

Commentary

Does not compute: Avoiding pitfalls assessing the Internet's energy and carbon impacts

Jonathan Koomey^{1,*} and Eric Masanet²

Jonathan Koomey is president of Koomey Analytics and has in the past been a visiting professor at Stanford University, Yale University, and UC Berkeley. He's one of the leading international experts on the economics of climate solutions and the energy and environmental effects of information technology. Dr. Koomey holds M.S. and Ph.D. degrees from the Energy and Resources Group at UC Berkeley and an A.B. in History and Science from Harvard University. He is the author or coauthor of more than 200 articles and reports and nine books, including *Turning Numbers into Knowledge: Mastering the Art of Problem Solving* and *Cold Cash, Cool Climate: Science-Based Advice for Ecological Entrepreneurs*. More at <http://www.koomey.com>.



Eric Masanet is the Mellichamp Chair in Sustainability Science for Emerging Technologies at the University of California, Santa Barbara, where he holds appointments in the Bren School of Environmental Science and Management and the Department of Mechanical Engineering. He has authored more than 130 scientific publications on sustainability modeling of energy and materials demand systems, with particular focuses on data centers and IT systems. He holds a Ph.D. in mechanical engineering from UC Berkeley, with a focus on sustainable manufacturing.

Introduction

Accurately assessing information technology's (IT's) energy and environmental effects is not easy. In this commentary, we describe four common pitfalls analysts face when creating, evaluating, sharing, or publishing such statistics and offer recommendations for avoiding them.

A recurrent theme is that well-intentioned research often overestimates IT's electricity use and climate impacts, sometimes by orders of magnitude. These results then become "factoids" that spread quickly as people share them and the media report them.¹

The problem, of course, is that incorrect numbers can have real-world consequences when widely believed. Consider the growth of global Internet traffic in the 1990s, as described by Coffman and Odlyzko.² Those data flows doubled every year or so for years but doubled every 100 days for parts of 1995 and 1996. Growth reverted to doubling every year after that, but extrapolations based on "doubling every 100 days" led to vast overinvestment in network capacity around 2000. In 2002, more than 97% of fiber capacity sat unused,³ and because electricity used by operating network equipment is roughly constant, regardless of whether it is used or not, that overcapacity had significant energy implications.

Consequential mistakes can result when analysts' inherent curiosity about an important topic area collides with a pervasive lack of accurate and up-to-date information. IT changes so quickly that most data characterizing it are obsolete in short order, and people's inability to accurately predict the effects of exponential change in multiple dimensions can make things even worse. The best data on IT's energy and emissions characteristics are closely held proprietary secrets among tech companies, which compounds these problems.

When misinformation takes root, it can be difficult to correct due to an asymmetry often called "Brandolini's Law." It states (paraphrased for a family audience) that "the amount of energy needed to refute misinformation is ten times greater than the amount of energy needed to create




Figure 1. Annual energy use and network data flows for two large network providers, expressed as an index relative to 2016 = 1.0

Telefonica data start in 2015, and Cogent data start in 2016. Although data traffic increased by more than 3-fold since 2016, network energy use dropped by 2.4% (Telefonica) and 50% (Cogent) over the same time period, demonstrating that historical network energy use is not directly proportional to data traffic. Notably, Telefonica’s data traffic jumped by 45% in 2020 due in part to COVID (compared to 2019), with no reported increase in network energy use. Cogent’s electricity use dropped 21% from 2019 to 2020 even as data traffic increased 38%. Sources: <https://www.cogentco.com/en/about-cogent/corporate-responsibility/our-environment> and <https://www.telefonica.com/en/web/responsible-business/report-2020>.

misinformation.” Rigorous research to produce accurate data will always trail misinformation, because it’s harder and takes longer to get the numbers right. That systemic bias has become even stronger with the rise of social media, which has accelerated the pace of sharing factoids.

Four common research pitfalls

Perhaps the most common pitfall is that analysts conducting retrospective analysis or making projections ignore, misunderstand, or mischaracterize changes in key parameters over time. The challenge is three-fold:

- Rapid changes and pervasive data gaps make accurate, up-to-date assessments difficult or impossible.
- Analytical errors, failure to fact-check cited statistics, and inadequate documentation make accurate assessments less common than they ought to be.
- Analysts and the media often recycle published data long after they are no longer valid.

The past decade saw highly cited top-down projections of data-center energy use that either underestimated subsequent energy-efficiency gains⁴ or ignored them altogether,⁵ leading to predictions of massive energy-demand growth by the decade’s end. Subsequent retrospective analysis⁶ revealed that these simplified projections erred by failing to anticipate the huge energy-efficiency improvements that occurred in servers, storage, network, and power and cooling systems in the interim.

Sometimes efficiency increases will exceed growth in service demand, for example, in the case of Cogent, a major network provider that has tripled data traffic over the last 4 years while network energy use declined by 50%. Network operator Telefonica more than tripled data traffic while energy use stayed roughly constant over that same period (Figure 1). Sometimes service demand will increase more rapidly than efficiency, leading to growth in energy use, as is the case for cloud/hyperscale data centers

globally.⁶ However, it is almost never true that Internet electricity use scales exactly proportionately to growth in service demand.

A recent example of using outdated information is the BBC using an emissions intensity of email published in 2010⁷ to estimate the carbon emissions effects of email in 2020, without correcting for the large efficiency gains in processing, sending, and storing data over that period^{6,8,9} or for improvements in the emissions intensity of electricity. In the UK, each kilowatt-h of electricity emitted less than half as much carbon dioxide in 2020 as it did in 2010. Many technology companies also purchased renewable power in recent years, which likely further reduced relevant emissions intensities.

Citing previous work without checking its rigor and accuracy is common in the face of pervasive data gaps, but it can accelerate the dissemination of questionable factoids that might otherwise remain obscure. This tendency has been documented for investment research reports in IT electricity use,¹⁰ but it happens in the academic literature as well.

Sometimes research is peer reviewed but still contains significant errors.¹¹ Sometimes research is accurately critiqued after publication, but disputed results are still highly cited.^{12,13} A complicating factor is that not every study provides enough information to enable easy replication, which makes it difficult for other analysts to properly assess analysis results.

A second pitfall is assuming that short-term changes in computing services must lead to proportional and immediate changes in electricity use. A recent example relates to the increase in data flows experienced by many networks from 2019 to 2020.¹⁴

Telefonica showed an increase in data demand above the trend in 2020 due

in part to COVID with virtually no change in network energy use, and Cogent registered a similar jump in data demand in 2020 while energy use declined 21% from 2019 (Figure 1). Models displaying the effect of increased short-term data demand that fail to account for non-proportionality between energy and data flows in network equipment risk yielding inflated environmental-impact results (like, for example, in Obringer et al.¹⁵). Most other IT equipment is not perfectly energy proportional either, a fact that analysts ignore at their peril, but this issue is especially acute for network equipment.

A third pitfall is making long-term projections, even if these projections explicitly account for technological change in efficiency and other key drivers of electricity use. IT changes so quickly¹⁶ that even projections extending only a few years are highly uncertain (but can sometimes be valid and useful). Applying exponential growth rates in demand growth for more than a few years can result in eye-popping projected changes,^{4,11,12} which invariably lead to media attention. However, the possibility of errors in such long-term forecasts is large, and small differences in assumptions can lead to significant variation in results. Researchers should resist the temptation to project IT electricity use beyond a few years (which is about how far into the future manufacturer product roadmaps extend) and should use caution in drawing conclusions from such extrapolations.

A fourth pitfall is drawing broad conclusions based on trends in only one part of the IT system. Analytical rigor comes from analyzing a whole system, which is often difficult due to pervasive data gaps. However, just focusing on one highly visible part of the system can give a mistaken impression about what's happening with the electricity used by the whole system. That is

because changes in one part of the IT system often offset changes in other parts of the IT system, and such substitution effects are real and powerful.

For example, the tremendous growth in the cloud data-center segment (which includes the world's largest "hyper-scale" data centers) has led some to predict massive future growth in global data-center energy use.¹⁷ Between 2010 and 2018, the workloads hosted by this segment increased by 2,600%, whereas its estimated electricity use increased by 500%.⁶ Despite the rapid growth of this segment, the global energy use of *all* data centers grew far more modestly, rising by less than 10%.⁶

Much of the compute output from cloud/hyperscale data centers has displaced traditional in-house data centers that use several times more electricity to perform the same tasks.⁶ As long as there are still inefficient facilities to displace, large increases in demands for computing need not drive large increases in total electricity use.

How can we do better?

Industry must improve data sharing

Each of the four pitfalls traces back to pervasive data gaps, which still pose a substantial problem for future research. Most such data gaps result from reluctance among tech companies to share the latest proprietary information.

Some broadband and mobile internet providers have released network energy-intensity data over time. Some data-center companies have released data on their infrastructure efficiency, use of renewable power, and sometimes even overall data-center energy use. However, comparatively few global companies have released such data, and those data often lack the granularity necessary for analysts to understand technology trends. The industry still needs to develop consistent cross-industry protocols for assessing

drivers of low emissions. More companies need to report such information, and they need to focus measurements of zero-emissions electricity on hour-by-hour accounting instead of the annual accounting now common for most companies.

The analysis community can also help. Performance and efficiency metrics for IT are difficult to create, and where they exist are a closely guarded proprietary secret for virtually all companies. Information technology workloads vary greatly, and developing benchmarks is difficult even within companies (and is even harder for enabling cross-industry comparisons). One way for analysts to help industry overcome its reluctance to release metrics of IT efficiency and performance would be to develop *indices* of progress in IT efficiency over time for servers, storage, and networking. This approach could allow companies to share internal trends in IT efficiency without revealing sensitive information like numbers of transactions or other business drivers. Efficiency would need to be measured with standard protocols for assessing service demand and rely on measured energy use for each type of IT equipment, with safeguards for keeping key data proprietary. More research is also needed to help develop such metrics.

Analysts must report results more precisely and transparently

Analysts should always separately report electricity use, emissions intensities, and absolute emissions, giving exact dates and locations to which these estimates apply. This practice helps avoid confusion when numbers are used by other researchers. Analysts should avoid averaging key parameters over long periods for these fast-moving technologies. And of course, analysts should release complete information to enable replication of calculations by others, which is a hallmark of an open and rigorous scientific process.

Analysts must exercise restraint

It is common practice when data aren't available to make assumptions for key parameters, based on intuition, common knowledge, or physical principles. The problem for new and fast-changing systems like IT is that such assumptions are often incorrect. Where key parameters are unknown, it is better to collect data instead of making assumptions, and if data aren't available, it is better not to publish estimates at all, or to couch the results in an appropriately cautious way, recognizing the limitations of assumptions in this fast-moving space.

That lesson suggests that analysts should not project IT electricity use more than a few years into the future and should avoid the use of simple extrapolations altogether. It also implies that analysts should not use estimates and data from even a few years ago to assess current IT systems without correcting them for changes over time. As a corollary, analysts should exercise caution when *citing* factoids about IT electricity use and encourage the media to exercise similar restraint.

Journals must improve peer review

Journal editors need to recognize that papers claiming to estimate internet energy use must be evaluated carefully by people with real subject-matter expertise, especially when researchers from outside the small community of experts who study IT electricity use delve into this field. One way to improve the accuracy of peer review is to require full disclosure of all assumptions, data, and models for every article, including separate reporting of electricity use and emissions intensities, so that results can be applied correctly to other times and locations. Journal editors need to ensure thoroughness of reviews, and that might require changes to the incentives editors offer

to reviewers, but that of course is a problem beyond our purview here.

Conclusions

IT engenders endless fascination. It has the power to reshape society at a rapid pace, but understanding developments in this fast-moving field requires topic knowledge, restraint, and serious due diligence. Analysts need to engage this subject matter with respect and care, otherwise consequential errors in public policy will result.

ACKNOWLEDGMENTS

The authors wrote this article on their personal time as a labor of love. It had no funding sources.

AUTHOR CONTRIBUTIONS

Both authors contributed equally to the writing of the article.

DECLARATION OF INTERESTS

Jonathan Koomey is also a senior advisor to the World Resource Institute (WRI). Koomey declares modest ongoing research funding on efficient computing technologies from Sony Interactive Entertainment, modest occasional funding on efficient computing technologies from AMD, and modest ongoing research funding from WRI on climate solutions. Eric Masanet also holds an adjunct professorship in the McCormick School of Engineering and Applied Science at Northwestern University and is a faculty scientist at Lawrence Berkeley National Laboratory.

1. Koomey, J. (2014). *Separating Fact from Fiction: A Challenge For The Media*. IEEE Consumer Electronics Magazine 3, 9–11.
2. Coffman, K.G., and Odlyzko, A.M. (2001). Growth of the Internet (AT&T). <http://www.dtc.umn.edu/~odlyzko/doc/oft.internet.growth.pdf>.
3. Drazan, Y.J. (2002) Behind the Fiber Glut. The Wall Street Journal. September 26, 2002. B1.
4. Andrae, A.S.G., and Edler, T. (2015). On Global Electricity Usage of Communication

Technology: Trends to 2030. Challenges 6, 117–157.

5. Belkhir, L., and Elmeligi, A. (2018). Assessing ICT global emissions footprint: Trends to 2040 & recommendations. J. Clean. Prod. 177, 448–463.
 6. Masanet, E., Shehabi, A., Lei, N., Smith, S., and Koomey, J. (2020). Recalibrating global data center energy-use estimates. Science 367, 984–986.
 7. Berners-Lee, M. (2010). *How Bad Are Bananas?: The carbon footprint of everything* (Profile Books).
 8. Koomey, J., and Naffziger, S. (2016). Energy efficiency of computing: What's next? Electronic Design. <https://www.electronicdesign.com/electronic-design/whitepaper/21804096/energy-efficiency-of-computing-whats-next-pdf-download>.
 9. Aslan, J., Mayers, K., Koomey, J.G., and France, C. (2018). Electricity Intensity of Internet Data Transmission: Untangling the Estimates. J. Ind. Ecol. 22, 785–798.
 10. Koomey, J. (2017). *Turning Numbers into Knowledge: Mastering the Art of Problem Solving, Third Edition* (Analytics Press).
 11. Mora, C., Rollins, R.L., Taladay, K., Kantar, M.B., Chock, M.K., Shimada, M., and Franklin, E.C. (2018). Bitcoin emissions alone could push global warming above 2°C. Nat. Clim. Chang. 8, 931–933.
 12. Masanet, E., Shehabi, A., Lei, N., Vranken, H., Koomey, J., and Malmodyn, J. (2019). Implausible projections overestimate near-term Bitcoin CO2 emissions. Nat. Clim. Chang. 9, 653–654.
 13. Mora, C., Rollins, R.L., Taladay, K., Kantar, M.B., Chock, M.K., Shimada, M., and Franklin, E.C. (2019). Mora et al. reply. Nature Climate Change 9, 658–659.
 14. Malmodyn, J. (2020). The power consumption of mobile and fixed network data services—The case of streaming video and downloading large files. Electronics Goes Green 2020.
 15. Obringer, R., Rachunok, B., Maia-Silva, D., Arbabzadeh, M., Nateghi, R., and Madani, K. (2021). The overlooked environmental footprint of increasing Internet use. Resour. Conserv. Recycling 167, 105389.
 16. Koomey, J.G., Berard, S., Sanchez, M., and Wong, H. (2011). Implications of Historical Trends in The Electrical Efficiency of Computing. IEEE Ann. Hist. Comput. 33, 46–54.
 17. Mills, M.P. (2020). *Digital Cathedrals* (Encounter Books).
- ¹Koomey Analytics, Burlingame, California, USA
²University of California, Santa Barbara, Santa Barbara, CA, USA
 *Correspondence: jgkoomey@gmail.com
<https://doi.org/10.1016/j.joule.2021.05.007>