumers* of data-centre services (through their own AI footprint) and *financiers* of rapidly expanding data-centre investments – the latter being reflected in financed emissions and warranting separate analysis.

5 ACCOUNTING FOR THE BANKING AI FOOTPRINT: A PROPOSED METHODOLOGICAL FRAMEWORK

5.1 Accounting boundaries and allocation

The proposed framework takes as its starting point a three-level typology of the banking AI portfolio according to the mode of access, as this determines where the emissions are accounted for:

1. *AI operated on the bank's own infrastructure (on-premise)*: consumption can be measured (through server-level power monitoring and GPU-hour tracking); emissions are a Scope 2 item, calculated using the carbon intensity of the bank's own electricity contracts (location-based and market-based).


2. *Cloud-hosted, self-managed workloads (IaaS/PaaS)*: consumption can be estimated from customer-level reports provided by the cloud service provider or from resource-usage data (vCPU-hours, GPU-hours); emissions are classified as Scope 3, Category 1, and estimated using the provider's region-specific carbon intensity and PUE.

3.1 *Model serving via API (SaaS/MaaS)*: the most difficult level to measure; estimates may be based on energy-intensity factors per token or per query, which should be derived from the literature (de Vries, 2023; Luccioni et al., 2024) or, as they become available, from service-provider disclosures. Scope 3, Category 1.

The accounting classification of embodied emissions varies depending on the mode of access, and failure to classify them correctly can lead to systematic underestimation of the inventory. In the case of own infrastructure, the bank owns the hardware; accordingly, emissions arising from the manufacture and transportation of servers, GPUs, and associated networking and cooling equipment fall under Scope 3 Category 2 (capital goods) and are accounted for in full in the year of acquisition, while end-of-life treatment falls under Category 5. This is therefore not merely a qualitative consideration but an emission source that should be quantified and included in the inventory if it is considered material. In the case of cloud- and API-based use, by contrast, the hardware is owned by the service provider, meaning that the embodied emissions arise within the provider's capital goods and, if at all, are passed back to the bank through the provider's allocation under Scope 3 Category 1; the provider methodologies discussed in Section 2.4 therefore differ materially in whether they include the embodied component and what amortisation assumptions they apply. At all three levels, it is also appropriate – depending on materiality – to track water consumption as an indicator (Li et al., 2023; Mytton, 2021), and to explicitly document data gaps relating to embodied emissions (Luccioni et al., 2023; Wu et al., 2022). The proposed accounting framework is summarised in *Figure 4*.

Figure 4

Three-tier accounting framework for the banking AI footprint: access modes, measurement approaches and GHG accounting categories



Source: author's compilation

5.2 Illustrative estimation model

The applicability of the framework is illustrated by a stylised numerical example. Consider a medium-sized banking group employing 5,000 white-collar employees that introduces a general-purpose AI assistant. The parameters required for the estimate are the number of queries per employee per day (q), the average electricity requirement per query for the full-service chain, including data-centre auxiliary consumption (PUE) (e), the number of working days (d), and the carbon intensity of the relevant electricity mix (c). Annual AI inference emissions are then:

$$E = (Nqdec)/10^6, \quad (1)$$

where E is annual emissions in 't' CO₂e, ' N ' is the number of users (persons), ' q ' is the number of daily queries (queries/person/day), ' d ' is the number of working days (days), ' e ' is the specific energy requirement (Wh/query), and ' c ' is the grid carbon intensity (kg CO₂e/kWh). The divisor of 10^6 combines two unit conversions: Wh→kWh (10^3) and kg→tonnes (10^3). The interpretation of the ' e ' parameter requires an independent methodological decision, as the choice of system boundary can affect the result by orders of magnitude. In this study, ' e ' includes: (i) the actual electricity consumption of the server and accelerator serving the model during query processing; (ii) data-centre auxiliary consumption through the

PUE multiplier (cooling and power-supply losses); and (iii) the allocated share of idle and standby consumption, as well as consumption arising from redundancy and reserved but unused capacity. It does not, however, include (iv) the energy requirement of network data traffic between the endpoint and the data centre, (v) consumption by the user's end device, or (vi) broader elements of the provider's infrastructure that cannot be directly attributed to the workload. It should be emphasised that specific values taken from the literature have not been derived using uniform system boundaries: some measurements record only accelerator consumption, others that of the entire server, while still others include PUE, and the treatment of idle consumption also varies. For example, the lifecycle assessment of the BLOOM model demonstrated that idle consumption alone can be material in magnitude (Luccioni et al., 2023). When documenting an estimate, therefore, the system boundary of the ' e ' value used must be specified in detail, and specific values from different sources cannot be combined without harmonisation. The parameters used here are: $q=30$ queries/day, $d=230$ working days, $N=5,000$ persons, corresponding to approximately 34.5 million queries annually; and $c=0.2-0.4$ kg CO₂e/kWh, depending on the electricity mix of the serving region.

The choice of the specific e value, by contrast, opens up an uncertainty range spanning three orders of magnitude and therefore cannot be treated as a single "average" value. Two distinctly different measurement families can be identified in the literature. One – the literature based on controlled measurements of open models – gives a specific value of approximately 0.05 Wh/query for text generation (Luccioni et al., 2024; see Figure 2). The other – the literature using top-down estimates for large, mass-deployed general-purpose assistants – uses values two orders of magnitude higher, in the range of several watt-hours (IEA, 2024). This difference is not a measurement error, but the combined result of three factors: *model size* (the measured open models are substantially smaller than the models underlying market-leading assistants), *system boundary* (the differing treatment of idle capacity, redundancy and PUE components discussed above), and *serving efficiency* (batching, caching and quantisation).

The present study therefore does not use a single range, but four explicitly derived scenarios:

1. *Small, task-specific model*: – $e \approx 0.05$ Wh: classification, summarisation and short-text generation using a distilled model, consistent with controlled measurements (Luccioni et al., 2024).
2. *Medium-sized model*: – $e \approx 0.5$ Wh: a larger context window, longer response, medium-sized model; one order of magnitude above the measured baseline.
- 3.1 *Large, general-purpose assistant*: – $e \approx 3$ Wh: consistent with estimates reported for mass-deployed assistants (IEA, 2024).

4. *Agentic workflow*: – $e \approx 30$ Wh: here, a single user query does not correspond to a single model call. On a conservative estimate, an agentic workflow typically involves 10–20 model calls (planning, tool use, and iterative verification), and because each step processes the growing context anew, the number of tokens processed increases at each step. Taking the large-model call in Scenario 3 as the basis, 10 calls $\times$ an average context size equal to the initial context size ≈ 30 Wh; this value is therefore not an assumption, but an estimate derived from Scenario 3 based on the number of calls and steps.

The parameters and results of the four scenarios are summarised in *Table 2*.

Table 2

Four scenarios of the stylised numerical example (N=5,000 persons; q=30 queries/day; d=230 days; total of 34.5 million queries/year)

Scenario	Typical banking use	e (Wh/query)	Annual consumption (MWh)	Annual emissions (t CO ₂ e, where c = 0.2–0.4)
1. Small, task-specific model	classification, summarisation, and short template-based generation	0.05	1.7	0.3–0.7
2. Medium-sized model	standard text generation, longer context	0.5	17.3	3.5–6.9
3.1 Large, general-purpose assistant	corporate co-pilot, open-ended questions	3	103.5	20.7–41.4
4. Agentic workflow	10–20 model calls, iterative use of tools	30	1 035.0	207.0–414.0

Source: Author's own calculations. The system boundary in each scenario comprises server and accelerator consumption + PUE + allocated idle and redundancy components.


Under these scenarios, annual energy consumption ranges from 1.7 MWh to 1,035 MWh, while annual emissions range from 0.3 t CO₂e to 414 t CO₂e (*Figure 5*). The three-order-of-magnitude variation is itself one of the study's principal findings: the magnitude of the banking AI footprint is determined not by the number of users, but by the choice of model and process architecture. The same institution with 5,000 employees, at the same intensity of use, can produce footprints differing by orders of magnitude depending on its model-sizing and agentification

decisions; this is what gives practical significance to the efficiency lever discussed in Section 5.3.

The example is nevertheless incomplete for three reasons. First, it covers only the use of a general-purpose assistant and does not include document-processing or code-generation applications. Second, it does not include embodied emissions or fine-tuning components. Third, the number of queries itself increases dynamically. The estimate should therefore always be reported as a range, specifying the scenario and system boundary and accompanied by a sensitivity analysis (Henderson et al., 2020; OECD, 2022).

Figure 5

Sensitivity analysis of the stylised numerical example across four scenarios: annual emissions (t CO₂e, logarithmic scale) as a function of specific energy requirement (e) and grid carbon intensity (c) (N = 5,000 persons; q = 30 queries/day; d = 230 days)



Source: Author's own calculations and compilation; the sources and derivation of the 'e' values are provided in Section 5.2.

The essence of the methodological proposal therefore lies not in a point estimate, but in the institutionalisation of the *process*: (i) establishing an energy-focused inventory of the AI application portfolio; (ii) using actual data for measurable (on-premise, IaaS) items and range-based estimates for non-measurable (API-based) items; (iii) updating the estimate annually and incorporating it into sustainability reporting – at least at the level of methodological notes; and (iv) applying energy-related criteria in AI procurement and architecture decisions (model-size pro-

portionality, region selection, and contractual provisions requiring energy-data provision by service providers).

5.3 Model sizing as an efficiency lever

From the perspective of banking practice, the greatest short-term potential for reducing emissions lies in *matching tasks and models*. A significant proportion of use cases – classification, summarisation and template-based text generation – can be served at high quality by smaller, task-specific or distilled models whose specific energy requirements are a fraction of those of the largest general-purpose models (Luccioni et al., 2024; Schwartz et al., 2020). Instead of a “use the largest model for everything” default, model portfolio management – which incorporates the energy dimension alongside risk, data-protection and cost considerations – is a natural extension of banking AI governance and is consistent with the application of efficiency techniques documented in the Green AI literature, including quantisation, distillation, caching and batch processing (Verdecchia et al., 2023). The elimination of redundant calls and optimisation of the number of steps in agentic workflows also fall within this approach.

6 NET IMPACT: GREENING BY AI AND GREEN AI

The environmental balance of AI applications in banking is not one-sided. The research community has systematically mapped the potential applications of machine learning for climate action, from energy-system optimisation to climate modelling (Rolnick et al., 2022; Vinuesa et al., 2020), and promising applications can also be identified in the financial sector: the literature and central-bank practice alike indicate progress in the automation of Scope 3 data collection, climate-risk analysis, energy-efficiency optimisation and the detection of greenwashing (BIS, 2024; Nishant et al., 2020). A review of fintech and blockchain-based green financial solutions in Hungary also highlights the sustainability potential of the technology (Horváth, 2022).

The “greening by AI” narrative, however, cannot methodologically be accounted for at the expense of the “greening of AI” requirement. This pair of concepts – AI as a tool for sustainability and AI as a sustainability problem – is also a well-established analytical framework in the international ethics literature (Cowls et al., 2023; Taddeo et al., 2021; van Wynsberghe, 2021), and the methodological conditions for climate policy assessment were systematised by Kaack et al. (2022): the system-level climate impact of AI can only be assessed by taking into account the combined effects of direct (computational), indirect (application)

and structural (behavioural, market) impacts. At the level of corporate accounting, the two effects fall into different categories: emissions reductions achieved through AI applications generally constitute avoided emissions, which cannot be netted against actual Scope 2–3 items in the corporate inventory. In the communications of credit institutions, therefore, the two dimensions should be kept strictly separate; failure to do so carries a greenwashing risk, particularly in the context of increasingly stringent supervisory scrutiny of sustainability claims and the phenomenon of AI washing (Prisznayák, 2023). Critical literature further cautions that the “green” narrative surrounding AI often obscures the technology’s own environmental costs and the political economy of its supply chains (Brevini, 2020; Dauvergne, 2022), while it remains an open question whether AI will ultimately accelerate or slow the net-zero transition (Luers et al., 2024). The rebound effect provides a further reason for caution: according to the logic of the Jevons paradox, efficiency gains may increase total consumption by driving growth in the volume of use (Alcott, 2005; Sorrell, 2009), a phenomenon also supported by experience from earlier efficiency cycles in the ICT sector (Freitag et al., 2021).

7 WHERE ARE WE HEADING? EXPECTED TRAJECTORIES, CONCLUSIONS AND RECOMMENDATIONS

7.1 The demand trajectory: structural growth

Available forecasts consistently point to rapid, structural growth in the banking AI footprint. On the demand side, the projected near-doubling of data-centre electricity consumption by 2030 (IEA, 2025) and the data-centre share of national electricity consumption in the United States potentially reaching 12% (Shehabi et al., 2024) define the macro trajectory, within which the centre of gravity of consumption is expected to shift permanently from training towards inference (de Vries, 2023; Desislavov et al., 2023). In the banking sector, this trajectory is further steepened by the sharp increase in the share of agentic applications – from 15% to nearly one-third within a single quarter (Evident, 2026) – as the computational requirements of agentic workflows per task unit are several times those of conventional queries. Against this background, a plausible scenario is that credit institutions’ AI-related Scope 2–3 items will expand at an annual rate of several tens of per cent over the remainder of the decade, while other components of their operational footprint (building energy, travel and paper use) stagnate or decline – in other words, the share of AI within operational emissions will continue to increase even if its absolute magnitude remains moderate for the time being.

7.2 Technological countervailing forces: efficiency and carbon-aware operation

The growth trajectory may be moderated by technological countervailing forces. Improvements in hardware and algorithmic efficiency have historically been substantial (Kooimey et al., 2011), and the Green AI research agenda – from model-size reduction through quantisation to distillation – is continuing to expand the efficiency toolkit at a rapid pace (Schwartz et al., 2020; Verdecchia et al., 2023). A promising direction is *carbon-aware computing*: rescheduling workloads in time and space in line with the real-time carbon intensity of the electricity grid, an approach already being deployed operationally by hyperscale providers (Radovanović et al., 2023), while at the level of data-centre design, the joint optimisation of embodied and operational emissions has also emerged as an area of focus (Acun et al., 2023). From the perspective of the banking AI footprint, this means that provider selection and contractual terms (requiring carbon-aware scheduling and favouring regions with low carbon intensity) may become genuine, measurable differentiating factors in the coming years. Efficiency gains alone, however, do not guarantee an absolute reduction in emissions: the literature on the rebound effect (Alcott, 2005; Sorrell, 2009) and the current debate on the climate impact of AI (Luers et al., 2024) alike caution that cheaper computation has historically been accompanied by even faster growth in the volume of use.

7.3 Regulatory trajectory and implications for banks

The regulatory environment is expected to become predictably more stringent over the next three to five years. The development of harmonised energy-efficiency standards under Article 40 of the AI Act and the AI energy and emissions label being prepared by the Commission (European Commission, 2026; European Parliament and Council, 2024) are expected to make standardised supplier data available over the medium term. The timetable is also known: by 2 August 2028, the Commission is required to report on progress in standardisation, and this report may provide for legally binding corrective measures if the voluntary approach proves insufficient (Green Software Foundation, 2026). This standardised supply of data, together with the CSRD data chain (European Parliament and Council, 2022), will convey the same expectation to banks from the perspective of large corporate customers; supervisory interest in the concentration and operational risks associated with the use of AI in banking is already evident (ECB, 2024). Three implications follow for banks. First, an institution that builds its transparency infrastructure (AI inventory, estimation methodology, and supplier data connections) in advance can comply with disclosures that may become man-

datory at low marginal cost, whereas those that wait will be forced to undertake costly retrospective reconstruction. Second, energy costs and the carbon price will gradually become incorporated into the business case for AI applications, meaning that the currently voluntary principle of model-size proportionality will, over the medium term, become a matter of straightforward cost rationality. Third, a bank capable of measuring and communicating its own AI footprint can present itself more credibly as a financier of the sustainability transition – and conversely, the environmental variant of AI washing (Prisznayák, 2023), namely green AI communication without actual measurement, carries growing supervisory and reputational risks.

7.4 Conclusions and recommendations

The study's argument can be summarised in five points, which also constitute recommendations for banking practice and regulation.

First, the banking AI footprint is still small in absolute terms, but it is the sector's fastest-growing operational emissions item, and its structure (cloud-based, Scope 3 in nature and dependent on supplier data) cannot be managed using conventional operational-emissions management tools. Credit institutions should establish an energy-focused inventory of their AI application portfolio now, while adoption is still in its growth phase, because retrospective reconstruction is more costly and less accurate.

Second, the AI governance frameworks – which banks are in any event developing in order to comply with the AI Act (Harkácsi–Szegfű, 2021) – should be expanded to include an energy and environmental dimension: the principle of model-size proportionality, criteria for region and provider selection, and contractual requirements for suppliers to provide energy and emissions data.

Third, sustainability reporting should explicitly extend the double materiality assessment to the digital footprint, including the AI-related footprint; where the materiality threshold has not yet been reached, the trend should nevertheless be presented voluntarily and with the methodology documented.

Fourth, at the level of the regulatory and professional community, the most important missing element is *standardised energy data at the deployer level*. The preparation of an EU AI energy label (European Commission, 2026) and the documentation obligations imposed on GPAI providers point in the right direction, but the collective bargaining power of banks as large corporate consumers – for example, through the formulation of uniform supplier data-reporting requirements at banking-association level – could in itself act as a catalyst for transparency.

Fifth, an open question for research is the empirical measurement of the banking AI footprint: applying the framework proposed in this study to a sample of Hungarian banks, validating specific energy values, and exploring the consumption profiles of agentic applications constitute the next steps.

8 SUMMARY

As a financier of the climate transition, the banking sector has learned to measure what was previously considered immeasurable: its customers' emissions. The next, smaller in scale but symbolically significant task is to transparently account for the environmental footprint of its own digital operations. This study has proposed a conceptual and methodological framework for this purpose: a three-part typology based on the mode of access, the corresponding Scope classification and system boundaries, and the practice of providing range-based estimates with sensitivity analysis. The principal argument of the analysis is that the current constraint on measurability is not theoretical but relates to data access: provider allocation methodologies are most deficient precisely where the banking AI footprint is growing most rapidly, namely in API-based model services. Banks' own AI footprint is currently just a footnote in the emissions inventory; at the current pace of adoption, the footnote may become a chapter within a few years.

Statement on the use of AI

In preparing the original of this study in Hungarian language, I used the OpenAI ChatGPT service, powered by the GPT-5.6 SOL Thinking model, to support source collection and to review the language and style of the text. I used the tool to identify relevant academic literature search terms, research directions and potential sources, as well as to refine the language of individual sentences, improve their style and support the editing of the text. I independently verified all sources suggested by ChatGPT, consulted them at their original locations, and used only sources that I had verified and considered professionally appropriate. I likewise reviewed the linguistic and stylistic modifications suggested by the artificial intelligence. I take responsibility for the content of the study, its analysis, the interpretation of its findings and its conclusions.

REFERENCES

- Acun, B. – Lee, B. – Kazhamiaka, F. – Maeng, K. – Gupta, U. – Chakkaravarthy, M. – Brooks, D. – Wu, C.-J. (2023): Carbon Explorer: A holistic framework for designing carbon aware data-centers. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS 2023)*, 118–132. ACM. <https://doi.org/10.1145/3575693.3575754>.
- Alcott, B. (2005): Jevons' paradox. *Ecological Economics*, 54(1), 9–21. <https://doi.org/10.1016/j.ecolecon.2005.03.020>.
- Aldasoro, I. – Gambacorta, L. – Korinek, A. – Shreeti, V. – Stein, M. (2024): Intelligent financial system: How AI is transforming finance. Bank for International Settlements. *BIS Working Papers* No. 1194. <https://www.bis.org/publ/work1194.htm>.
- Andrae, A. S. G. – Edler, T. (2015): On global electricity usage of communication technology: Trends to 2030. *Challenges*, 6(1), 117–157. <https://doi.org/10.3390/challe6010117>.
- AWS – Amazon Web Services (2025). AWS customer carbon footprint methodology – Model 3.0. *Amazon*. <https://sustainability.aboutamazon.com/aws-customer-carbon-footprint-methodology.pdf>.
- Bagó, P. (2023a): A mesterséges intelligencia lehetőségei a pénzügyekben. *Gazdaság és Pénzügy*, 10(1), 21–38. <https://doi.org/10.33926/GP.2023.1.2>.
- Bagó, P. (2023b): Kiberbiztonság és a mesterséges intelligencia. *Gazdaság és Pénzügy*, 10(2), 196–221. <https://doi.org/10.33926/GP.2023.2.5>.
- Bank of China Limited Magyarország Fióktelepe (2026): Különálló pénzügyi kimutatások az Európai Unió által befogadott IFRS-ekkel összhangban, 2025. december 31.
- Belkhir, L. – Elmeligi, A. (2018): Assessing ICT global emissions footprint: Trends to 2040 & recommendations. *Journal of Cleaner Production*, 177, 448–463. <https://doi.org/10.1016/j.jclepro.2017.12.239>.
- Bender, E. M. – Gebru, T. – McMillan-Major, A. – Shmitchell, S. (2021): On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*, 610–623. ACM. <https://doi.org/10.1145/3442188.3445922>.
- Benedek, B. – Nagy, B. Zs. (2023): Hagyományos versus MI-alapú csalásfelderítés: Költséghatékonyság a gépjármű-biztosítások területén. *Hitelintézeti Szemle*, 22(2), 77–100. <https://doi.org/10.25201/HSZ.22.2.77>.
- BIS – Bank for International Settlements. (2024). Project Gaia: Enabling climate risk analysis using generative AI. *BIS Innovation Hub*. <https://www.bis.org/publ/othp84.htm>.
- Bódi, R. – Fáykiss, P. – Nyikes, Á. (2023): A technológiai óriásvállalatok rendszerkockázatai és szabályozásuk – „Too Big(Tech) To Fail?” *Hitelintézeti Szemle*, 22(1), 5–21. <https://doi.org/10.25201/HSZ.22.1.5>.
- Brevini, B. (2020): Black boxes, not green: Mythologizing artificial intelligence and omitting the environment. *Big Data & Society*, 7(2). <https://doi.org/10.1177/2053951720935141>.
- Cao, L. (2022): AI in finance: Challenges, techniques, and opportunities. *ACM Computing Surveys*, 55(3) 64. <https://doi.org/10.1145/3502289>.
- CIB Bank Zrt. (2026). *Fenntarthatósági jelentés 2025*.
- Cowls, J. – Tsamados, A. –, Taddeo, M. – Floridi, L. (2023): The AI gambit: Leveraging artificial intelligence to combat climate change – opportunities, challenges, and recommendations. *AI & Society*, 38(1), 283–307. <https://doi.org/10.1007/s00146-021-01294-x>.
- Crawford, K. (2021): Atlas of AI: Power, politics, and the planetary costs of artificial intelligence. *Yale University Press*.

- Dauvergne, P. (2022): Is artificial intelligence greening global supply chains? Exposing the political economy of environmental costs. *Review of International Political Economy*, 29(3), 696–718. <https://doi.org/10.1080/09692290.2020.1814381>.
- de Vries, A. (2023): The growing energy footprint of artificial intelligence. *Joule*, 7(10), 2191–2194. <https://doi.org/10.1016/j.joule.2023.09.004>.
- Desislavov, R. – Martínez-Plumed, F. – Hernández-Orallo, J. (2023): Trends in AI inference energy consumption: Beyond the performance-vs-parameter laws of deep learning. *Sustainable Computing: Informatics and Systems*, 38, 100857. <https://doi.org/10.1016/j.suscom.2023.100857>.
- Deutsche Bank (2026): Sustainable Finance Framework (hatályos 2026. január 1-től).
- Dhar, P. (2020): The carbon impact of artificial intelligence. *Nature Machine Intelligence*, 2(8), 423–425. <https://doi.org/10.1038/s42256-020-0219-9>.
- Doumpos, M. – Zopounidis, C. – Gounopoulos, D. – Platanakis, E. – Zhang, W. (2023): Operational research and artificial intelligence methods in banking. *European Journal of Operational Research*, 306(1), 1–16. <https://doi.org/10.1016/j.ejor.2022.04.027>.
- EBA – European Banking Authority (2020): Report on big data and advanced analytics (EBA/REP/2020/01). <https://www.eba.europa.eu>.
- ECB – European Central Bank (2024): The rise of artificial intelligence: Benefits and risks for financial stability. In *Financial Stability Review, May 2024*. <https://www.ecb.europa.eu>.
- Európai Bizottság (2026): Targeted consultation on measuring energy consumption and emissions of AI models and systems. *Európai Bizottság, DG CNECT*. Közzététel: 2026. 04.07. A konzultáció lezárva: 2026.05.10. <https://digital-strategy.ec.europa.eu/en/consultations/targeted-consultation-measuring-energy-consumption-and-emissions-ai-models-and-systems>.
- Európai Parlament és Tanács (2022): Az (EU) 2022/2464 irányelv a fenntarthatósággal kapcsolatos vállalati beszámolásról (CSRD). Az *Európai Unió Hivatalos Lapja*, L 322.
- Európai Parlament és Tanács (2024): Az (EU) 2024/1689 rendelet a mesterséges intelligenciára vonatkozó harmonizált szabályok megállapításáról (AI-rendelet). Az *Európai Unió Hivatalos Lapja*, L sorozat.
- Evident (2025): The Evident AI Index for Banks: Key findings report 2025. *Evident Insights*. <https://evidentinsights.com/insights>.
- Evident (2026): AI use case trends in banking: Q1 2026. *Evident Insights*. <https://evidentinsights.com/insights/banking-use-case-trends-q1-2026>.
- Fares, O. H. – Butt, I. – Lee, S. H. M. (2023): Utilization of artificial intelligence in the banking sector: A systematic literature review. *Journal of Financial Services Marketing*, 28(4), 835–852. <https://doi.org/10.1057/s41264-022-00176-7>.
- Freitag, C. – Berners-Lee, M. – Widdicks, K. – Knowles, B. – Blair, G. S – Friday, A. (2021): The real climate and transformative impact of ICT: A critique of estimates, trends, and regulations. *Patterns*, 2(9), 100340. <https://doi.org/10.1016/j.patter.2021.100340>.
- FSB – Financial Stability Board: (2017). Artificial intelligence and machine learning in financial services: Market developments and financial stability implications. <https://www.fsb.org>.
- Google Cloud (2026): Carbon Footprint reporting methodology. Google. <https://docs.cloud.google.com/carbon-footprint/docs/methodology>.
- GRÁNIT Bank Nyrt. (2026): Az Igazgatóság jelentése a 2025. évi üzleti tevékenységről: konszolidált üzleti jelentés és fenntarthatósági nyilatkozat.
- Green Software Foundation (2026): Software Carbon Intensity for AI and EU AI Act environmental compliance: A framework for measuring and reducing AI's climate impact (1.0, 2026.02.15.). *Software Standards Working Group, Green Software Foundation*. <https://greensoftware.foundation/policy/research/sci-ai-eu-ai-act/>.

- Gupta, U. – Kim, Y. G. – Lee, S. – Tse, J. – Lee, H.-H. S. – Wei, G.-Y. – Brooks, D. – Wu, C.-J. (2021): Chasing carbon: The elusive environmental footprint of computing. In *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)* (854–867). IEEE. <https://doi.org/10.1109/HPCA51647.2021.00076>.
- Harkácsi, G. J. – Szegfű, L. P. (2021): A megfélelősbiztosítási funkció szerepe a digitalizáció, mesterséges intelligencia és robotizáció idején a pénzügyi szektorban. *Hitelintézeti Szemle*, 20(1), 152–170. <https://doi.org/10.25201/HSZ.20.1.152170>.
- Henderson, P. – Hu, J. – Romoff, J. – Brunskill, E. – Jurafsky, D. – Pineau, J. (2020): Towards the systematic reporting of the energy and carbon footprints of machine learning. *Journal of Machine Learning Research*, 21(248), 1–43. <https://doi.org/10.48550/arXiv.2002.05651>.
- Hentzen, J. K. – Hoffmann, A. – Dolan, R. – Pala, E. (2022): Artificial intelligence in customer-facing financial services: A systematic literature review and agenda for future research. *International Journal of Bank Marketing*, 40(6), 1299–1336. <https://doi.org/10.1108/IJBM-09-2021-0417>.
- Horváth, D. (2022): FinTech és blokklánc alapú megoldások alkalmazási lehetőségei a zöld pénzügyekben. *Vezetéstudomány / Budapest Management Review*, 53(4), 41–54. <https://doi.org/10.14267/VEZTUD.2022.04.04>.
- IEA – International Energy Agency (2024): Electricity 2024: Analysis and forecast to 2026. IEA. <https://www.iea.org>.
- IEA – International Energy Agency (2025). *Energy and AI*. IEA. <https://www.iea.org>.
- ING Bank N.V. Magyarországi Fióktelepe. (2026). IFRS egyedi beszámoló, 2025. december 31.
- Jones, N. (2018): How to stop data centres from gobbling up the world's electricity. *Nature*, 561(7722), 163–166. <https://doi.org/10.1038/d41586-018-06610-y>.
- Kaack, L. H. – Donti, P. L. – Strubell, E. – Kamiya, G. – Creutzig, F. – Rolnick, D. (2022): Aligning artificial intelligence with climate change mitigation. *Nature Climate Change*, 12(6), 518–527. <https://doi.org/10.1038/s41558-022-01377-7>.
- Königstorfer, F. – Thalmann, S. (2020): Applications of artificial intelligence in commercial banks – A research agenda for behavioral finance. *Journal of Behavioral and Experimental Finance*, 27, 100352. <https://doi.org/10.1016/j.jbef.2020.100352>.
- Koomey, J. – Berard, S. – Sanchez, M. – Wong, H. (2011): Implications of historical trends in the electrical efficiency of computing. *IEEE Annals of the History of Computing*, 33(3), 46–54. <https://doi.org/10.1109/MAHC.2010.28>.
- Lacoste, A. – Luccioni, A. – Schmidt, V. – Dandres, T. (2019): Quantifying the carbon emissions of machine learning. *arXiv*. <https://doi.org/10.48550/arXiv.1910.09700>.
- Li, P. – Yang, J. – Islam, M. A. – Ren, S. (2023): Making AI less “thirsty”: Uncovering and addressing the secret water footprint of AI models. *arXiv*. <https://doi.org/10.48550/arXiv.2304.03271>.
- Ligozat, A.-L. – Lefèvre, J. – Bugeau, A. – Combaz, J. (2022): Unraveling the hidden environmental impacts of AI solutions for environment: Life cycle assessment of AI solutions. *Sustainability*, 14(9), 5172. <https://doi.org/10.3390/su14095172>.
- Luccioni, A. S. – Jernite, Y. – Strubell, E. (2024): Power hungry processing: Watts driving the cost of AI deployment? In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency (FAccT '24)* (pp. 85–99). ACM. <https://doi.org/10.1145/3630106.3658542>.
- Luccioni, A. S. – Viguier, S. – Ligozat, A.-L. (2023): Estimating the carbon footprint of BLOOM, a 176B parameter language model. *Journal of Machine Learning Research*, 24(253), 1–15. <https://doi.org/10.48550/arXiv.2211.02001>.
- Luers, A. – Koomey, J. – Masanet, E. – Gaffney, O. – Creutzig, F. – Lavista Ferres, J. – Horvitz, E. (2024): Will AI accelerate or delay the race to net-zero emissions? *Nature*, 628(8009), 718–720. <https://doi.org/10.1038/d41586-024-01137-x>.

- Lülök, G. – Sebestyén, Z. (2025): A mesterséges intelligencia legújabb alkalmazási trendjei a bankszektorban. *Hitelintézeti Szemle*, 24(2), 47–74. <https://doi.org/10.25201/HSZ.24.2.47>.
- MagNet Bank Zrt. (2026). A MagNet Bank Zrt. egyedi és csoportszintű, valamint a Polgári Bank Zrt. egyedi szintű nyilvánosságra hozatali dokumentuma az 575/2013/EU rendelet alapján, 2025. december 31.
- Malmodin, J. – Lundén, D. (2018): The energy and carbon footprint of the global ICT and E&M sectors 2010–2015. *Sustainability*, 10(9), 3027. <https://doi.org/10.3390/su10093027>.
- Masanet, E. – Shehabi, A. – Lei, N. – Smith, S. – Koomey, J. (2020): Recalibrating global data center energy-use estimates. *Science*, 367(6481), 984–986. <https://doi.org/10.1126/science.aba3758>.
- MBH Bank Nyrt. (2026). Éves jelentés a 2025. december 31-vel zárult évről (benne: Konszolidált fenntarthatósági jelentés).
- McKinsey & Company (2024). How generative AI can help banks manage risk and compliance. *McKinsey Risk & Resilience Practice*.
- Microsoft (2026): Emissions Impact Dashboard for Azure: Methodology and access. *Microsoft*. <https://www.microsoft.com/en-us/sustainability/emissions-impact-dashboard>.
- Milana, C. – Ashta, A. (2021): Artificial intelligence techniques in finance and financial markets: A survey of the literature. *Strategic Change*, 30(3), 189–209. <https://doi.org/10.1002/jsc.2403>.
- MNB – Magyar Nemzeti Bank. (2024): Zöld pénzügyi jelentés. *Magyar Nemzeti Bank*. <https://www.mnb.hu>.
- MNB – Magyar Nemzeti Bank (2025): Fintech és digitalizációs jelentés. *Magyar Nemzeti Bank*. <https://www.mnb.hu>.
- Mytton, D. (2021): Data centre water consumption. *npj Clean Water*, 4, Article 11. <https://doi.org/10.1038/s41545-021-00101-w>.
- Mytton, D. – Ashtine, M. (2022): Sources of data center energy estimates: A comprehensive review. *Joule*, 6(9), 2032–2056. <https://doi.org/10.1016/j.joule.2022.07.011>.
- Nishant, R. – Kennedy, M. – Corbett, J. (2020): Artificial intelligence for sustainability: Challenges, opportunities, and a research agenda. *International Journal of Information Management*, 53, 102104. <https://doi.org/10.1016/j.ijinfomgt.2020.102104>.
- Oberbank AG. (2026): Bond investor presentation, June 2026.
- OECD (2021): Artificial intelligence, machine learning and big data in finance: Opportunities, challenges, and implications for policy makers. *OECD Publishing*. <https://www.oecd.org>.
- OECD (2022): Measuring the environmental impacts of artificial intelligence compute and applications: The AI footprint (OECD Digital Economy Papers No. 341). *OECD Publishing*. <https://doi.org/10.1787/7babf571-en>.
- Patterson, D. – Gonzalez, J. – Le, Q. – Liang, C. – Munguia, L.-M – Rothchild, D. – So, D. – Texier, M. – Dean, J. (2021): Carbon emissions and large neural network training. *arXiv*. <https://doi.org/10.48550/arXiv.2104.10350>.
- PCAF – Partnership for Carbon Accounting Financials (2022): The global GHG accounting and reporting standard for the financial industry (2nd ed.). *PCAF*.
- Prisznyák, A. (2023): A bankrobotika horizontális és vertikális értékkeremtése és az AI-washing jelenség. *Hitelintézeti Szemle*, 22(3), 97–122. <https://doi.org/10.25201/HSZ.22.3.97>.
- Radovanović, A. – Koningstein, R. – Schneider, I. – Chen, B. – Duarte, A. – Roy, B. – Xiao, D. – Haridasan, M. – Hung, P. – Care, N. – Talukdar, S. – Mullen, E. – Smith, K. – Cottman, M. – Cirne, W. (2023): Carbon-aware computing for datacenters. *IEEE Transactions on Power Systems*, 38(2), 1270–1280. <https://doi.org/10.1109/TPWRS.2022.3173250>.
- Rolnick, D. – Donti, P. L. – Kaack, L. H. – Kochanski, K. – Lacoste, A. – Sankaran, K. – Ross, A. S. – Milojevic-Dupont, N. – Jaques, N. – Waldman-Brown, A. – Luccioni, A. S. – Maharaj, T.

- Sherwin, E. D. – Mukkavilli, S. K. – Kording, K. P. – Gomes, C. P. – Ng, A. Y. – Hassabis, D. – Platt, J. C. – Bengio, Y. (2022): Tackling climate change with machine learning. *ACM Computing Surveys*, 55(2), Article 42. <https://doi.org/10.1145/3485128>.
- Schwartz, R. – Dodge, J. – Smith, N. A. – Etzioni, O. (2020): Green AI. *Communications of the ACM*, 63(12), 54–63. <https://doi.org/10.1145/3381831>.
- Sevilla, J. – Heim, L. – Ho, A. – Besiroglu, T. – Hobbhahn, M. – Villalobos, P. (2022): Compute trends across three eras of machine learning. In *2022 International Joint Conference on Neural Networks (IJCNN)* (pp. 1–8). IEEE. <https://doi.org/10.1109/IJCNN55064.2022.9891914>.
- Shehabi, A. – Smith, S. J. – Hubbard, A. – Newkirk, A. – Lei, N. – Siddik, M. A. B. – Holecek, B. – Koomey, J. – Masanet, E. – Sartor, D. (2024): 2024 United States data center energy usage report (LBNL-2001637). *Lawrence Berkeley National Laboratory*.
- Siddik, M. A. B. – Shehabi, A. – Marston, L. (2021): The environmental footprint of data centers in the United States. *Environmental Research Letters*, 16(6), 064017. <https://doi.org/10.1088/1748-9326/abfbai>.
- Sorrell, S. (2009): Jevons’ Paradox revisited: The evidence for backfire from improved energy efficiency. *Energy Policy*, 37(4), 1456–1469. <https://doi.org/10.1016/j.enpol.2008.12.003>.
- Strubell, E. – Ganesh, A. – McCallum, A. (2019): Energy and policy considerations for deep learning in NLP. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics* (3645–3650). ACL. <https://doi.org/10.18653/v1/P19-1355>.
- Taddeo, M. – Tsamados, A. – Cows, J. – Floridi, L. (2021): Artificial intelligence and the climate emergency: Opportunities, challenges, and recommendations. *One Earth*, 4(6), 776–779. <https://doi.org/10.1016/j.oneear.2021.05.018>.
- Vőneki, T. Zs. – Lamanda, G. (2020): Az ESG-kockázatokkal kapcsolatos banki közzétételek tartalom-elemzése – Pillanatkép a hazai nagybankok közzétételi gyakorlatáról. *Gazdaság és Pénzügy*, 7(4), 420–433. <https://doi.org/10.33926/GP.2020.4.3>.
- Thompson, N. C. – Greenewald, K. – Lee, K. – Manso, G. F. (2021): Deep learning’s diminishing returns: The cost of improvement is becoming unsustainable. *IEEE Spectrum*, 58(10), 50–55. <https://doi.org/10.1109/MSPEC.2021.9563954>.
- UniCredit Bank Hungary Zrt. (2026): Éves jelentés és beszámoló a 2025-ös évről.
- van Wynsberghe, A. (2021): Sustainable AI: AI for sustainability and the sustainability of AI. *AI and Ethics*, 1(3), 213–218. <https://doi.org/10.1007/s43681-021-00043-6>.
- Verdecchia, R. – Sallou, J. – Cruz, L. (2023): A systematic review of Green AI. *WIREs Data Mining and Knowledge Discovery*, 13(4), e1507. <https://doi.org/10.1002/widm.1507>.
- Vinuesa, R. – Azizpour, H. – Leite, I. – Balaam, M. – Dignum, V. – Domisch, S. – Felländer, A. – Langhans, S. D. – Tegmark, M. – Fuso Nerini, F. (2020): The role of artificial intelligence in achieving the Sustainable Development Goals. *Nature Communications*, 11, Article 233. <https://doi.org/10.1038/s41467-019-14108-y>.
- Vives, X. (2019): Digital disruption in banking. *Annual Review of Financial Economics*, 11, 243–272. <https://doi.org/10.1146/annurev-financial-100719-120854>.
- WRI–WBCSD (2004). The greenhouse gas protocol: A corporate accounting and reporting standard (Rev. ed.). *World Resources Institute – World Business Council for Sustainable Development*.
- Wu, C.-J. – Raghavendra, R. – Gupta, U. – Acun, B. – Ardalani, N. – Maeng, K. – Chang, G. – Aga, F. – Huang, J. – Bai, C. – Gschwind, M. – Gupta, A. – Ott, M. – Melnikov, A. – Candido, S. – Brooks, D. – Chauhan, G. – Lee, B. – Lee, H.-H. S. – Hazelwood, K. (2022): Sustainable AI: Environmental implications, challenges and opportunities. In *Proceedings of Machine Learning and Systems* (Vol. 4, pp. 795–813). MLSys. <https://doi.org/10.48550/arXiv.2111.00364>.