

INTELLIGENZA 2.0

□ PENFL □ □ R 19

Privatsphäre im Zentrum


Marcel Neuhäusler



<https://www.flickr.com/photos/34811759@N08/29725977283>

Wie steht es um die Privatsphäre?

<https://www.bbc.com/news/technology-49343774>


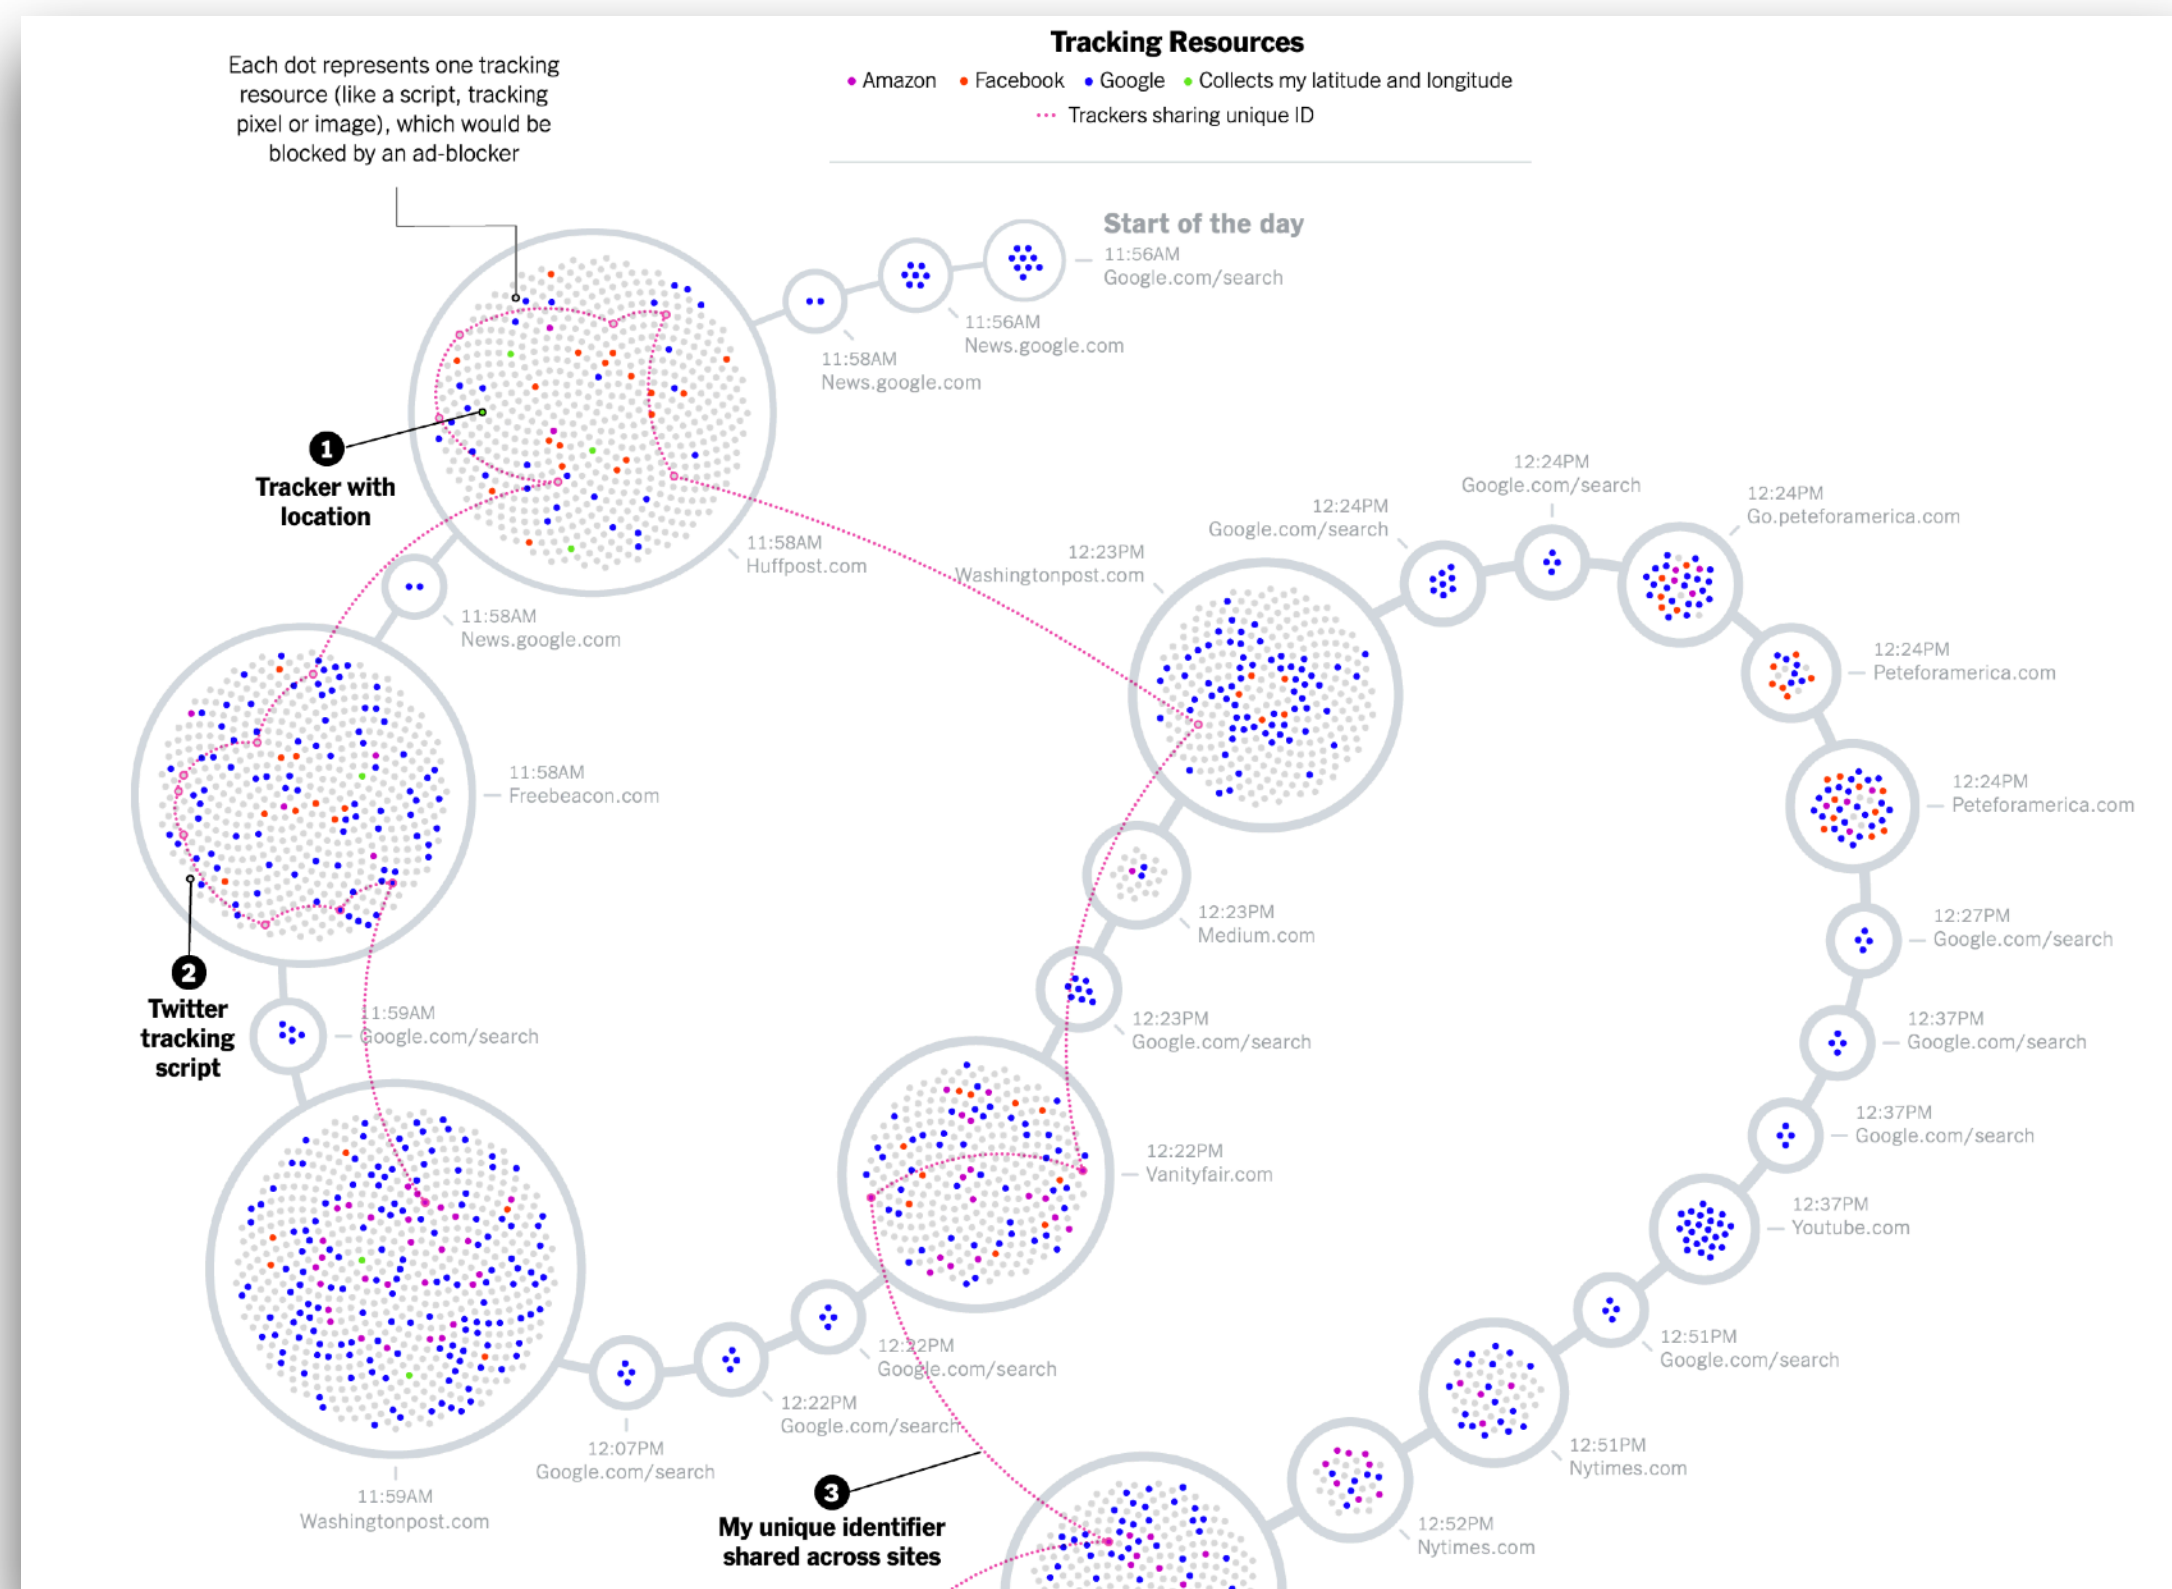


As well as fingerprint records, the researchers say they found photographs of people, facial recognition data, names, addresses, passwords, employment history and records of when they had accessed secure areas.



Wie steht es um die Privatsphäre?

<https://www.nytimes.com/interactive/2019/08/23/opinion/data-internet-privacy-tracking.html>




Tracker with browser details

Fingerprinting

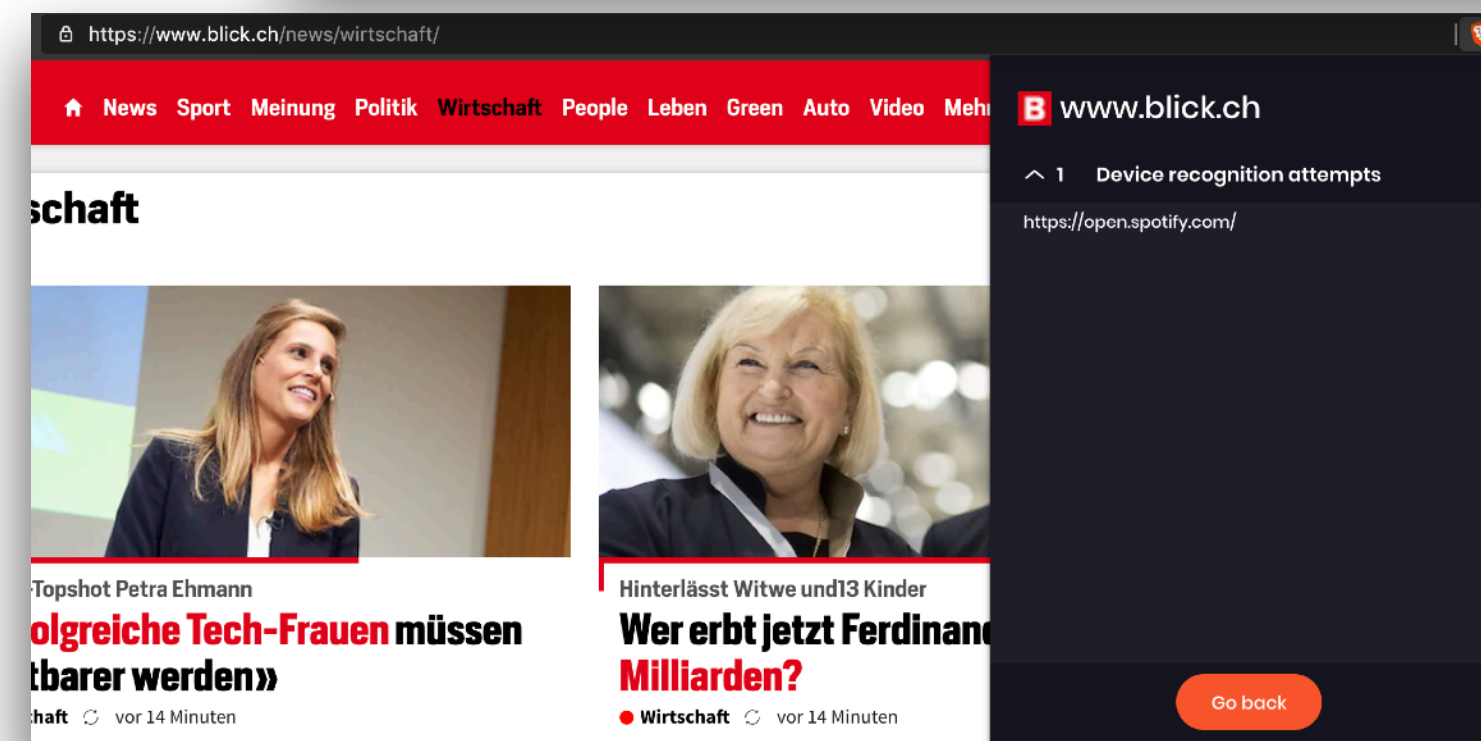
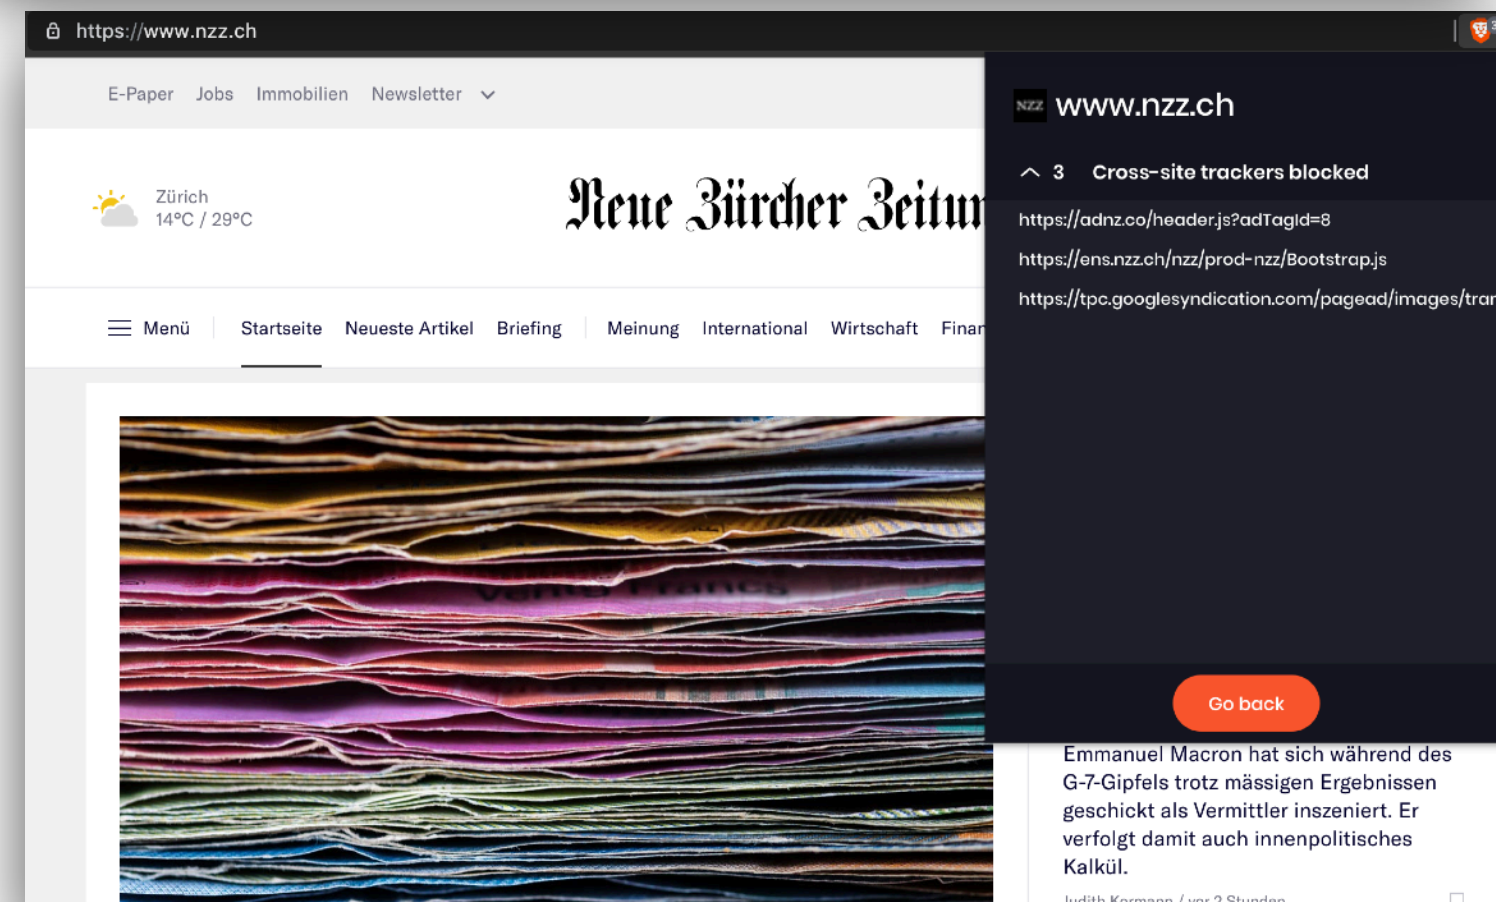
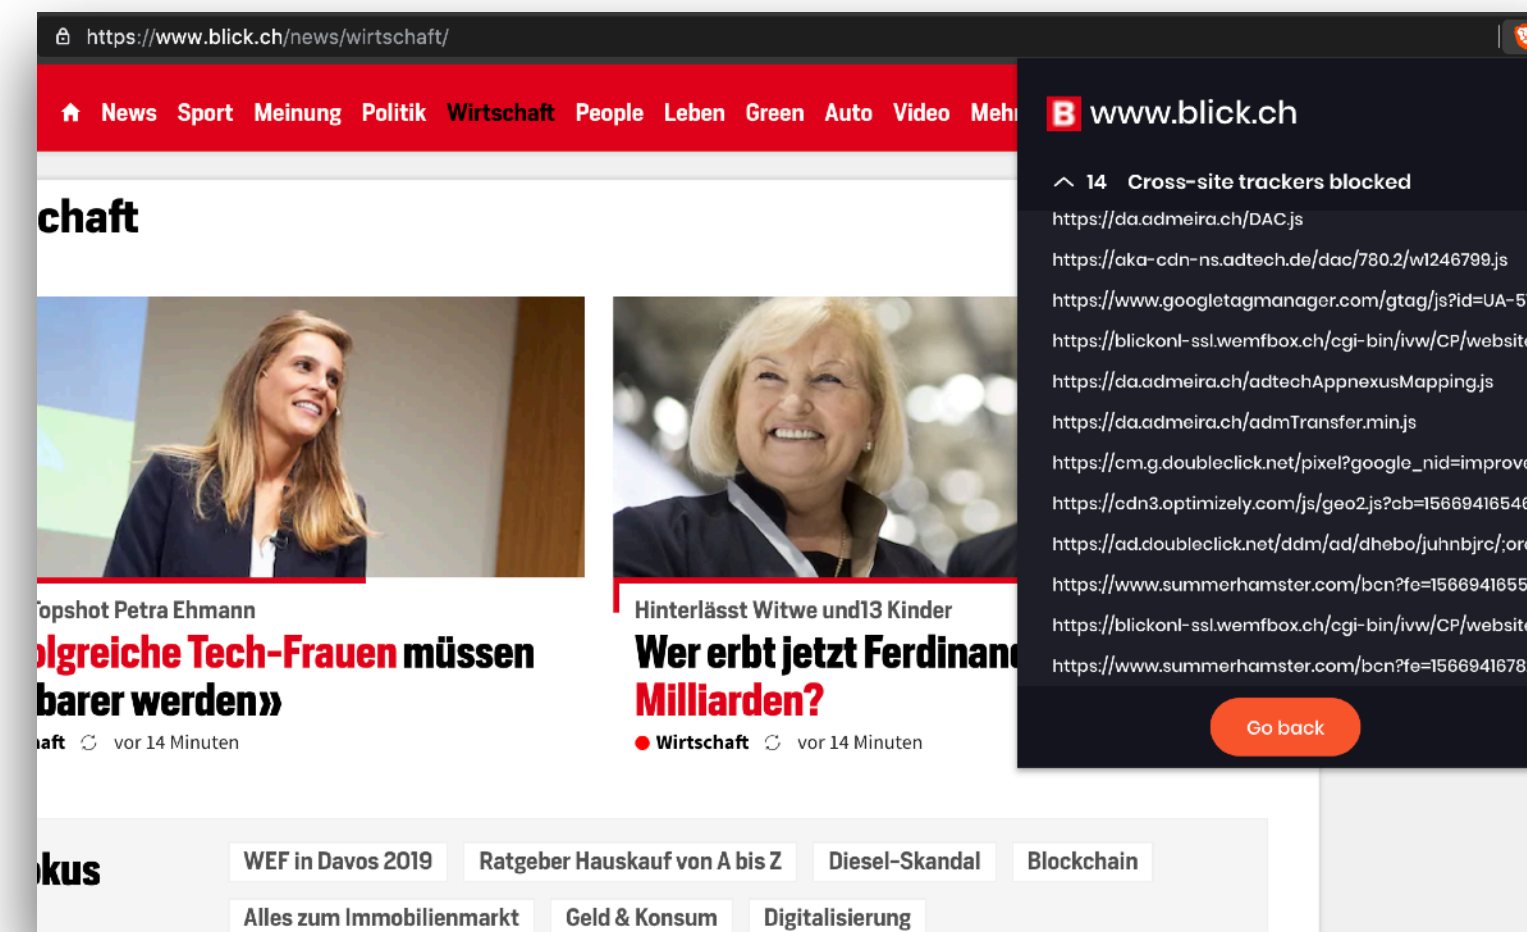
Even when companies don't have an ID to track me, they can use signals from my computer to guess who I am across sites. That's partly why trackers like this one received more information about my computer than you could imagine being useful, like my precise screen size. Other trackers received my screen resolution, browser information, operating system details, and more.

Google, Google, everywhere

Google's own domains don't contain that many trackers. The same is true for Facebook. But that's because they place most of their trackers on other websites. Google was present on every site I visited, collecting information on where I live, the device I used and everything I looked at.




Wie steht es um die Privatsphäre?



Wie steht es um die Privatsphäre?

<https://www.statista.com/statistics/266206/googles-annual-global-revenue/>



by J. Clement,

last edited May 22, 2019

This statistic displays Google's annual revenue from 2002 to 2018. In the most recently reported fiscal year, Google's revenue amounted to 136.22 billion US dollars. Google's revenue is largely made up by [advertising revenue](#), which amounted to 116 billion US dollars in 2018. As of May 2018, Google ranks first among worldwide internet companies, with a [market capitalization](#) of 510 billion U.S. dollars.

SHARE OF U.S. INTERNET USERS WHO USE GOOGLE SERVICES AND PRODUCTS ON A DAILY BASIS

69%

SHARE OF U.S. INTERNET USERS WHO HAVE A PERSONAL GOOGLE USER ACCOUNT

75%


SHARE OF U.S. INTERNET USERS WHO THINK GOOGLE IS A MODERN COMPANY

64%



Wie steht es um die Privatsphäre?

<https://www.statista.com/statistics/277229/facebook-annual-revenue-and-net-income/>



by J. Clement,

last edited Feb 4, 2019

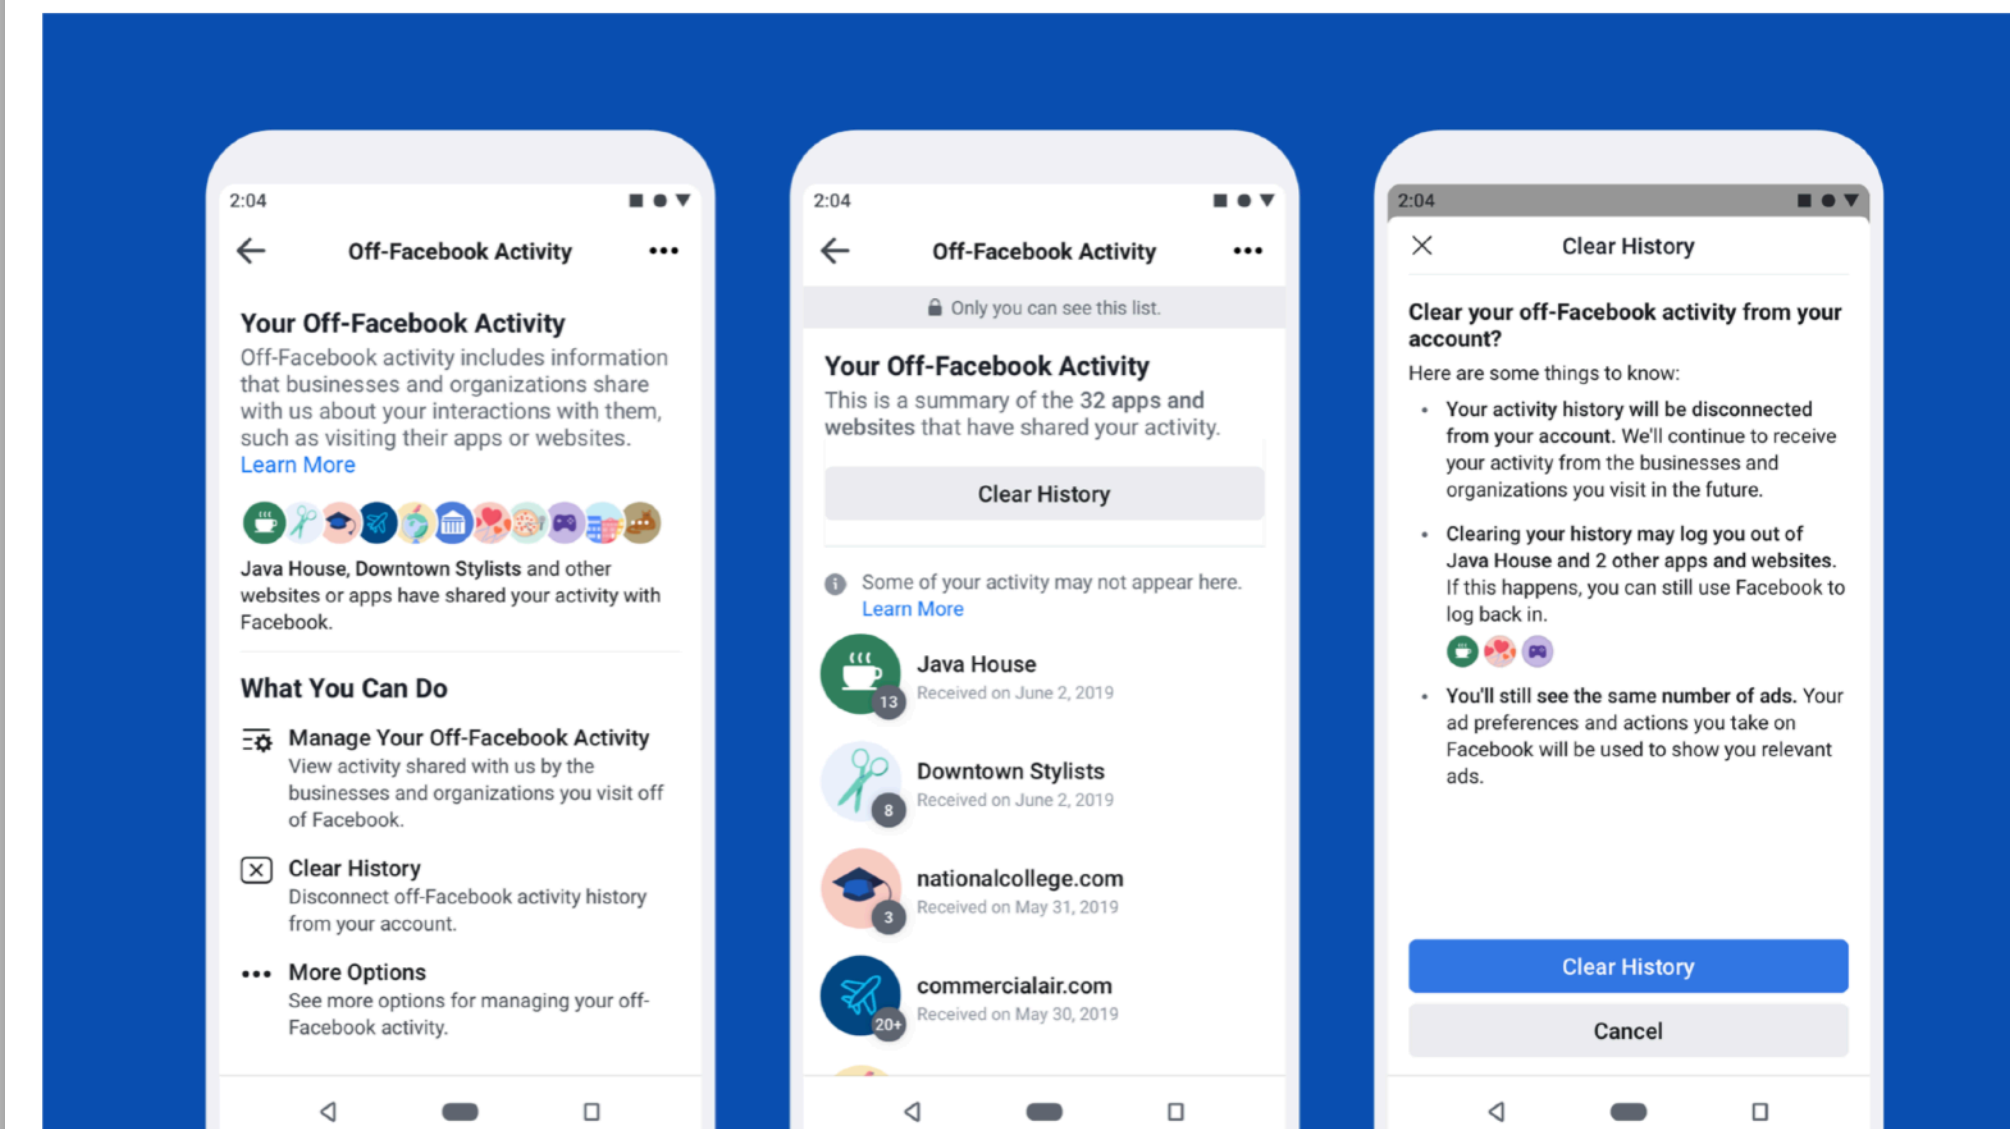
This statistic presents Facebook's annual revenue and net income from 2007 to 2018. Facebook's revenue grew from 7.87 billion in 2013 to 55.8 billion US dollars in 2018. That year, the social network accumulated a net income of 22.1 billion US dollars, ranking **first among social media companies in annual revenues**. The majority of Facebook revenues are generated via advertising - the shift towards e-commerce and online marketing has never been more evident than in 2016. Other revenue-generating factors are online marketing and online promotion during events such as the **U.S. election in 2016**, during which online advertising played a major role.



Wie steht es um die Privatsphäre?

<https://engineering.fb.com/data-infrastructure/off-facebook-activity/>

Redesigning our systems to provide more control over Off-Facebook activity



By Mayur Patel Scott Renfro Benjamin Strahs



Wie steht es um die Privatsphäre?

<https://www.nytimes.com/interactive/2019/04/16/opinion/facial-recognition-new-york-city.html>



We were able to turn those cameras into a facial recognition-powered tracking system for less than \$100 — using a service available to anyone with an internet connection and a credit card. The process was completely legal.

To demonstrate how easy it is to track people without their knowledge, we collected public images of people who worked near Bryant Park (available on their employers' websites, for the most part) and ran one day of footage through Amazon's commercial facial recognition service. Our system detected 2,750 faces from a nine-hour period (not necessarily unique people, since a person could be captured in multiple frames). It returned several possible identifications, including one frame matched to a head shot of Richard Madonna, a professor at the SUNY College of Optometry, with an 89 percent similarity score. The total cost: about \$60.

89% match




Image from
captured video



Image from SUNY
College of Optometry



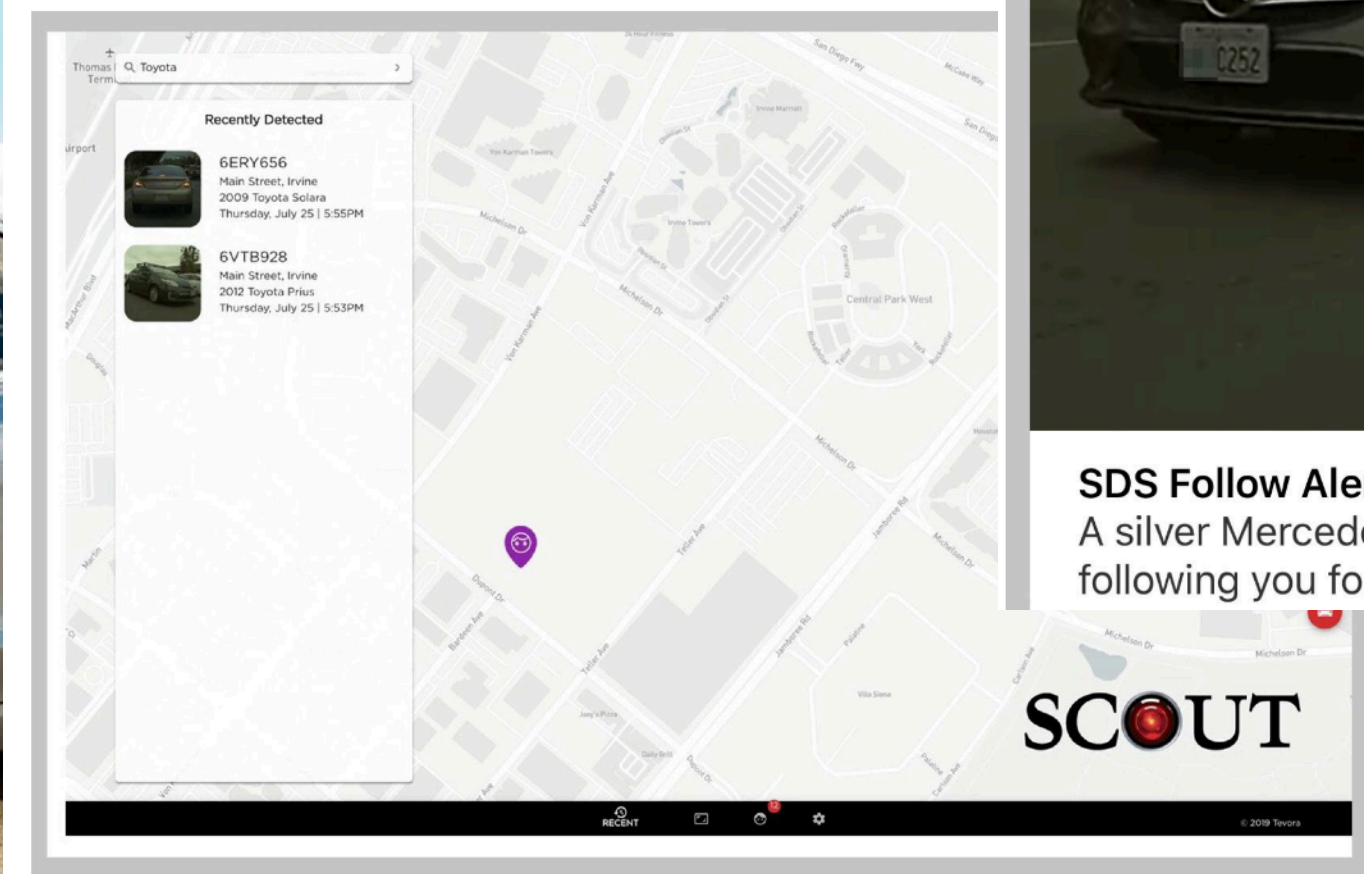

Wie steht es um die Privatsphäre?

<https://www.wired.com/story/tesla-surveillance-detection-scout/>

WIRED BUSINESS CULTURE GEAR IDEAS SCIENCE SECURITY TRANSPORTATION SIGN IN SUBSCRIBE

This Tesla Mod Turns a Model S Into a Mobile 'Surveillance Station'

The Surveillance Detection Scout can track license plates and faces near your Tesla—with all the privacy concerns that implies.



SDS Follow Alert
A silver Mercedes-Benz C300 has been following you for 7 minutes.

SCOUT


A screenshot of Surveillance Detection Scout's interface, showing recently detected license plates.
COURTESY TRUMAN KAIN

Wie steht es um die Privatsphäre?

<https://www.engadget.com/2019/08/09/ring-rewarded-suspicious-activity-reports/>

Ring rewarded users for reporting 'suspicious activities'

The company promised users discounts and swag bags.



Mariella Moon, @mariella_moon
08.09.19 in [Security](#)

Comments

192
Shares



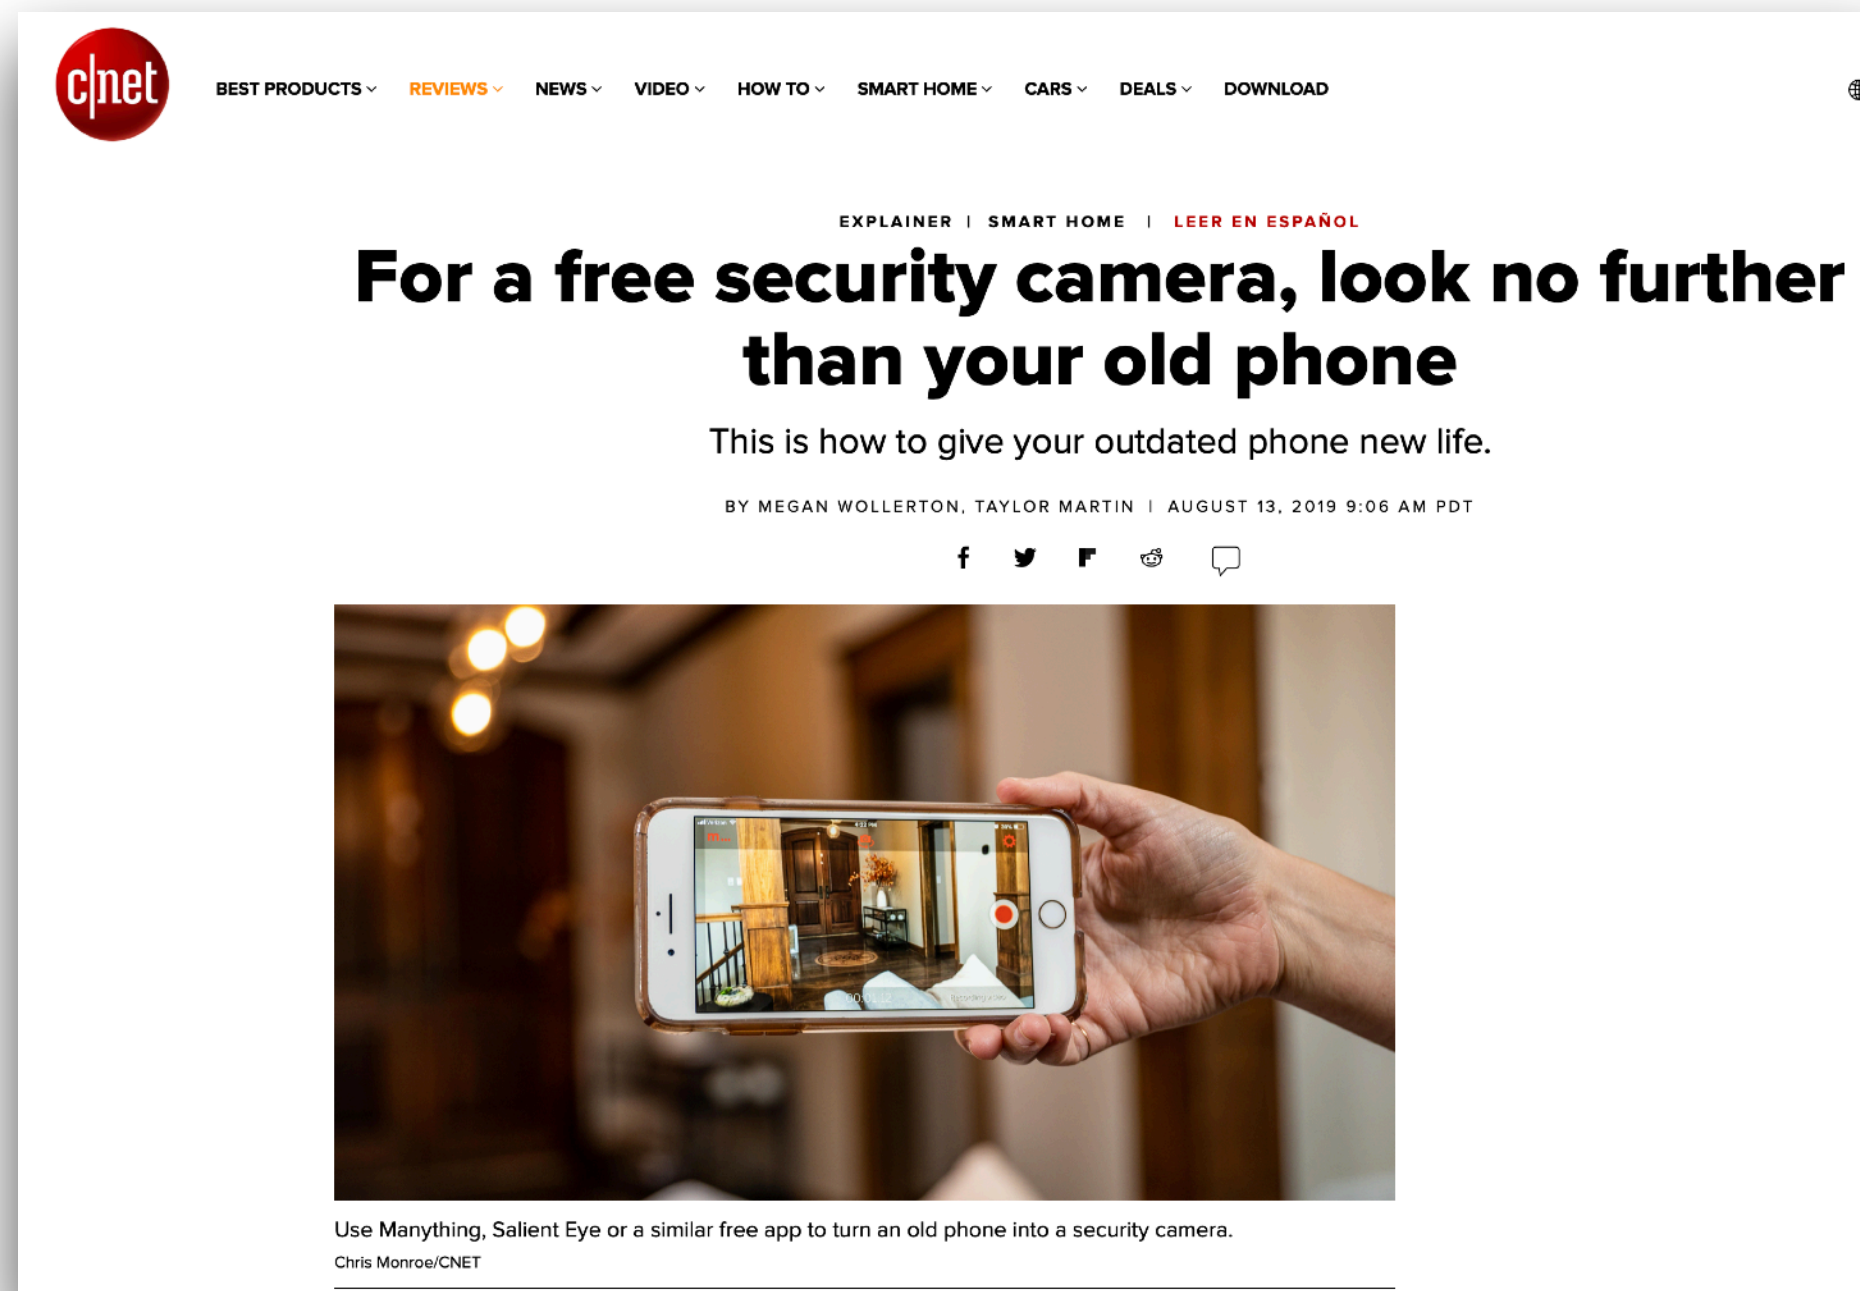
Ring encouraged users to narc on their neighbors in exchange for swag bags and discounts, according to [Motherboard](#). The publication found an internal company slide presentation that shows the instructions the [controversial](#) Amazon-owned home security service gave communities on how to form "Digital Neighborhood Watch" groups.

Apparently, the company encouraged users to form Watch teams and to post videos on social media to receive promo codes for future devices. It also promised free swag to anyone who recruits 10 new users and to those who blog about Ring "in a positive way," as well as 50 percent discounts on Ring products to those who can solve a crime with the help of local cops. The slides didn't elaborate on what those swag bags contained.



Wie steht es um die Privatsphäre?

<https://www.cnet.com/how-to/for-a-free-security-camera-look-no-further-than-your-old-phone/>



1. Alfred Premium - enjoy every moment in clarity

Alfred Premium is a monthly subscription service that offers **HD viewing/recording**, **Ad-removal**, **30-day Event storage**, **Motion Detection Schedule** ^{1.3k}, **Motion Events up to 120-sec long**, as well as **Zoom!**

Key Features

© HD viewing/recording

To be more specific, we mean both **HD viewing** AND **recording**. We don't like to brag, but the making, as well as maintenance, of an HD version is monumental work. Offering HD video entails much more than switching on a button. Just to give you an idea, here is what Alfred Premium really is all about:

- A more powerful infrastructure working continuously to provide the most advanced Alfred experience to date, including a larger server for HD Motion Detection videos.
- Superior technology for compressing, encoding, and decoding the videos.

2. Alfred Plus - say goodbye to ads

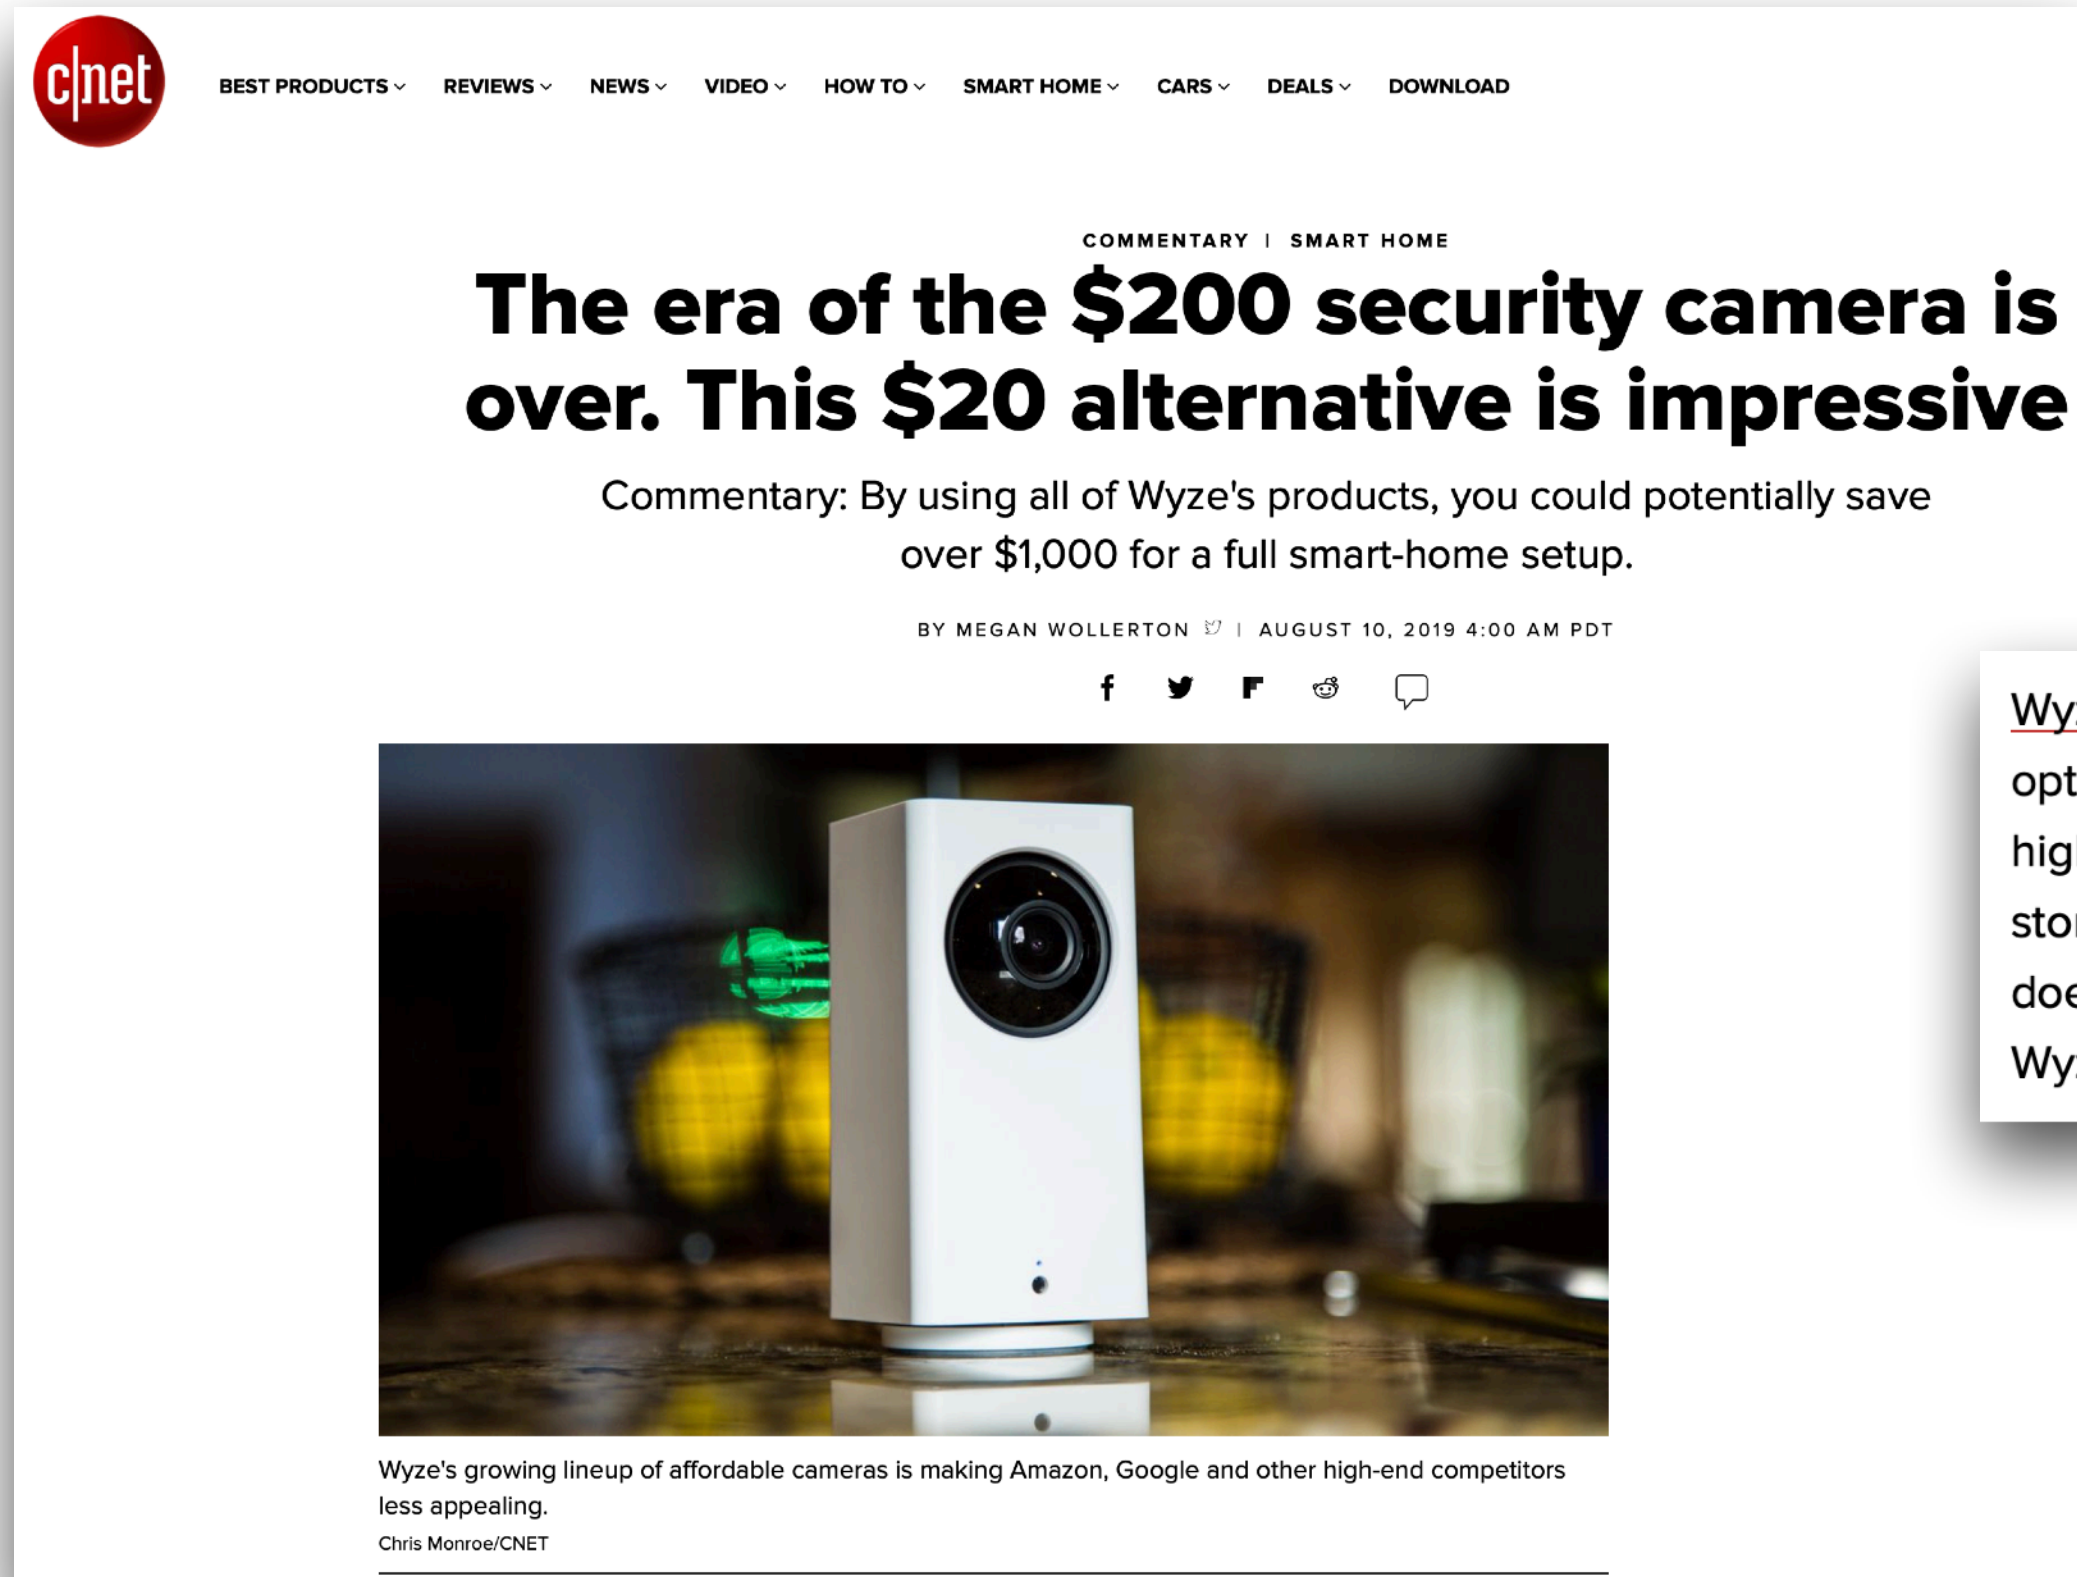
If you only want to remove ads, you can consider **Alfred Plus**.

Alfred Plus is a one-time payment that removes all ads from all the devices under the same account. While we try to make ads as least aesthetically displeasing as possible, we understand that some users prefer not to see them at all. By choosing Alfred Plus, you will never ever see another ad on any of your devices!



Wie steht es um die Privatsphäre?

<https://www.cnet.com/news/the-era-of-the-200-security-camera-is-over/>



The screenshot shows a CNET article page. At the top is the CNET logo and a navigation bar with links: BEST PRODUCTS, REVIEWS, NEWS, VIDEO, HOW TO, SMART HOME, CARS, DEALS, and DOWNLOAD. The article title is "The era of the \$200 security camera is over. This \$20 alternative is impressive" in bold black text. Below the title is a sub-headline: "Commentary: By using all of Wyze's products, you could potentially save over \$1,000 for a full smart-home setup." The author is listed as "BY MEGAN WOLLERTON" and the date as "AUGUST 10, 2019 4:00 AM PDT". There are social media sharing icons for Facebook, Twitter, Facebook, and a comment icon. Below the text is a photograph of a white Wyze security camera on a wooden surface. At the bottom of the article, there is a short paragraph: "Wyze's growing lineup of affordable cameras is making Amazon, Google and other high-end competitors less appealing." followed by the byline "Chris Monroe/CNET".


The era of the \$200 security camera is over. This \$20 alternative is impressive

Commentary: By using all of Wyze's products, you could potentially save over \$1,000 for a full smart-home setup.

BY MEGAN WOLLERTON | AUGUST 10, 2019 4:00 AM PDT

Wyze's growing lineup of affordable cameras is making Amazon, Google and other high-end competitors less appealing.
Chris Monroe/CNET

Wyze added person detection to its security cameras in July, available for free as an optional feature. Basically, Wyze cameras have most, if not all of the things, that many high-end DIY cameras offer -- but Wyze offers them for free. Nest charges for cloud storage, motion zones and person alerts via its Nest Aware subscription service; it doesn't have local storage at all. It also can't listen for special frequencies, like the Wyze cameras can.



Intelligenz 0.1

Termin: Dienstag, 11. März 1986, 9.30 – 17.00 Uhr

Ort: ETH Zürich-Zentrum, Hauptgebäude, Rämistrasse 101, Hörsaal E3

Leitung: Dr. H.-J. Appelrath, ETH Zürich, Institut für Informatik, 8092 Zürich

Ziel: Seit einigen Monaten existiert ein Arbeitskreis «Wissensbasierte Systeme», in dem sich Vertreter aus Industrie, Wirtschaft und Hochschulen zu einem regelmässigen Erfahrungsaustausch über ihre Forschungs- und Entwicklungsarbeit auf dem Gebiet wissensbasierter Expertensysteme treffen.


Diesem Arbeitskreis gehören an:
BBC Baden
Ciba Geigy Basel
Contraves AG Zürich
Elektrowatt Zürich
ETH Zürich
Hasler AG Bern
Landis & Gyr AG Zug
Mettler Instrumente AG Greifensee
Schweizerische Bankgesellschaft Zürich
Schweizerische Kreditanstalt Zürich
Universität Bern
Universität Zürich.

Auf der Informationstagung werden in kompakter Form einige Aktivitäten von Mitgliedern des Arbeitskreises vorgestellt.

Anmeldung: bis zum 28. 2. 1986 mit anhängender Anmeldekarte

Auskunft: Frau S. Papp, ETH Zürich, Institut für Informatik, 8092 Zürich, Tel. (01) 256 22 53

**«Wissensbasierte Systeme:
Forschungs- und Entwicklungstendenzen»**



Dienstag, 11. März 1986

ETH Zürich

PROGRAMM


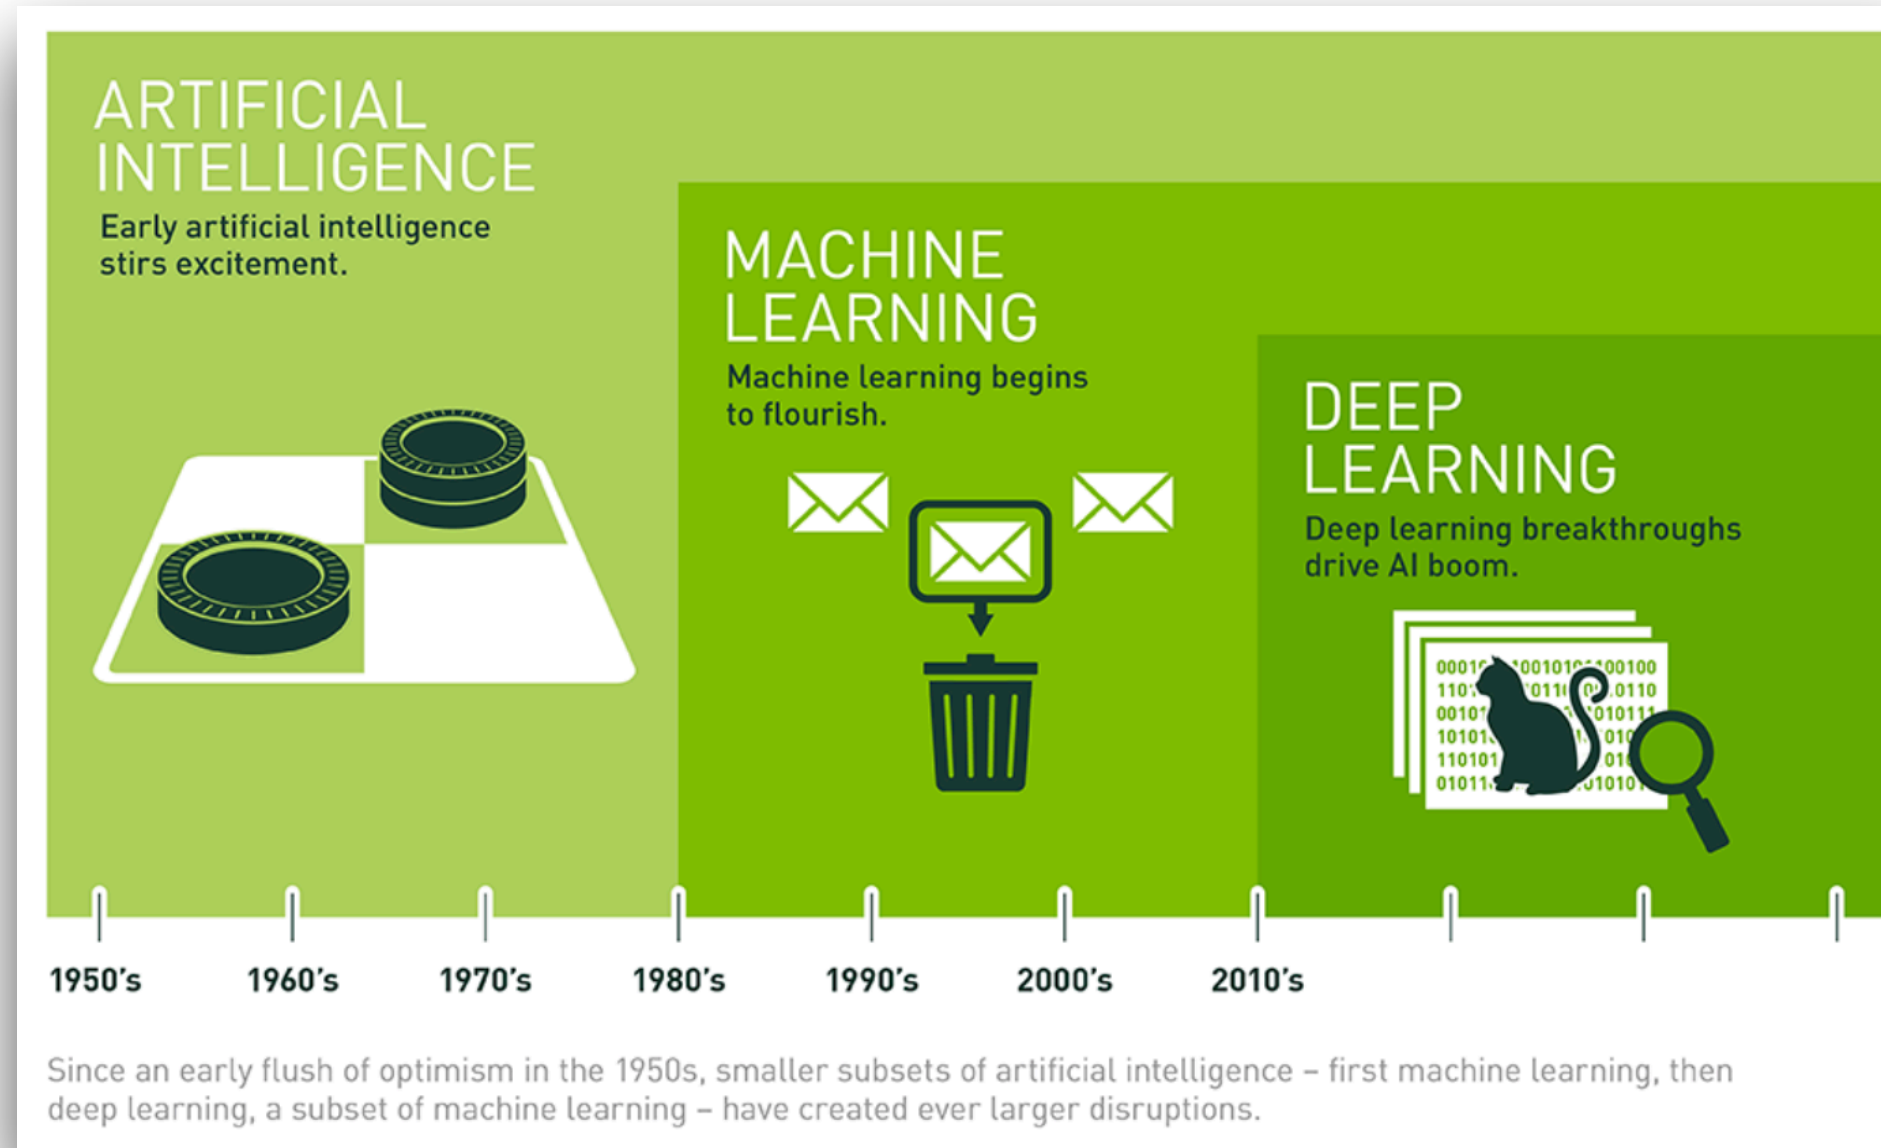
9.30 – 11.00 Uhr	Begrüssung und Eröffnung (C.A. Zehnder)	14.15 – 15.15 Uhr	Mettler Instrumente AG Greifensee (N. E. Fuchs): «Expertensysteme bei Mettler»
	BBC Baden (J. Kriz, C. Muller, M. Vitins, H. Sugaya): «Wissensbasierte Systeme für Konfiguration und Diagnose»		Schweizerische Bankgesellschaft, Zürich (M. Wietlisbach, A. Eisendle): «Erfahrungen mit dem AI-Shell KEE auf Lisp-Maschinen»
	ETH Zürich, Institut für Elektronik (J. Dähler): «Ausführbare Spezifikation nebenläufiger Systeme mit Prolog»	15.15 – 15.45 Uhr	Kaffeepause
	ETH Zürich, Institut für Informatik (H.-J. Appelrath, M. Ester, H. Jasper, M. Kiener, A. Ultsch): «Entwurf und Implementierung portabler Prolog-Werkzeuge»	15.45 – 17.00 Uhr	Universität Bern, Institut für Informatik und angewandte Mathematik (H. Bunka, F. Davis, E. Gmür, F. Grimm, D. Jiang): «Wissensbasierte Systeme und Computer-Sehen»
11.00 – 11.30 Uhr	Kaffeepause		Universität Bern, Institut für Wirtschaftsinformatik (E. Brakemeier, M. Roth, H. Sailer, A. von Weissenfluh): «Vergleich von Expertensystemtools auf PCs»
11.30 – 12.45 Uhr	ETH Zürich, Institut für Kommunikationstechnik (Frau E. Egeli, F. Ade, F. Klein): «Verstehen grafischer Strukturen mit Prolog»		Universität Zürich (M. Hess, R. Pfeifer, T. Rothenfluh): «Natürlichsprachliche Interfaces» «Fallbasiertes Problemlösen» «Ein Expertensystem für Differentialdiagnose in der Psychiatrie»
	Hasler AG Bern (B. Humpert, E. Peter): «Experten- und Deduktions-Systeme bei Hasler»		
	Landis & Gyr AG Zug (H. Lienhard, J. Olnhoff, G. Brun, A. Schürch): «Ein wissensbasiertes System zur Anlagenkonfiguration»		
12.45 – 14.15 Uhr	Mittagpause		



Intelligenz 1.8

<https://www.mathworks.com/discovery/deep-learning.html>

<https://mapr.com/blog/demystifying-ai-ml-dl/>



While deep learning was first theorized in the 1980s, there are two main reasons it has only recently become useful:

1. Deep learning requires large amounts of **labeled data**. For example, driverless car development requires millions of images and thousands of hours of video.
2. Deep learning requires substantial **computing power**. High-performance GPUs have a parallel architecture that is efficient for deep learning. When combined with clusters or cloud computing, this enables development teams to reduce training time for a deep learning network from weeks to hours or less.



Figure 3. Comparing a machine learning approach to categorizing vehicles (left) with deep learning (right).

Intelligenz 1.8


<https://www.nytimes.com/2018/11/25/business/china-artificial-intelligence-labeling.html?smid=nytcore-ios-share>



A.I. has to be taught. It must digest vast amounts of tagged photos and videos before it realizes that a black cat and a white cat are both cats. This is where the data factories and their workers come in.

Taggers helped AInnovation, a Beijing-based A.I. company, fix its automated cashier system for a Chinese bakery chain. Users could put their pastry under a scanner and pay for it without help from a human. But nearly one-third of the time, the system had trouble telling muffins from doughnuts or pork buns thanks to store lighting and human movement, which made images more complex. Working with photos from the store's interior, the taggers got the accuracy up to 99 percent, said Liang Rui, an AInnovation project manager.

"All the artificial intelligence is built on human labor," Mr. Liang said.



Intelligenz 1.8

<https://arxiv.org/abs/1906.02243>

Energy and Policy Considerations for Deep Learning in NLP

Emma Strubell, Ananya Ganesh, Andrew McCallum

(Submitted on 5 Jun 2019)

Recent progress in hardware and methodology for training neural networks has ushered in a new generation of large networks trained on abundant data. These models have obtained notable gains in accuracy across many NLP tasks. However, these accuracy improvements depend on the availability of exceptionally large computational resources that necessitate similarly substantial energy consumption. As a result these models are costly to train and develop, both financially, due to the cost of hardware and electricity or cloud compute time, and environmentally, due to the carbon footprint required to fuel modern tensor processing hardware. In this paper we bring this issue to the attention of NLP researchers by quantifying the approximate financial and environmental costs of training a variety of recently successful neural network models for NLP. Based on these findings, we propose actionable recommendations to reduce costs and improve equity in NLP research and practice.


Comments: In the 57th Annual Meeting of the Association for Computational Linguistics (ACL). Florence, Italy. July 2019

Consumption	CO ₂ e (lbs)
Air travel, 1 passenger, NY↔SF	1984
Human life, avg, 1 year	11,023
American life, avg, 1 year	36,156
Car, avg incl. fuel, 1 lifetime	126,000
Training one model (GPU)	
NLP pipeline (parsing, SRL)	39
w/ tuning & experimentation	78,468
Transformer (big)	192
w/ neural architecture search	626,155



Ist die Privatsphäre noch zu retten?

<https://www.nytimes.com/2019/08/21/technology/personaltech/alexa-siri-google-assistant-listen.html>



In two separate reports in The Guardian and Bloomberg News, whistle-blowers recently said they had listened in on [Siri recordings](#) and [Alexa activations](#) that inadvertently recorded couples having sex and criminals making drug deals. Another publication, VRT, chronicled how a Google subcontractor shared more than [1,000 excerpts from Google recordings](#), which journalists then used to identify some individuals.

In the tech industry, it's an open secret that artificial intelligence isn't all that smart yet. It takes [lots of people manually sifting through data](#) to train the computing systems. That means humans occasionally cull through voice recordings to train Alexa, Siri and Google to understand the nuances of speech, such as distinguishing spoken words like "Austin" from "Boston," or "U2" from "YouTube."



Ist die Privatsphäre noch zu retten?

<https://www.nau.ch/news/europa/ist-die-privatsphäre-noch-zu-retten-65471459>

 Suchwort oder Ortschaft eingeben

Home > News > Europa

Ist die Privatsphäre noch zu retten?

Sprachassistenten erfreuen sich grosser Beliebtheit. Doch Datenschützer warnen, dass die Konzerne so mitbekommen, was beim Verbraucher zuhause passiert.



Datenschützer befürchten, dass Konzerne wie Amazon oder Google durch Sprachassistenten wertvolle Daten über Konsumenten beziehen. - Keystone

Das Wichtigste in Kürze

- Millionen von Menschen benutzen Sprachassistenten bei sich zuhause.
- Was die Konzerne mit den Daten der Konsumenten machen, ist nicht bekannt.
- Deshalb sollen die Daten in Zukunft auf lokalen Routern gespeichert werden.

Daten könnten auch vor Ort bleiben

Es gibt längst Möglichkeiten, das Material der Nutzer auch vor Ort zu belassen. Eine Methode heisst Edge Computing – Rechenleistung am Rande (Edge) der Cloud. Die Daten werden dabei dezentral nah am Nutzer gesammelt und verarbeitet.

Möglich sei aber auch, dass solche Prozesse direkt auf der Hardware ablaufen, auf der die Sprachassistenten installiert sind. Eine Firma, die diesen Weg geht, ist das Start-up Snips mit Sitz in Paris und New York. Die Privatsphäre der Nutzer habe dabei oberste Priorität, betonen die Verantwortlichen.

«Es fühlt sich einfach falsch an, dass ein Konzern Zugriff auf ein Mikrofon hat, das in der eigenen Wohnung steht», meint Snips-Mitgründer Rand Hindi. Stimme und Hintergrundgeräusche bleiben deshalb an Ort und Stelle – ohne dass jemand ungewollt darauf zugreifen kann.

Nicht nur für den Heimeinsatz ist Datenverarbeitung jenseits der Cloud attraktiv. Auch die Industrie ist längst aufmerksam geworden.



Intelligent Edge

<https://www.techopedia.com/definition/32559/intelligent-edge>



Definition - What does *Intelligent Edge* mean?

Intelligent edge is a term describing a process where data is analyzed and aggregated in a spot close to where it is captured in a network. The intelligent edge, also described as “intelligence at the edge,” has important ramifications for distributed networks including the internet of things (IoT).





Intelligent Edge


<https://www.sony.com/electronics/cyber-shot-compact-cameras/dsc-rx100m7>

Real-time Tracking for moving subjects


Real-time Tracking uses subject-recognition algorithms (AI)⁸ that draw on color, pattern (brightness), subject distance (depth) and face/eye information to track your subjects wherever they move. Maintaining a constant, steady focus on a moving subject helps you capture every important moment of action as it happens, whether you're taking stills or movies.



Real-time Eye AF for both stills and movies



Real-time Eye AF for Animals



Intelligent Edge

<https://devblogs.nvidia.com/jetson-nano-ai-computing/>


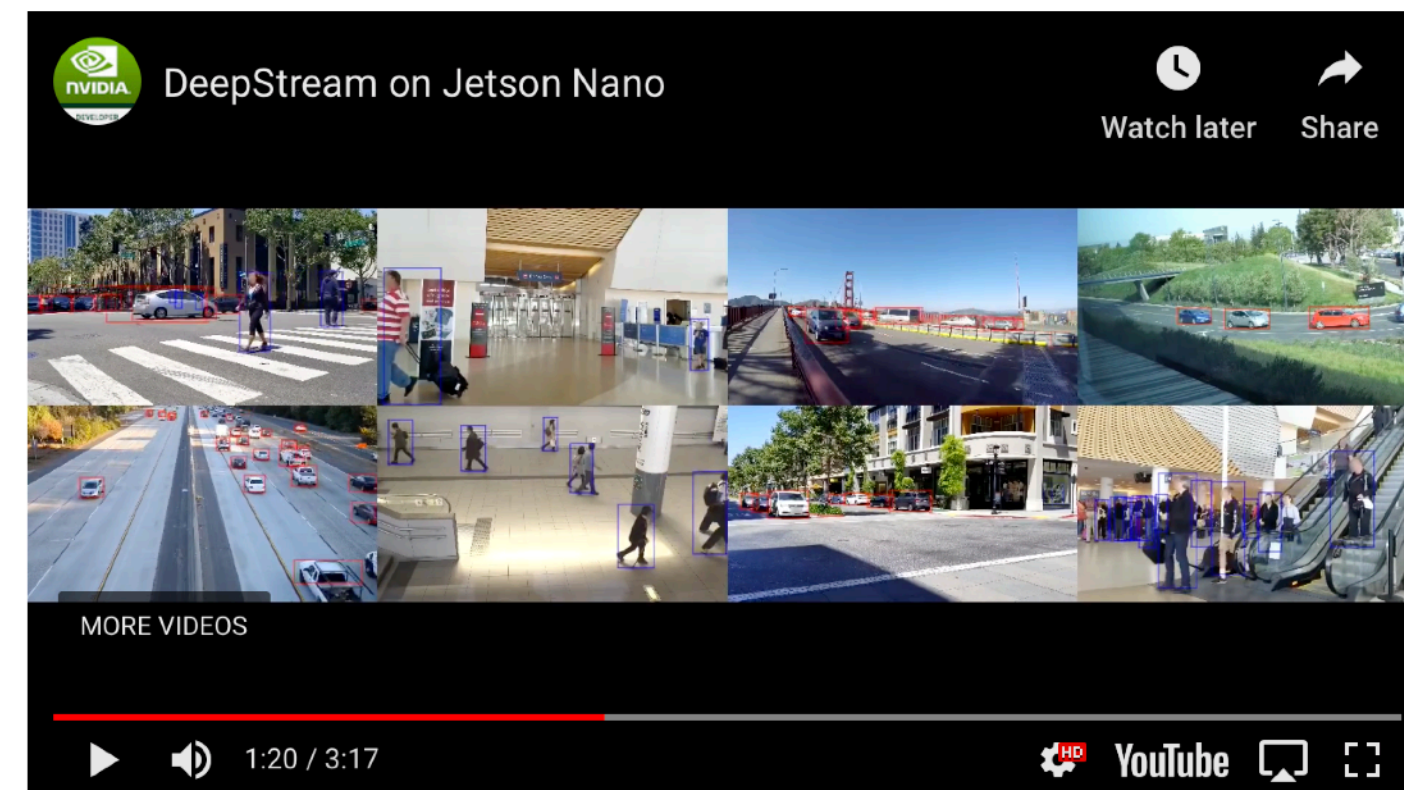


Figure 1. Jetson Nano Developer Kit (80x100mm), available now for \$99

Multi-Stream Video Analytics

Jetson Nano processes up to eight HD full-motion video streams in real-time and can be deployed as a low-power edge intelligent video analytics platform for Network Video Recorders (NVR), smart cameras, and IoT gateways. NVIDIA's [DeepStream SDK](#) optimizes the end-to-end inferencing pipeline with ZeroCopy and TensorRT to achieve ultimate performance at the edge and for on-premises servers. The video below shows Jetson Nano performing object detection on eight 1080p30 streams simultaneously with a ResNet-based model running at full resolution and a throughput of 500 megapixels per second (MP/s).



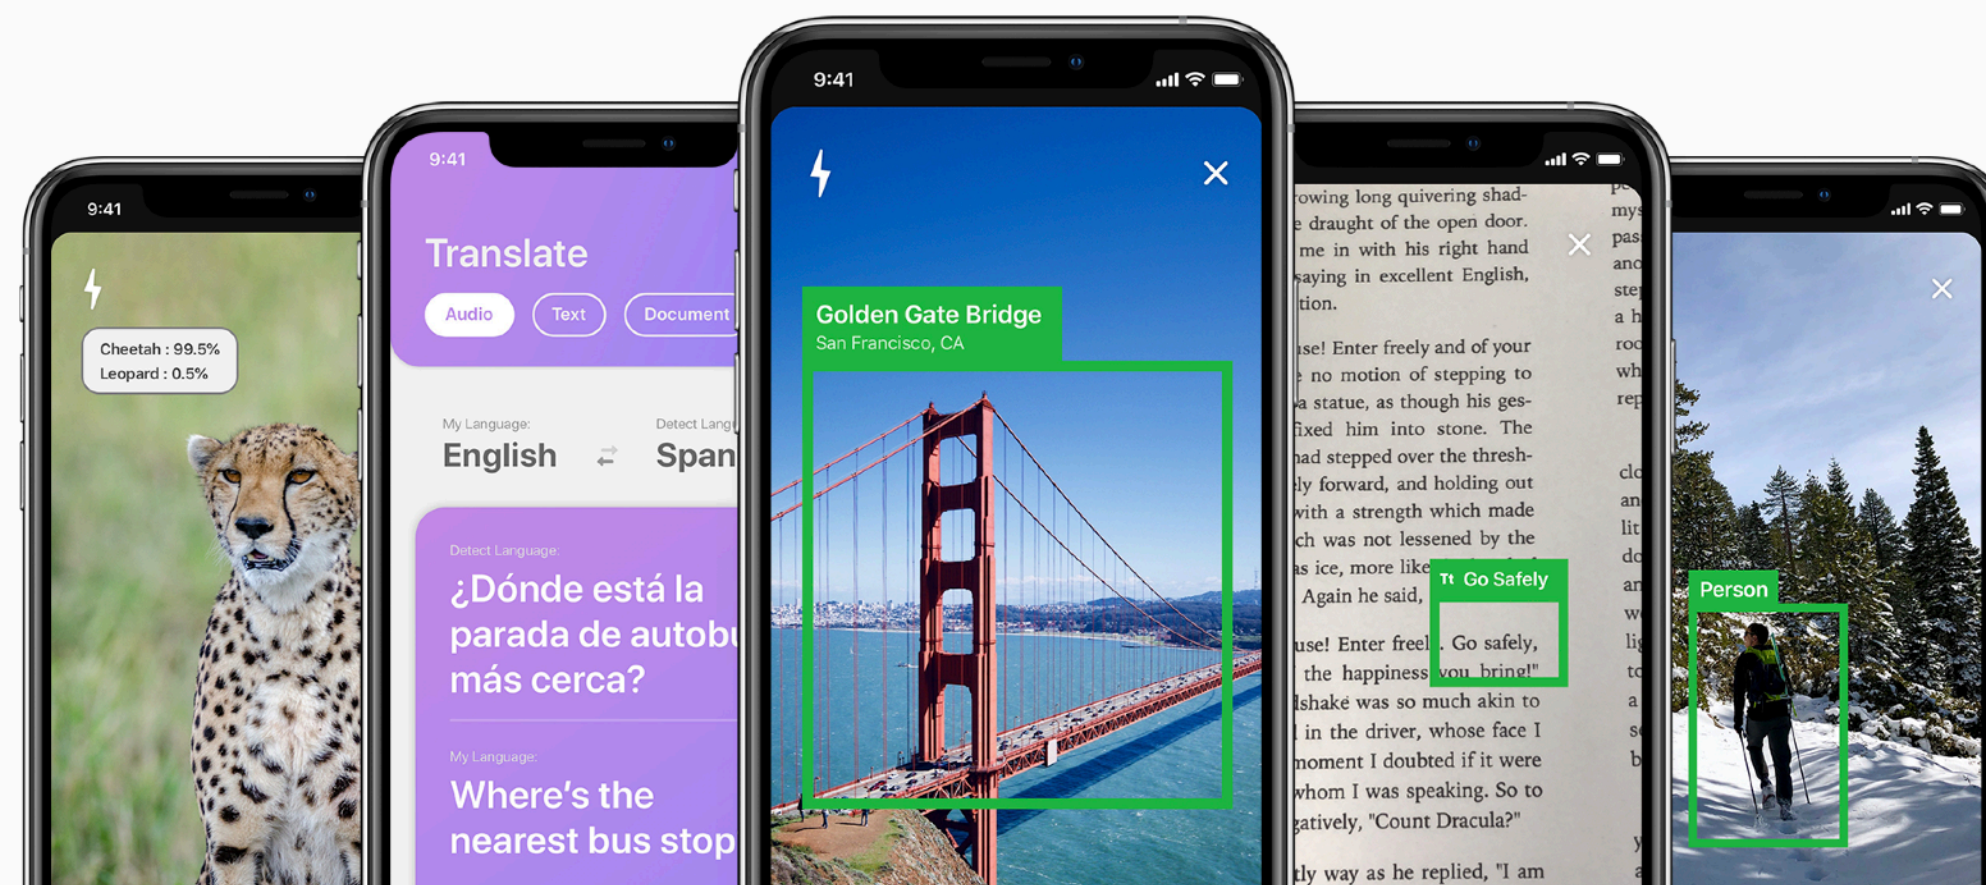
Intelligent Edge

<https://developer.apple.com/machine-learning/>

<https://www.apple.com/iphone-xs/a12-bionic/>

Machine Learning

Build realtime, personalized experiences with industry-leading, on-device machine learning using Core ML 3, Create ML, the powerful A-series chips, and the Neural Engine. Core ML 3 supports more advanced machine learning models than ever before. And with Create ML, you can now build machine learning models right on your Mac with zero code.



Apple-designed Neural Engine. Our Neural Engine is built for advanced, real-time machine learning. That means iPhone Xs can recognize patterns, make predictions, and learn from experience, similar to the way you do.

The Neural Engine is incredibly fast, able to perform five trillion operations per second. It's incredibly efficient, which enables it to do all kinds of new things in real time. And it's incredibly smart, so you can do things like jump right into immersive AR experiences. And quickly find all your pictures of dogs. Or beaches. Or babies.

We've also opened the Neural Engine up to the Core ML platform, so developers can bring powerful, real-time machine learning to their apps — learning that happens right on your iPhone.

8-core
architecture

5 trillion
operations per second



Intelligent Edge

<https://developer.ibm.com/articles/transfer-learning-for-deep-learning/>

Transfer learning for deep learning

Repurpose models for new problems with less data for training

by M. Tim Jones | Updated June 18, 2019 - Published June 20, 2019

Artificial intelligence Data science Deep learning Machine learning

In the early days of AI, a common problem was the lack of general intelligence. Models could be built to do things like play checkers, but the knowledge encapsulated in that model was restricted to that domain. This problem is being explored today under the name *transfer learning*— with the goal of building a model that can be applied to multiple related problem areas.

Transfer learning

Since the introduction of deep learning, there's been a renewed interest in neural networks for a range of applications. Deep learning has solved problems viewed as impossible not more than a decade ago. But deep learning neural networks require large clusters of compute servers, large amounts of training data, and a large amount of time to train the deep neural network.




Figure 1. The problem with deep learning from scratch

The deep neural network is represented by tens of millions of weights that connect the many layers of neurons of the networks together. These weights (typically real values) are adjusted during the training process and applied to inputs (including inputs from intermediary layers) to feed forward to an output classification. The basic idea of transfer learning is then to start with a deep learning network that is pre-initialized from training of a similar problem. Using this network, a smaller duration of training is required for the new, but related, problem.




Figure 2. Transfer learning with a pre-trained network

Transfer learning is the method of starting with a pre-trained model and training it for a new — related — problem domain. The pre-trained network serves as transferred knowledge to be applied in another domain. But there are numerous options that can be used, including feature transfer and fine-tuning (which depend upon the similarity of the problems at hand), in addition to freezing certain layers of the network and retraining others.



Intelligent Edge


<https://machinethink.net/blog/coreml-training-part1/>

On-device training with Core ML – part 1

19 JULY 2019 · 14 minutes

Machine learning on mobile gets more popular every year! WWDC 2019 gave us [lots of new goodies](#) for adding ML into our apps.

One of the biggest announcements was that Core ML 3 now supports **training of models on the iPhone and iPad**. Who would have thought a few years ago that we'd be training convnets on our handheld devices!



Why train on the device at all?

Doing training on the device is a good idea for a few reasons:

- You can make a model that is tailored to each user's own wishes or needs.
- To be fair, that's also something you can do in the cloud, but it will require running expensive servers and sending the user's data across the net.
- Because training happens on-device, the user's data never has to leave the device, protecting the user's privacy and making government regulators happy.

You need a trained model to start with

In order to fine-tune a model, you already need to have a trained model.

Machine learning models start out life not knowing anything. An untrained model's brain — the *learned parameters* or weights — is made up of randomly chosen numbers. If you ask such an untrained model for a prediction, it will just make random guesses. To get sensible predictions, you need to train the model first.



Often, your users will have data that is very **similar in the broad strokes**, but that differs in the details. In that case, it makes sense to use on-device personalization to adapt the model to each user's specific usage.

To enable this, you need to provide a trained model that **already understands the data in general terms**. This is the baseline model that gets shipped with the app. On-device training can then be used to make the model learn new things about just this user and their data.



Intelligent Edge

<https://aiyprojects.withgoogle.com/vision>



TENSORFLOW MODEL COMPILER

To further customize your project, you can build a customized [TensorFlow](#) model that recognizes new types of objects. You can (re)train a TensorFlow model and use our [compiler](#) to convert the model into binary file to run on Vision Kit.

Due to limited hardware resources on Vision Bonnet, there are [constraints](#) on what type of models can run on device. We have tested and verified that the following model structures are supported on Vision Kit.



MODEL TYPE	SUPPORTED CONFIGURATION
MobileNetV1	input size: 160x160, depth multiplier = 0.5 input size: 192x192, depth multiplier = 1.0
MobileNetV1 + SSD	input size: 256x256, depth multiplier = 0.125
SqueezeNet	input size: 160x160, depth multiplier = 0.75

- Identify all kinds of plant and animal species
- See when your dog is at the back door
- See when your car left the driveway
- See that your guests are delighted by your holiday decorations
- See when your little brother comes into your room (sound the alarm!)



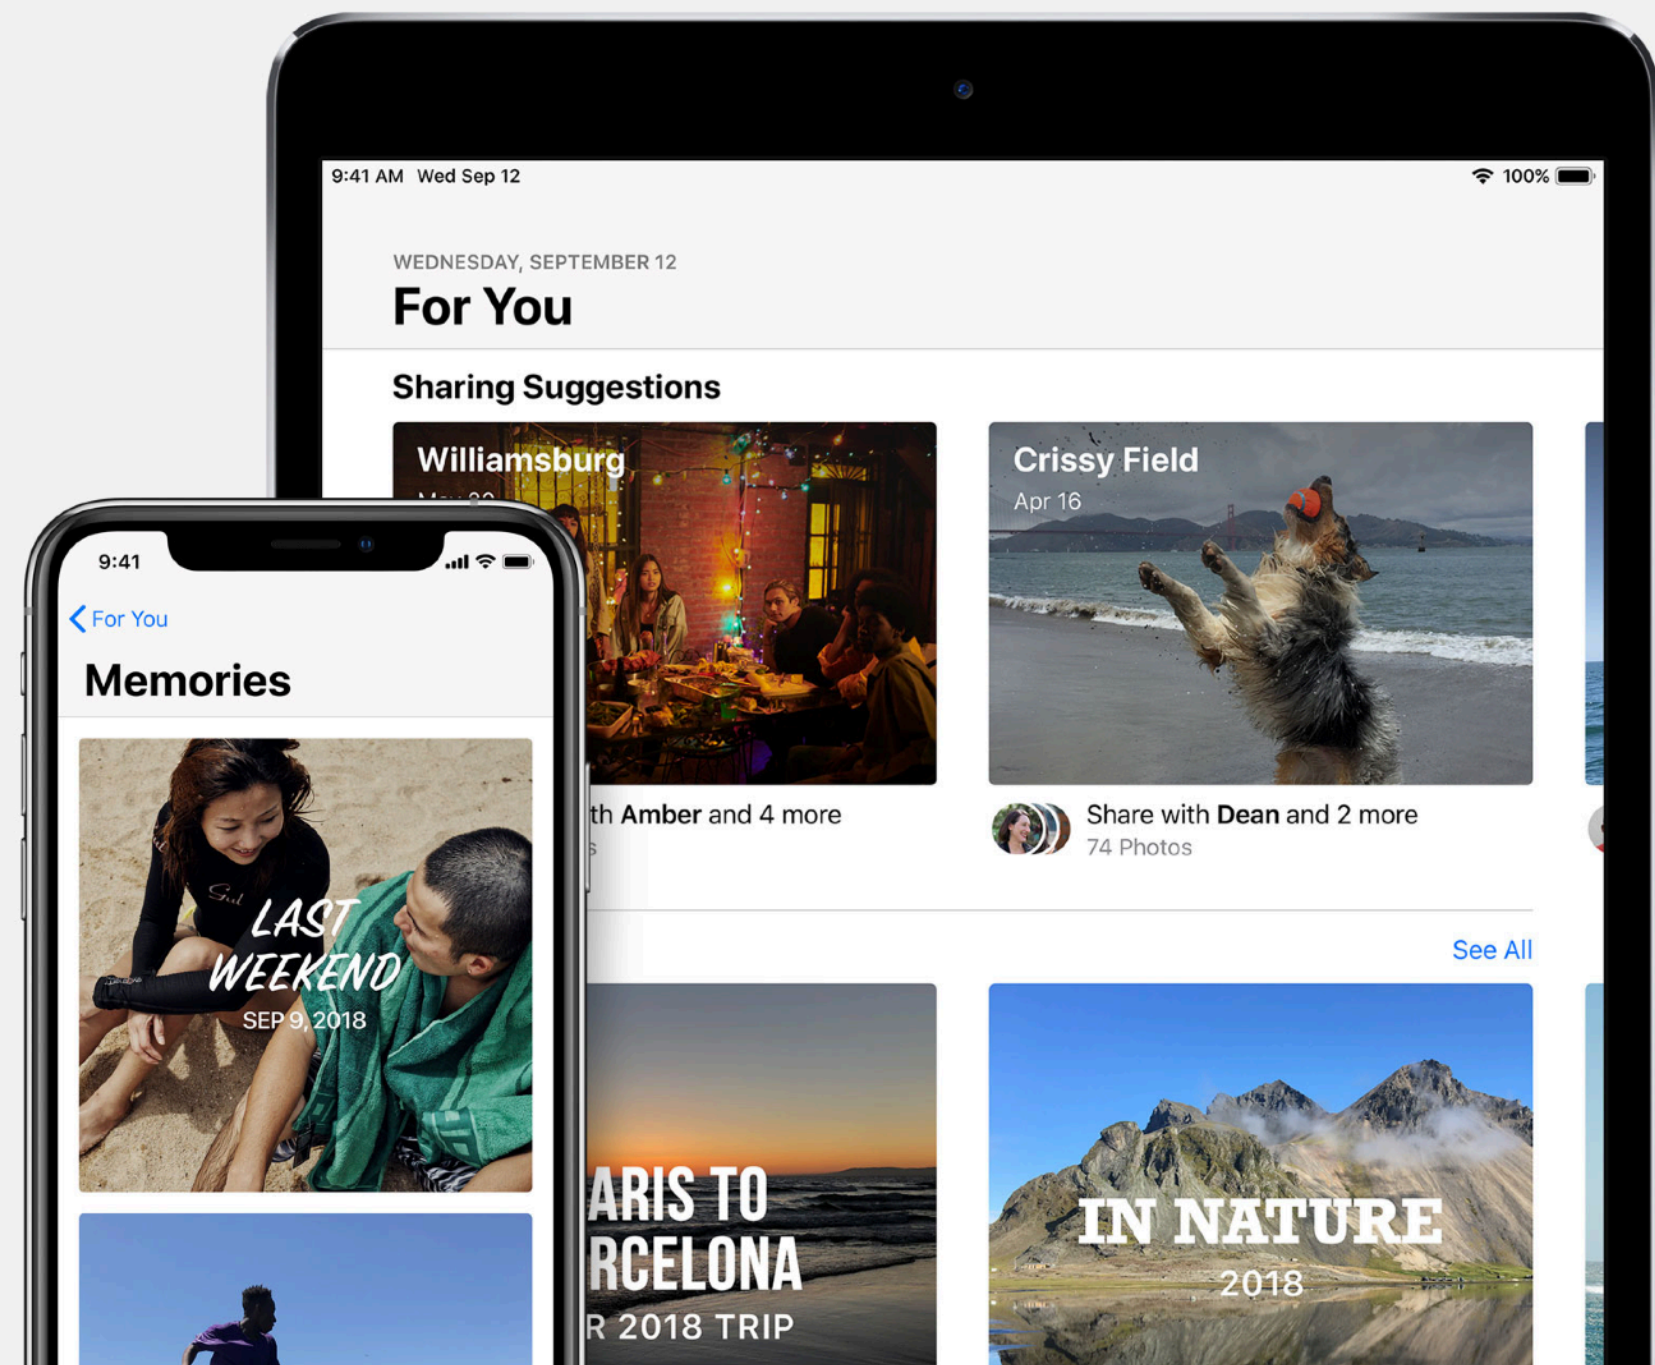
Intelligent Edge

<https://www.kickstarter.com/projects/roboLink/driving-into-the-world-of-ai-zumi/description>



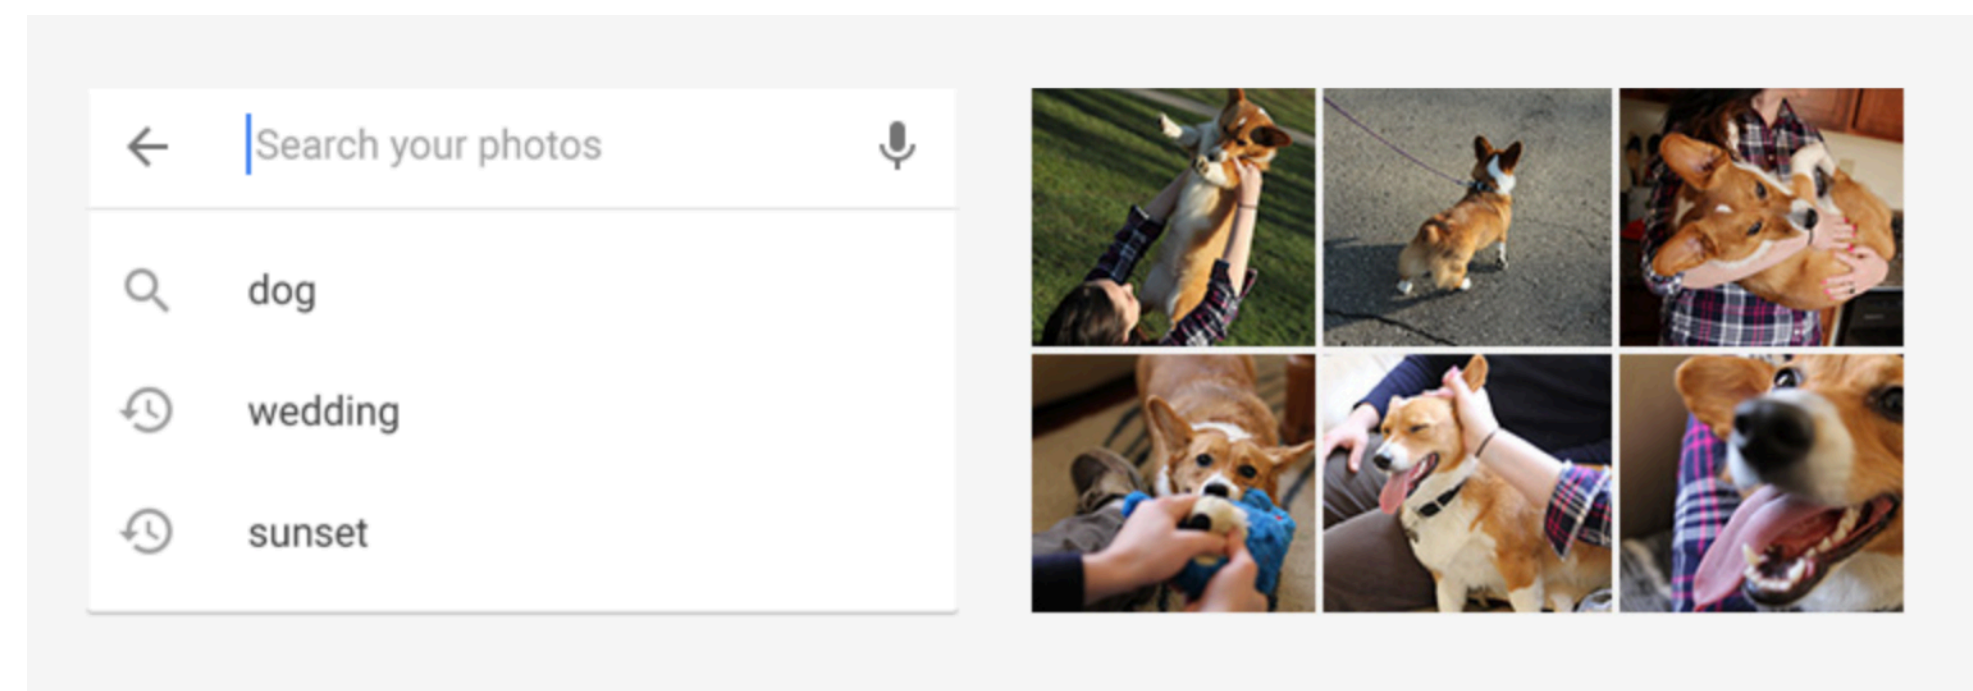
Intelligent Edge versus Intelligent Cloud

Photos for iOS is where all the amazing photos you take with your iPhone and iPad live. And now it's even smarter and more personal, with intelligent new features that help you find, edit, and share your best shots.

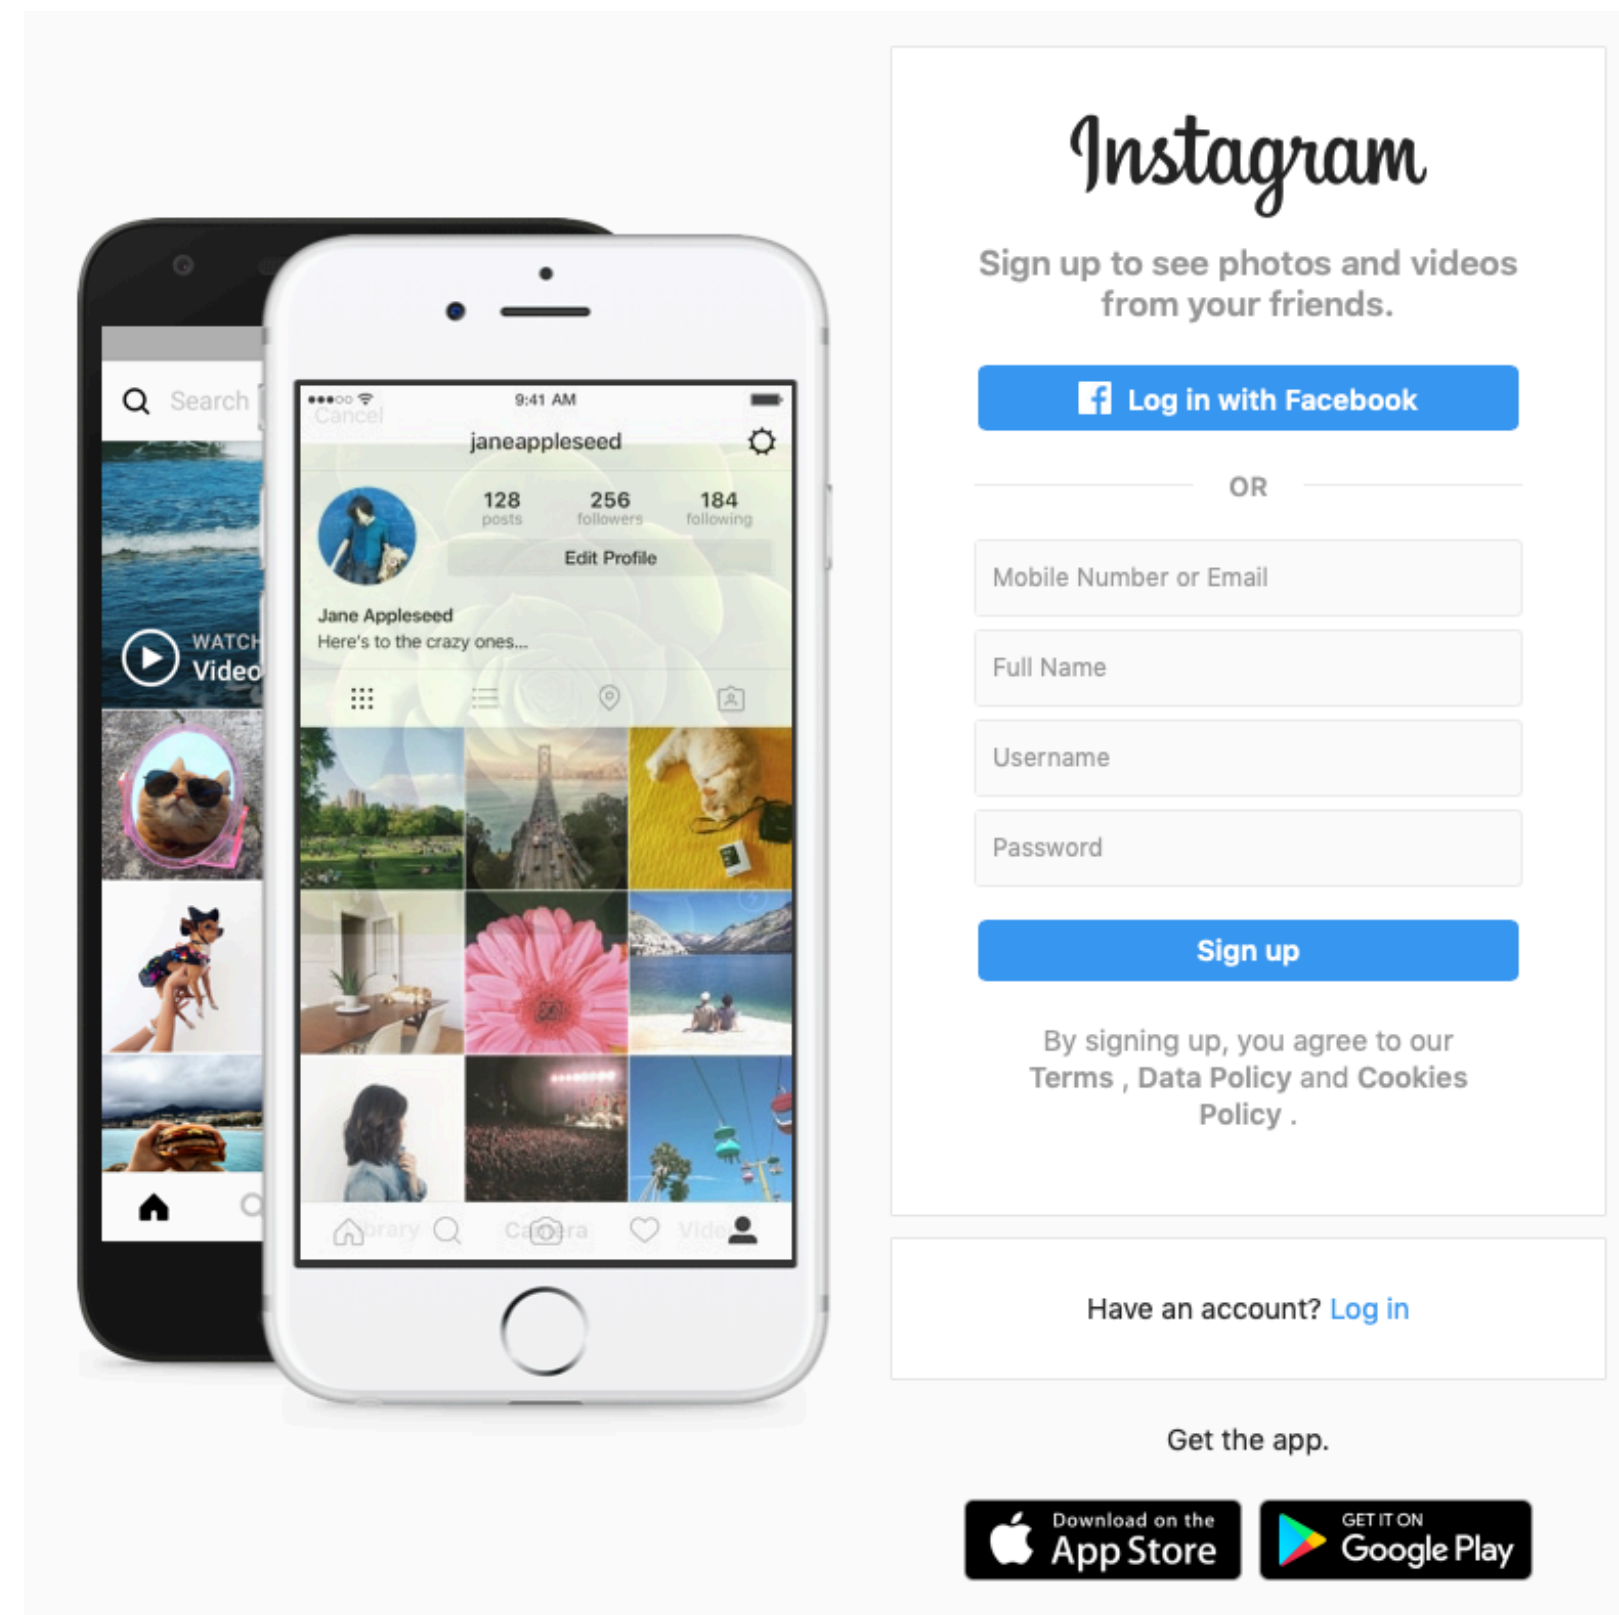


Google Photos

Free storage and automatic organization
for all your memories.

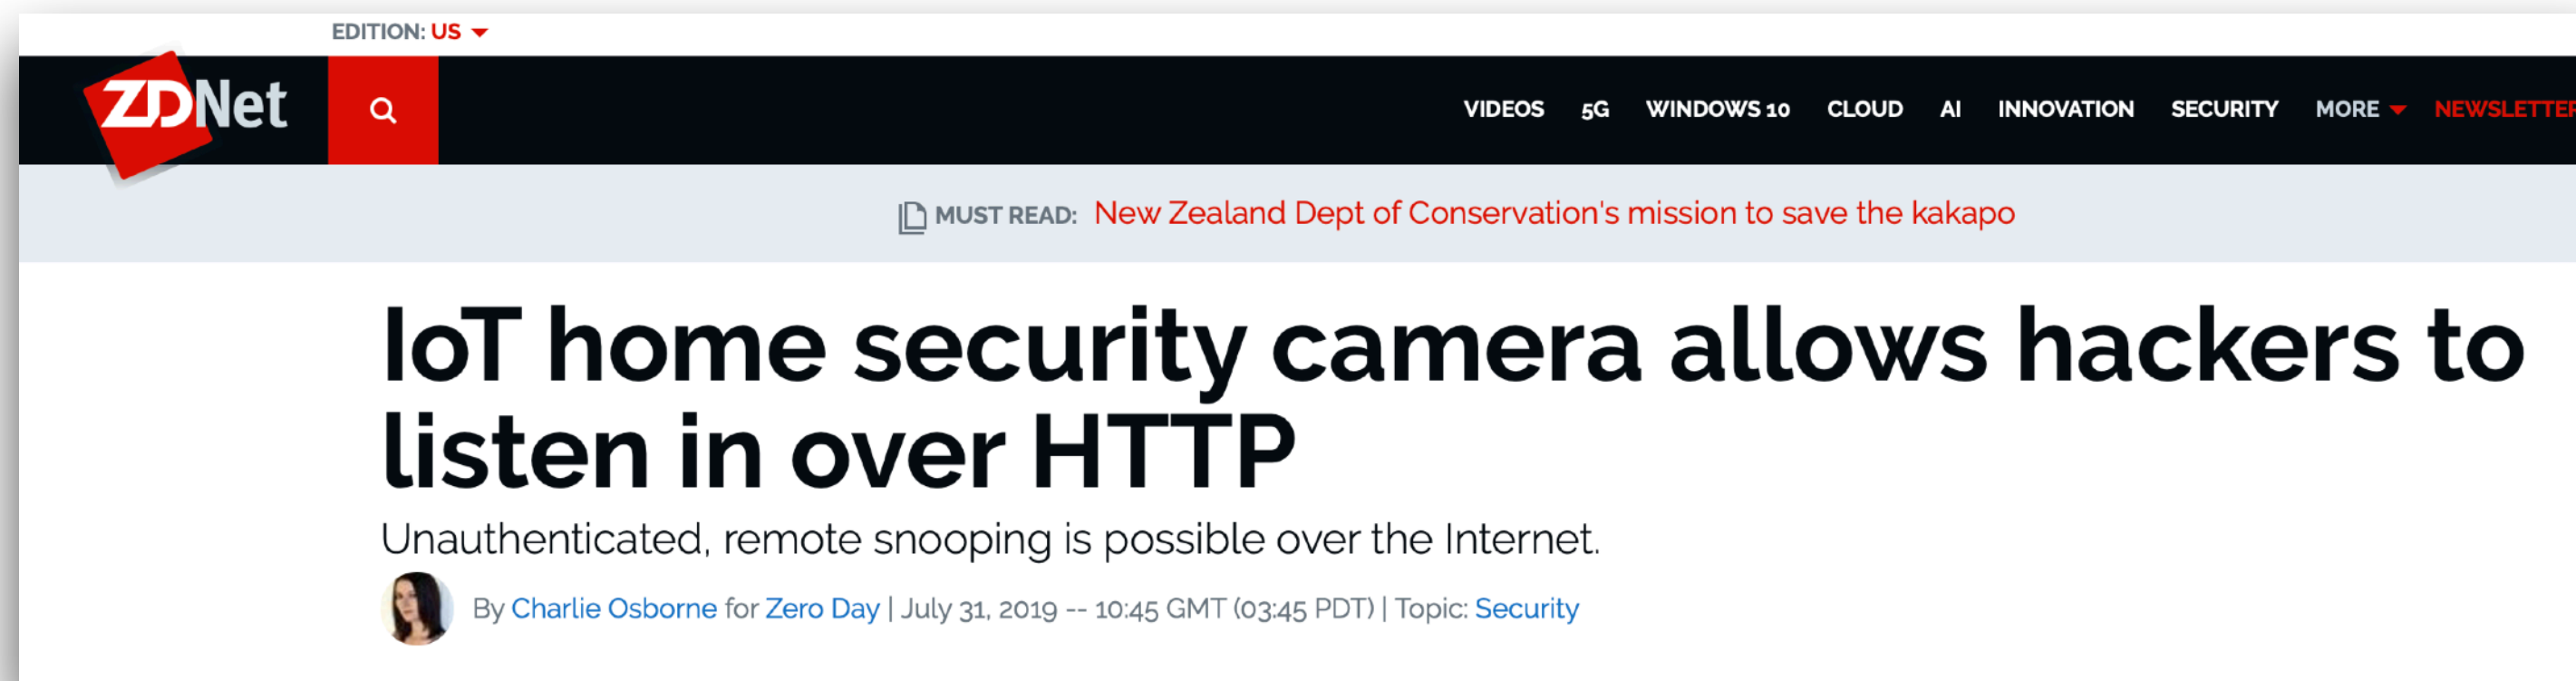


"Intelligent" Cloud




"Intelligent" Cloud

<https://www.zdnet.com/article/iot-home-security-camera-allows-hackers-to-listen-in-over-http/>



Security researchers have uncovered a security flaw in a popular home security camera which permits remote spying without any form of authentication.

This week, researchers from cybersecurity firm Tenable said the Amcrest IP2M-841B IP camera, [available on Amazon](#) and subject to 12,000 customer reviews -- many of which are positive -- contained a serious bug which is "trivial" to exploit.



"Intelligent" Cloud

<https://www.theverge.com/2019/8/14/20802774/june-smart-oven-remote-preheat-update-user-error>

<https://juneoven.com>

Smart ovens have been turning on overnight and preheating to 400 degrees

June Oven owners are reporting preheating incidents

By Ashley Carman | @ashleyrcarman | Aug 14, 2019, 2:54pm EDT

f t SHARE






Photo by Amelia Holowaty Krales / The Verge


GOOD DEALS



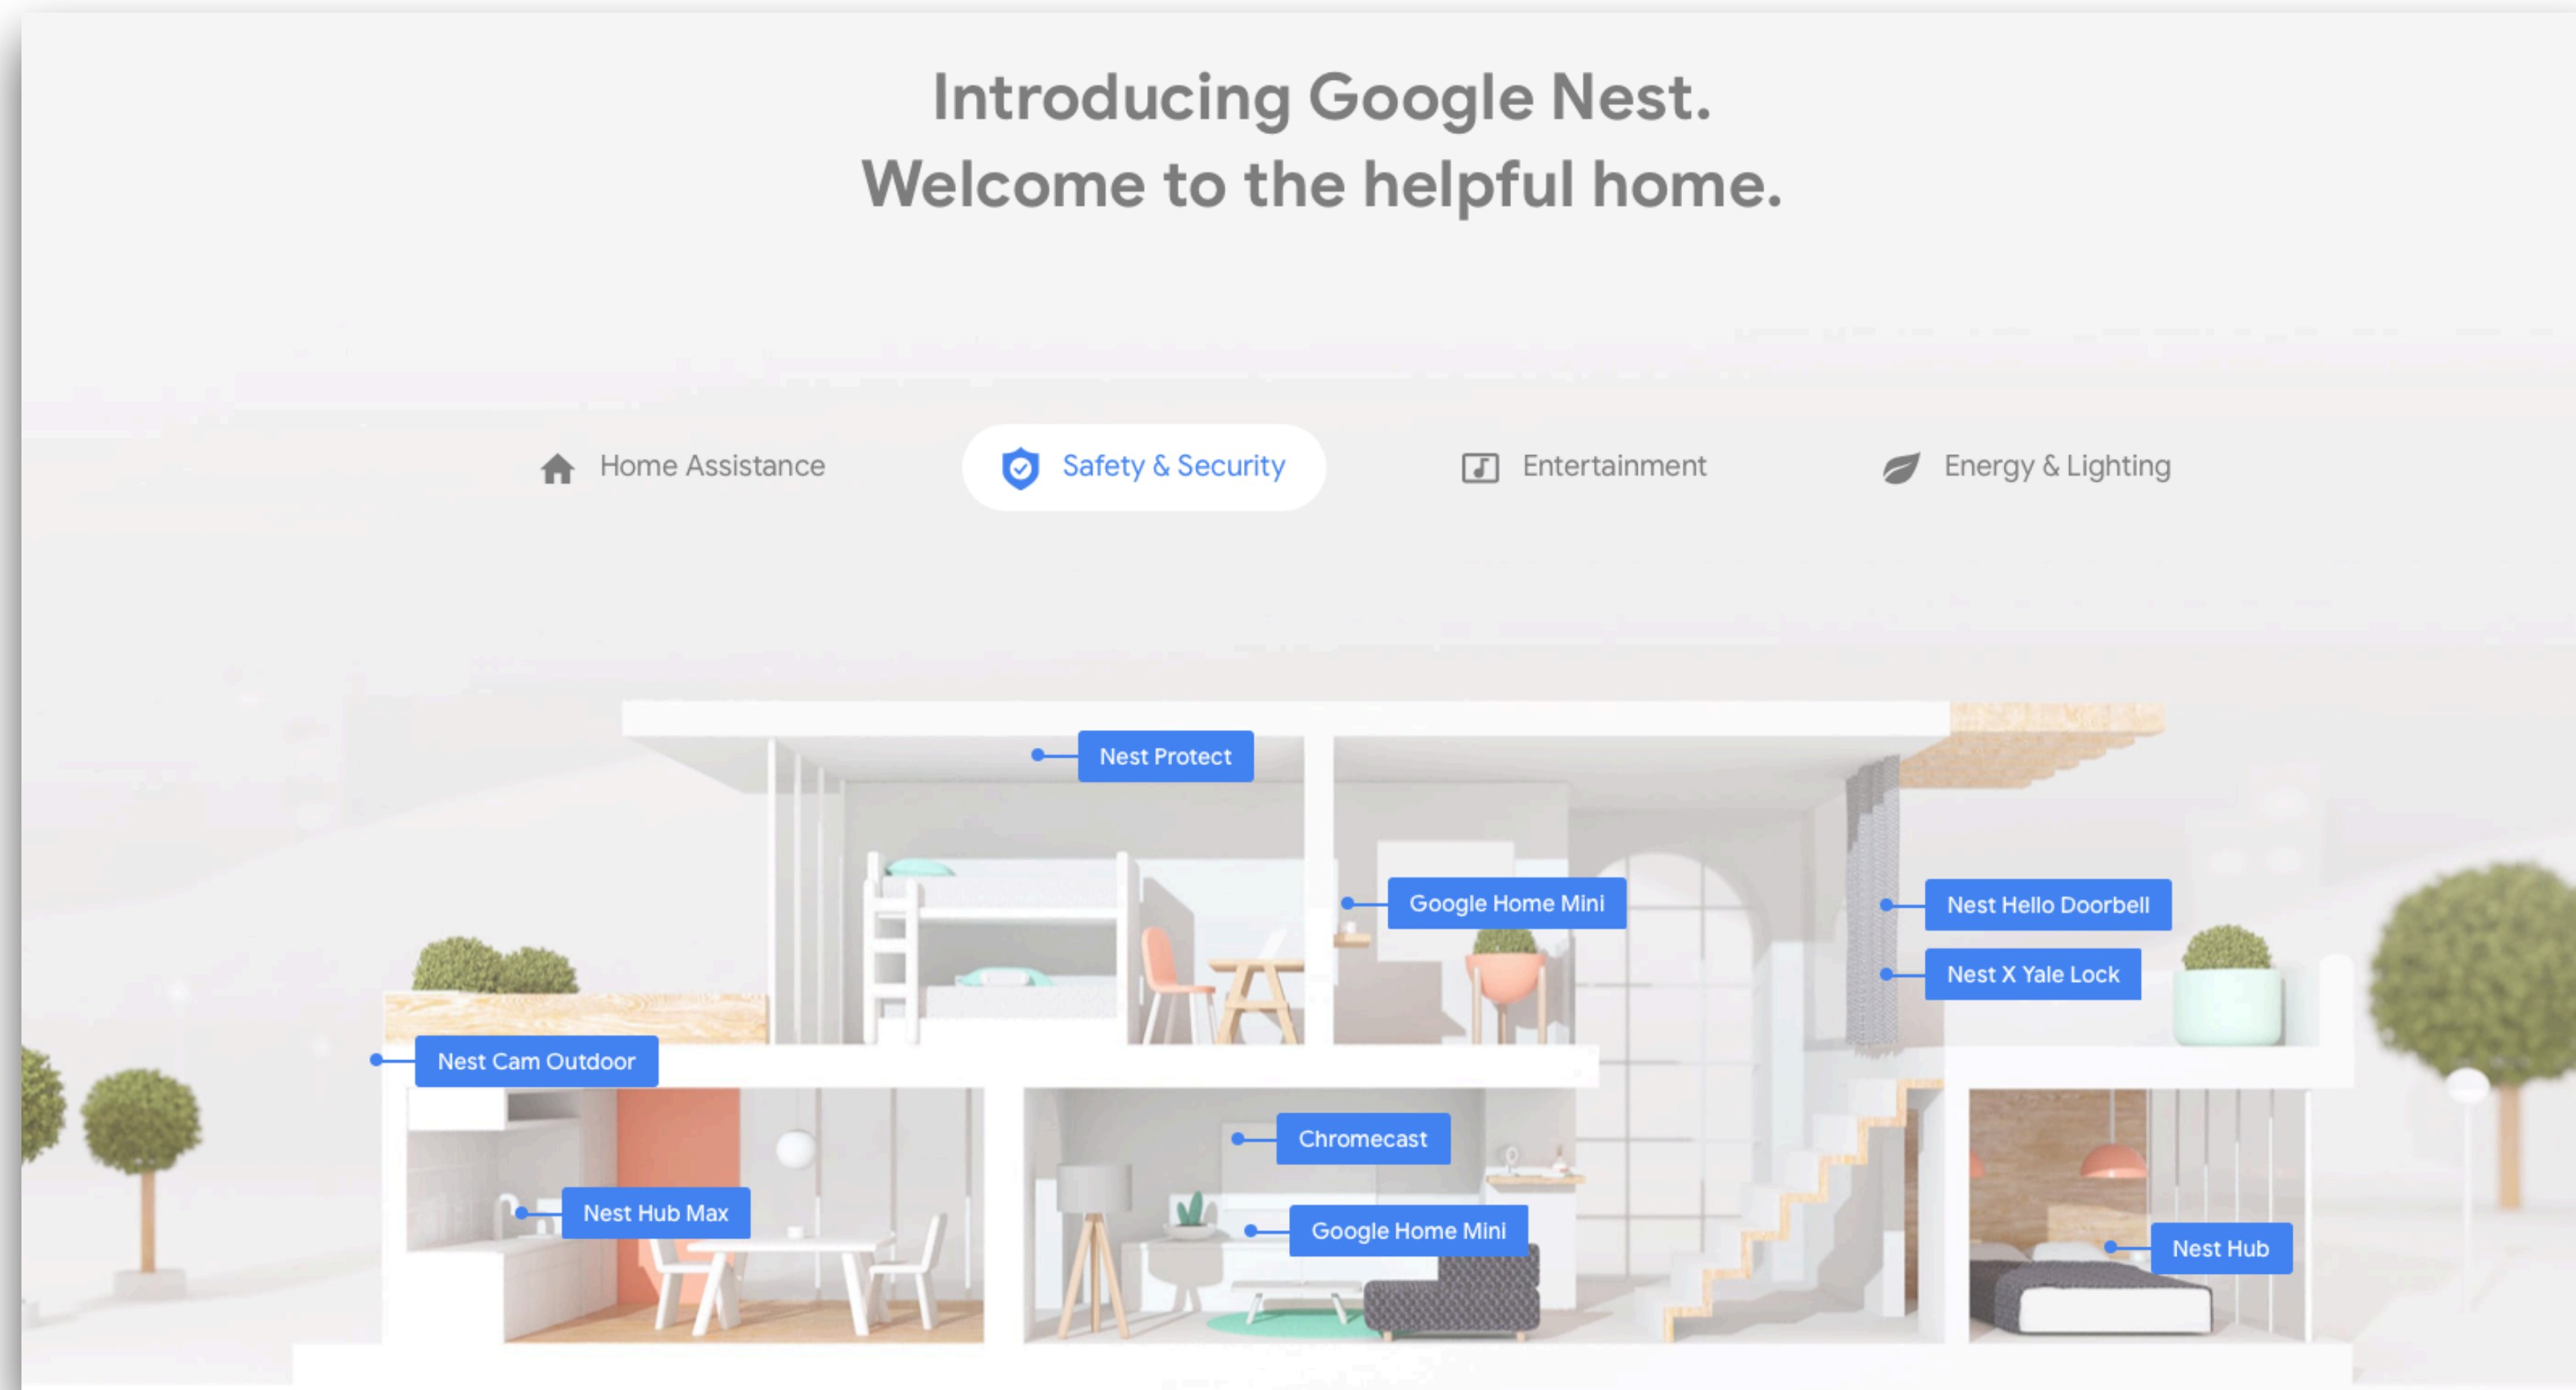
Control from anywhere

Use your phone to cook with your oven, watch live video of your food, check ETA and be notified when your meal is ready.

At least three smart June Ovens have turned on in the middle of the night and heated up to 400 degrees Fahrenheit or higher. The ovens' owners aren't sure why this happened, and June tells *The Verge* that user error is at fault. The company is planning an update that'll hopefully remedy the situation and prevent it from happening again, but that change isn't coming until next month.



"Intelligent" Cloud



Zusammenfassung

Es ist vorbei mit der Privatsphäre im Internet.

Verteidigung der Privatsphäre im Hause.

Daten Lokal halten.

Intelligence/Compute at the Edge - Encrypted Backup in the Cloud.

... Stop Feeding the AD/AI Giants (Facebook, Google) for Free.