

REFERENCES AND NOTES

- D. Musso, D. J. Gubler, *Clin. Microbiol. Rev.* **29**, 487–524 (2016).
- S. B. Halstead, *Science* **239**, 476–481 (1988).
- S. C. Kliks, S. Nimmamitya, A. Nisalak, D. S. Burke, *Am. J. Trop. Med. Hyg.* **38**, 411–419 (1988).
- D. W. Vaughn et al., *J. Infect. Dis.* **181**, 2–9 (2000).
- C. P. Simmons, J. J. Farrar, V. Nguyen, B. Wills, *N. Engl. J. Med.* **366**, 1423–1432 (2012).
- L. Dai et al., *IUBMB Life* **68**, 783–791 (2016).
- J. T. Roehrig, *Adv. Virus Res.* **59**, 141–175 (2003).
- L. Dai et al., *Cell Host Microbe* **19**, 696–704 (2016).
- D. Sirohi et al., *Science* **352**, 467–470 (2016).
- L. Priyamvada et al., *Proc. Natl. Acad. Sci. U.S.A.* **113**, 7852–7857 (2016).
- K. Stettler et al., *Science* **353**, 823–826 (2016).
- J. A. Swanstrom et al., *MBio* **7**, e01123-16 (2016).
- W. Dejnirattisai et al., *Nat. Immunol.* **17**, 1102–1108 (2016).
- L. M. Paul et al., *Clin. Transl. Immunol.* **5**, e117 (2016).
- L. R. Petersen et al., *Epidemiol. Infect.* **141**, 591–595 (2013).
- C. Chancey, A. Grinev, E. Volkova, M. Rios, *BioMed Res. Int.* **2015**, 376230 (2015).
- S. L. Stramer et al., *N. Engl. J. Med.* **353**, 451–459 (2005).
- S. L. Orton, S. L. Stramer, R. Y. Dodd, *Transfusion* **46**, 272–277 (2006).
- S. J. Balsitis et al., *PLOS Pathog.* **6**, e1000790 (2010).
- S. Kliks, *AIDS Res. Hum. Retroviruses* **6**, 993–998 (1990).
- B. J. Mady, I. Kurane, D. V. Erbe, M. W. Fanger, F. A. Ennis, *J. Gen. Virol.* **74**, 839–844 (1993).
- B. Guy et al., *Vaccine* **22**, 3563–3574 (2004).
- R. Jefferis, J. Lund, *Immunol. Lett.* **82**, 57–65 (2002).
- D. Kao et al., *Cell Rep.* **13**, 2376–2385 (2015).
- W. J. Liu et al., *J. Virol.* **79**, 1934–1942 (2005).
- A. Grant et al., *Cell Host Microbe* **19**, 882–890 (2016).
- J. Ashour, M. Laurent-Rolle, P. Y. Shi, A. Garcia-Sastre, *J. Virol.* **83**, 5408–5418 (2009).
- M. Mazzon, M. Jones, A. Davidson, B. Chain, M. Jacobs, *J. Infect. Dis.* **200**, 1261–1270 (2009).
- J. Ashour et al., *Cell Host Microbe* **8**, 410–421 (2010).
- S. Tripathi et al., *PLOS Pathog.* **13**, e1006258 (2017).
- G. W. Dick, S. F. Kitchen, A. J. Haddow, *Trans. R. Soc. Trop. Med. Hyg.* **46**, 509–520 (1952).
- D. H. Libraty et al., *J. Infect. Dis.* **185**, 1213–1221 (2002).
- W. K. Wang et al., *Virology* **305**, 330–338 (2003).
- E. B. Hayes, D. J. Gubler, *Annu. Rev. Med.* **57**, 181–194 (2006).
- M. Taketa-Graham et al., *Am. J. Trop. Med. Hyg.* **82**, 501–504 (2010).
- V. Vorndam, M. Beltran, *Am. J. Trop. Med. Hyg.* **66**, 208–212 (2002).
- T. Ishikawa, A. Yamanaka, E. Konishi, *Vaccine* **32**, 1326–1337 (2014).
- M. W. Gaunt et al., *J. Gen. Virol.* **82**, 1867–1876 (2001).
- S. S. Baba, A. H. Fagbami, O. D. Olaleye, *Rev. Inst. Med. Trop. Sao Paulo* **40**, 343–349 (1998).

ACKNOWLEDGMENTS

We thank R. Cadagan, O. Lizardo, and F. Amanat for excellent technical assistance; M. Evans's laboratory for providing the RNA standards for determination of ZIKV load by Taqman PCR assay; C. Schindler for providing the original *Stat2*^{-/-} mouse strain; A. Fernandez-Sesma for ZIKV and DENV strains; D. Bogunovic for invaluable evaluation of our manuscript; the microscopy core at Icahn School of Medicine at Mount Sinai for their support and technical assistance; and G. Estrella and J. Bethea at the Center for Comparative Medicine at the Icahn School of Medicine at Mount Sinai for their efforts to rapidly accommodate ZIKV studies in our vivarium. The following reagents were obtained through BEI Resources, NIAID (National Institute of Allergy and Infectious Diseases), NIH: Pre-Immune Dengue Virus Type 1 Sera, NR-41782; Pre-Immune Dengue Virus Type 2 Sera NR-41783; Pre-Immune Dengue Virus Type 3 Sera NR-41784; Pre-Immune Dengue Virus Type 4 Sera NR-41785; Early-Immune Dengue Virus Type 1 Antiserum NR-30247; Early-Immune Dengue Virus Type 2 Antiserum NR-29321; Early-Immune Dengue Virus Type 3 Antiserum NR-29323; Early-Immune Dengue Virus Type 4 Antiserum NR-29327; Late-Immune Dengue Virus Type 1 Antiserum NR-41786; Late-Immune Dengue Virus Type 2 Antiserum NR-41787; Late-Immune Dengue Virus Type 3 Antiserum NR-41788; Late-Immune Dengue Virus Type 4

Antiserum NR-41789; Pre-Immune Yellow Fever Virus Sera NR-42556, NR-42564, and NR-42565; Early-Immune Yellow Fever Virus Antiserum NR-29335, NR-29337, and NR-29338; and Late-Immune Yellow Fever Virus Antiserum NR-42567, NR-42575, and NR-42576. This work was partially supported by a supplement to NIAID grant U19AI118610 and NIAID grant R21AI130299. Human sample institutional review board protocols include Nucleic Acid Testing (NAT) for West Nile Virus (WNV) Protocol #2003-011 and Dengue Virus (Gen-Probe Procelex) Clinical Protocol #2012-016. All the data supporting our conclusions are contained in this manuscript.

SUPPLEMENTARY MATERIALS

www.sciencemag.org/content/356/6334/175/suppl/DC1
Materials and Methods
Figs. S1 to S12
Table S1
References (40–44)

18 November 2016; accepted 17 March 2017
Published online 30 March 2017
10.1126/science.aal4365

DEFAUNATION

The impact of hunting on tropical mammal and bird populations

A. Benítez-López,^{1*} R. Alkemade,^{2,3} A. M. Schipper,² D. J. Ingram,⁴ P. A. Verweij,⁵ J. A. J. Eikelboom,^{2,5,6} M. A. J. Huijbregts^{1,2}

Hunting is a major driver of biodiversity loss, but a systematic large-scale estimate of hunting-induced defaunation is lacking. We synthesized 176 studies to quantify hunting-induced declines of mammal and bird populations across the tropics. Bird and mammal abundances declined by 58% (25 to 76%) and by 83% (72 to 90%) in hunted compared with unhunted areas. Bird and mammal populations were depleted within 7 and 40 kilometers from hunters' access points (roads and settlements). Additionally, hunting pressure was higher in areas with better accessibility to major towns where wild meat could be traded. Mammal population densities were lower outside protected areas, particularly because of commercial hunting. Strategies to sustainably manage wild meat hunting in both protected and unprotected tropical ecosystems are urgently needed to avoid further defaunation.

Global biodiversity loss is occurring at an unprecedented rate (1). Few undisturbed areas remain in the tropics (2), but these are threatened by escalating road and infrastructure expansion, which promotes human accessibility to otherwise remote areas and facilitates illegal colonization and hunting (3–5). Hunting exerts a major pressure on wildlife, which can result in large population declines and local extirpations of wildlife populations in forests that appear structurally undisturbed (6). Overhunted “half-empty” or “empty ecosystems” are becoming common across the tropics (7). Indeed, the abundance of wildlife in natural ecosystems is more closely related to patterns of hunting than to factors such as forest type, habitat area, or habitat protection status (8). A growing body of research is focusing on defaunation and its far-reaching cascading effects, including disruptions in seed dispersal mutualisms and a decline

in total biomass (9, 10). However, hunting-induced defaunation is a cryptic phenomenon that is difficult to monitor and, to date, no large-scale estimates of the impact of hunting on wildlife abundances are available.

Here, we analyze the impact of hunting on bird and mammal populations at a pantropical scale, in terms of both magnitude (decline in abundance) and spatial extent (depletion distances). We collated 176 studies, including 384 and 1938 effect sizes for 97 bird and 254 mammal species, respectively (11) (Fig. 1), and estimated the overall reduction in mammal and bird abundance in hunted compared with unhunted sites with a mixed effects meta-analysis. As an effect size, we calculated response ratios (*RR*) between the abundance of each species in hunted (X_h) and unhunted sites (X_u) within each study [$RR = \log(X_h/X_u)$; (12)]. *RR* are therefore negative ($RR < 0$) or positive ($RR > 0$) if abundance estimates are lower or higher, respectively, because of hunting pressure. Based on the central-place foraging hypothesis, hunting intensity is generally higher in the proximity of hunters' access points (e.g., settlements and roads) (5, 10), generating gradients of increasing species densities up to a distance where no effect is observed (i.e., species depletion distances). We used single meta-regression models to estimate species-depletion distances and to quantify how the impact of hunting varied depending on accessibility to urban markets for trade [travel time to major towns (13)], region, type of hunting (commercial versus subsistence versus

¹Department of Environmental Science, Institute for Wetland and Water Research, Radboud University, Post Office Box 9010, NL-6500 GL, Nijmegen, Netherlands. ²PBL Netherlands Environmental Assessment Agency, Post Office Box 30314, 2500 GH The Hague, Netherlands. ³Environmental Systems Analysis Group, Wageningen University, Post Office Box 47, 6700 AA Wageningen, Netherlands. ⁴School of Life Sciences, University of Sussex, Falmer, Brighton, BN1 9QG, UK. ⁵Energy and Resources, Copernicus Institute of Sustainable Development, Utrecht University, Heidelberglaan 2, 3584 CS Utrecht, Netherlands. ⁶Resource Ecology Group, Wageningen University and Research, Droevendaalsesteeg 3a, 6708 PB Wageningen, Netherlands.

*Corresponding author. Email: a.benitez@science.ru.nl

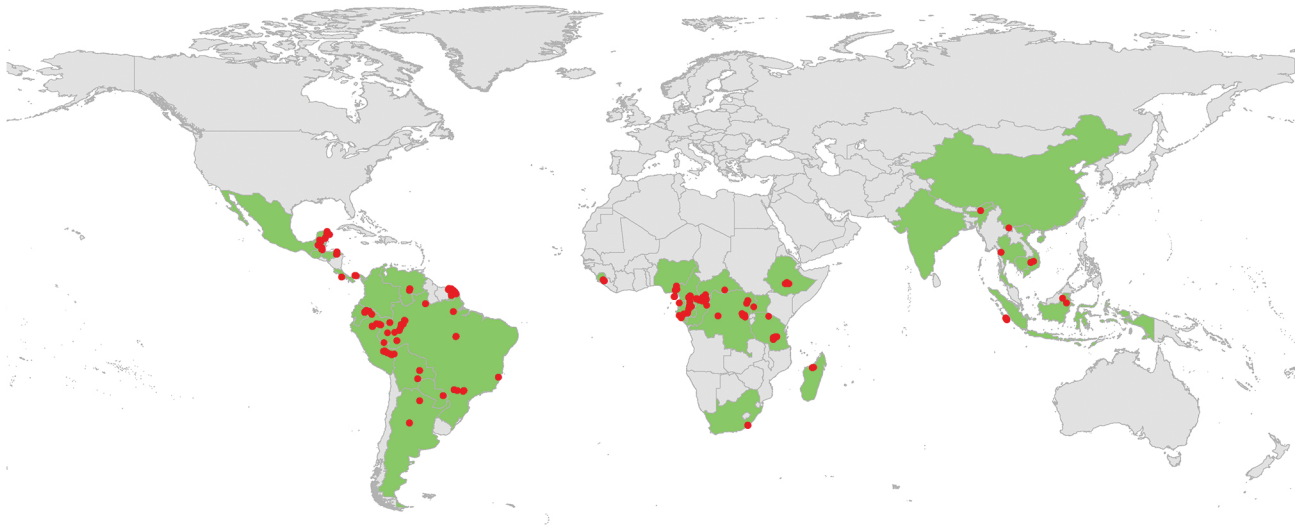


Fig. 1. Geographical location of the 176 studies included in the meta-analysis. Locations as red dots. Countries that contain at least one study are in green color. Red dots may represent multiple effect sizes.

both), protection status (protected versus unprotected area), species body size, and feeding guild. Finally, we tested the relative importance of these moderators using an information-theoretic approach of several multiple meta-regression models including first- and second-order interactions.

Overall, bird and mammal abundances were reduced by 58% [95% confidence interval (CI): 25, 76%] and 83% (95% CI: 72, 90%), respectively, in hunted areas (Fig. 2). Hunting pressure had a larger effect on mammals than on birds, probably because hunters preferentially target larger species (6). Results were robust to potential publication bias for mammals and to Geary diagnostic tests and differences in study quality for both groups (figs. S3 and S4). Hunting-induced abundance reductions varied with distance to hunters' access points (distance, hereafter), accessibility to urban markets, protected area status and type of hunting, with distance being the most important moderator (Figs. 3 and 4 and table S5). For birds, effect sizes were the lowest in proximity to hunters' access points ($RR_b = -3.17$, 95% CI = $-2.62, -3.71$, ~95% loss at 500 m) and approximated 0 at a distance of 7 km (Fig. 3A). For mammals, effect sizes first decreased from -0.76 ($-1.30, -0.23$) to -2.38 ($-2.84, -1.78$) within the first 700 m (~90% loss), and then increased steadily up to 0 at ~40 km from hunters' access points (Fig. 3B). This initial higher RR may reflect the replacement of large-bodied mammals by smaller ones. Indeed, we found evidence of size-differential mammal defaunation for frugivores, carnivores, herbivores, and insectivores (tables S6 and S7). Smaller mammals were consistently more abundant at higher hunting pressure than larger species (fig. S5), probably owing to release from predation pressure and competition as a result of (near) extirpation of medium- and large-sized mammals (14). Large-bodied frugivores, herbivores, and insectivores—including chimpanzees (*Pan troglodytes*), West-

ern gorillas (*Gorilla gorilla*), and giant armadillos (*Priodontes maximus*)—are largely hunted for wild meat consumption and trade (15). In turn, large carnivores, such as leopards (*Panthera pardus*) and jaguars (*Panthera onca*), are often persecuted because of livestock-wildlife conflicts, or their populations are reduced because of hunting-induced losses of prey species (16).

Bird and mammal population abundances were lower in hunted areas with higher accessibility to urban markets (Fig. 3, C and D). Effect sizes approached 0 within 1 to 2 days of travel time from the nearest major town. For mammals, this effect remained after controlling for other factors (table S6). Across the tropics, the majority of consumed and traded wild meat and body parts comes from mammals, whereas birds are generally killed for a hunter's own consumption (6, 17). However, for both species groups, the transition from subsistence to commercial hunting is having a massive impact on population densities (Fig. 4). Current prospects of infrastructure expansion in the Amazon, Africa, and Asia will facilitate accessibility to remote areas (3, 18, 19), boosting wild meat harvest and trade to meet urban demands (7) and, thus, increasing pressure on wildlife populations.

Mammal population densities were higher inside than outside protected areas (Fig. 4). However, hunting pressure reduced mammal abundances even within protected areas (Fig. 4). Overhunting within protected areas is ubiquitous across the Amazon, Africa, and Asia (8, 20). Although our results suggest that the effects within are less detrimental than outside reserves, gazettement of protected areas seems insufficient to safeguard wildlife populations if not accompanied with improved reserve management, effective law enforcement, and on-ground protection efforts (20).

Effect sizes were similar across regions for both taxa, although slightly lower in South America for birds (Fig. 4). This indicates that overhunting is




Fig. 2. Forest plots of 384 and 1938 effect size estimates for birds and mammals, respectively.

(A) Birds and (B) mammals. RR , response ratios (effect sizes), black dots with 95% confidence intervals (CI) as gray lines. Overall weighted mean effect size estimate, black dashed line and red diamond. 95% CI of weighted mean effect size, red line. $RR = 0$, dashed gray line. Extremely negative effect sizes indicate local extirpations.

affecting mammal and bird populations similarly across the tropics. However, we found more studies in South America and Africa than Asia or Central America (Fig. 1), which implies that our findings are more generalizable for the former two regions. It also points out an urgent need to focus research efforts in less-studied areas before wildlife populations are completely extirpated. Unfortunately, overhunting has already emptied most Asian

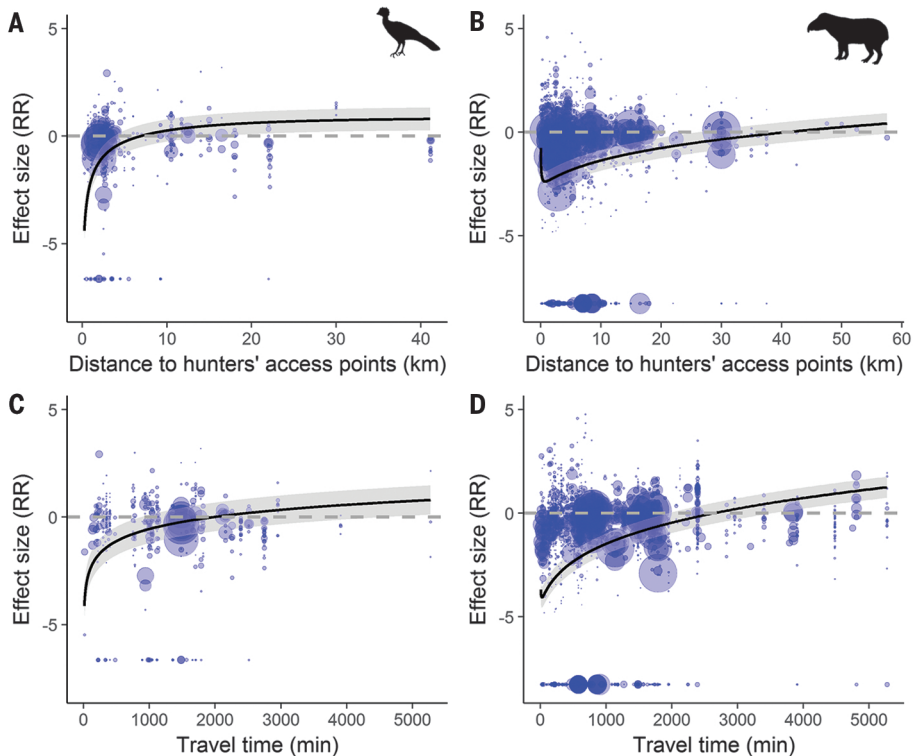


Fig. 3. Change in species abundance with distance to hunters' access points and travel time to major towns. (A and B) Distance to hunter's access points and (C and D) travel time to major towns; for birds (A and C) and mammals (B and D). RR, response ratios. $RR = 0$, dashed gray line; predicted mean effect size (with 95% CI in gray), black lines. Size of data points (in blue) is proportional to the sampling variance. Results obtained with single meta-regressions.

forests (7), leaving few un hunted control areas left for pairwise comparisons.

The most important terms retained in our multiple meta-regression models were distance for both groups (tables S6 and S7) and the interactions between guild, body size, and distance for mammals (see fig. S5 and explanations above). Our best models were significant according to omnibus tests (birds: $Q_M = 3157.5$, $P < 0.001$, McFadden pseudo- R^2 : 0.29; mammals: $Q_M = 19207.3$, $P < 0.001$, McFadden pseudo- R^2 : 0.18); however, residual heterogeneity was large (table S7), indicating that hunting is a multifaceted phenomenon influenced by additional factors, some of which were not included in our models (e.g., food security). Additionally, confounding variables such as small-scale habitat clearing and road disturbance are correlated with distance to settlements and roads (21). However, we minimized their influence as much as possible by avoiding pairwise comparisons where disturbances other than hunting were apparent.

Overexploitation is a long-established major driver of wildlife population declines and extinctions in terrestrial ecosystems which, to date, has not been successfully mitigated and rather shows an increasing trajectory in recent decades (22). Pleistocene extinctions were triggered in part by human hunters (23), and ongoing wildlife population declines and (near) extinctions of large-bodied species seem to share similar pathways. Consequently, defaunation is rendering tropical forests, savannahs, and grasslands "empty" (16),

with populations so sparse that the strength of species interactions is declining dramatically. The subtle nature of this process makes it undetectable by remote-sensing techniques, which are key to monitor deforestation but prove futile to track on-ground changes in biodiversity and ecological functioning (24). Matching the findings of many regionally specific studies (5, 10), our meta-analysis shows that large vertebrates of various functional groups are depleted in the vicinity of settlements and roads. Our estimated hunting-depletion distances can be used to assess ecosystem degradation as a result of current and future road developments and settlement establishment. Recently, Peres *et al.* (25) estimated that 32.4% of the remaining forest across the Brazilian Amazon (~1 million km^2) is affected by hunting on the basis of hunting distances of 6 km from settlements. Our results, however, indicate that the Amazon forest area affected by hunting-induced defaunation might be much larger. By 2050, with millions of kilometers of roads planned in developing countries (26), and human population and associated demand for wild meat increasing steadily, it is likely that the term "remoteness" will be a ghost of the past, with the last remnant half-depleted mammal and bird populations persisting in few protected areas. This can be ameliorated if we undertake coordinated strategies to expand the current network of protected areas, limit human encroachment around them, monitor hunting activities, and control overexploitation via law enforcement, if needed,

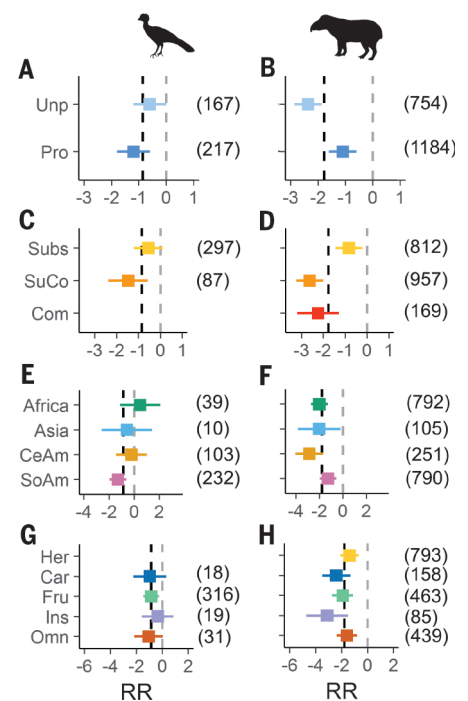


Fig. 4. Change in species abundance for different levels of protection, type of hunting, regions, and feeding guilds. (A to H) Parameters as labeled for birds (left) and mammals (right). Number of effect sizes is shown between brackets. RR, response ratios. Mean weighted effect size, dashed black line; $RR = 0$, dashed gray line; Unp, unprotected areas; Prot, protected areas; Subs, subsistence hunting; SuCo, subsistence and commercial hunting; Com, commercial hunting; CeAm, Central America; SoAm, South America; Her, herbivores; Car, carnivores; Fru, frugivores; Ins, insectivores; and Omn, omnivores. Results obtained with single meta-regressions. None of the studies reported on bird hunting for commercial purposes solely.

while implementing alternative livelihood programs for wild meat-dependent communities.

REFERENCES AND NOTES

- G. Ceballos *et al.*, *Sci. Adv.* **1**, e1400253 (2015).
- L. Gibson *et al.*, *Nature* **478**, 378–381 (2011).
- W. F. Laurance *et al.*, *Curr. Biol.* **25**, R259–R262 (2015).
- W. F. Laurance, A. Balmford, *Nature* **495**, 308–309 (2013).
- C. A. Peres, I. R. Lake, *Conserv. Biol.* **17**, 521–535 (2003).
- K. H. Redford, *Bioscience* **42**, 412–422 (1992).
- E. J. Milner-Gulland, E. L. Bennett, *Trends Ecol. Evol.* **18**, 351–357 (2003).
- R. D. Harrison, *Bioscience* **61**, 919–924 (2011).
- R. Dirzo *et al.*, *Science* **345**, 401–406 (2014).
- K. A. Abernethy, L. Coad, G. Taylor, M. E. Lee, F. Maisels, *Philos. Trans. R. Soc. London B Biol. Sci.* **368**, 20120303 (2013).
- Materials and methods are available as supplementary materials.
- L. V. Hedges, J. Gurevitch, P. S. Curtis, *Ecology* **80**, 1150–1156 (1999).
- A. Nelson, "Travel time to major cities: A global map of accessibility" (Global Environment Monitoring Unit—Joint Research Centre of the European Commission, Ispra, Italy, 2008).
- S. J. Wright, *Perspect. Plant Ecol. Evol. Syst.* **6**, 73–86 (2003).
- W. J. Ripple *et al.*, *R. Soc. Open Sci.* **3**, 160498 (2016).
- W. J. Ripple *et al.*, *Science* **343**, 1241484 (2014).

17. J. G. Robinson, E. L. Bennett, *Hunting for Sustainability in Tropical Forests* (Columbia Univ. Press, New York, 2000).
18. W. F. Laurance, S. Sloan, L. Weng, J. A. Sayer, *Curr. Biol.* **25**, 3202–3208 (2015).
19. G. R. Clements *et al.*, *PLOS ONE* **9**, e115376 (2014).
20. W. F. Laurance *et al.*, *Nature* **489**, 290–294 (2012).
21. A. Benítez-López, R. Alkemade, P. A. Verweij, *Biol. Conserv.* **143**, 1307–1316 (2010).
22. S. L. Maxwell, R. A. Fuller, T. M. Brooks, J. E. Watson, *Nature* **536**, 143–145 (2016).
23. C. Sandom, S. Faurby, B. Sandel, J.-C. Svenning, *Proc. Biol. Sci.* **281**, 20133254 (2014).
24. C. A. Peres, J. Barlow, W. F. Laurance, *Trends Ecol. Evol.* **21**, 227–229 (2006).
25. C. A. Peres, T. Emilio, J. Schiatti, S. J. Desmoulière, T. Levi, *Proc. Natl. Acad. Sci. U.S.A.* **113**, 892–897 (2016).
26. J. Dulac, “Global land transport infrastructure requirements: Estimating road and railway infrastructure capacity and costs to 2050” (International Energy Agency, Paris, 2013).

ACKNOWLEDGMENTS

Data reported in the paper are available at www.globo.info. We are grateful to S. Blake, A. Bowkett, J. Demmer, and T. Gray for kindly sharing their data with us. W. Viechtbauer helped

with data analyses. We thank the authors of the studies used in this meta-analysis. D.J.I. was funded by a Doctoral Training Grant from the University of Sussex.

SUPPLEMENTARY MATERIALS

www.sciencemag.org/content/356/6334/180/suppl/DC1
Materials and Methods
Figs. S1 to S6
Tables S1 to S7
References (27–159)

13 September 2016; accepted 9 March 2017
10.1126/science.aaj1891

COGNITIVE SCIENCE

Semantics derived automatically from language corpora contain human-like biases

Aylin Caliskan,^{1*} Joanna J. Bryson,^{1,2*} Arvind Narayanan^{1*}

Machine learning is a means to derive artificial intelligence by discovering patterns in existing data. Here, we show that applying machine learning to ordinary human language results in human-like semantic biases. We replicated a spectrum of known biases, as measured by the Implicit Association Test, using a widely used, purely statistical machine-learning model trained on a standard corpus of text from the World Wide Web. Our results indicate that text corpora contain recoverable and accurate imprints of our historic biases, whether morally neutral as toward insects or flowers, problematic as toward race or gender, or even simply veridical, reflecting the status quo distribution of gender with respect to careers or first names. Our methods hold promise for identifying and addressing sources of bias in culture, including technology.

We show that standard machine learning can acquire stereotyped biases from textual data that reflect everyday human culture. The general idea that text corpora capture semantics, including cultural stereotypes and empirical associations, has long been known in corpus linguistics (1, 2), but our findings add to this knowledge in three ways. First, we used word embeddings (3), a powerful tool to extract associations captured in text corpora; this method substantially amplifies the signal found in raw statistics. Second, our replication of documented human biases may yield tools and insights for studying prejudicial attitudes and behavior in humans. Third, since we performed our experiments on off-the-shelf machine learning components [primarily the Global Vectors for Word Representation (GloVe) word embedding], we show that cultural stereotypes propagate to artificial intelligence (AI) technologies in widespread use.

Before presenting our results, we discuss key terms and describe the tools we use. Terminology varies by discipline; these definitions are intended for clarity of the present article. In AI and ma-

chine learning, bias refers generally to prior information, a necessary prerequisite for intelligent action (4). Yet bias can be problematic where such information is derived from aspects of human culture known to lead to harmful behavior. Here, we will call such biases “stereotyped” and actions taken on their basis “prejudiced.”

We used the Implicit Association Test (IAT) as our primary source of documented human biases (5). The IAT demonstrates enormous differences in response times when subjects are asked to pair two concepts they find similar, in contrast to two concepts they find different. We developed our first method, the Word-Embedding Association Test (WEAT), a statistical test analogous to the IAT, and applied it to a widely used semantic representation of words in AI, termed word embeddings. Word embeddings represent each word as a vector in a vector space of about 300 dimensions, based on the textual context in which the word is found. We used the distance between a pair of vectors (more precisely, their cosine similarity score, a measure of correlation) as analogous to reaction time in the IAT. The WEAT compares these vectors for the same set of words used by the IAT. We describe the WEAT in more detail below.

Most closely related to this paper is concurrent work by Bolukbasi *et al.* (6), who propose a method to “debias” word embeddings. Our work is complementary, as we focus instead on rigorously

demonstrating human-like biases in word embeddings. Further, our methods do not require an algebraic formulation of bias, which may not be possible for all types of bias. Additionally, we studied the relationship between stereotyped associations and empirical data concerning contemporary society.

Using the measure of semantic association described above, we have been able to replicate every stereotype that we tested. We selected IATs that studied general societal attitudes, rather than those of subpopulations, and for which lists of target and attribute words (rather than images) were available. The results are summarized in Table 1.

Greenwald *et al.* introduced and validated the IAT by studying biases that they consider nearly universal in humans and about which there is no social concern (5). We began by replicating these inoffensive results for the same purposes. Specifically, they demonstrated that flowers are significantly more pleasant than insects, based on the reaction latencies of four pairings (flowers + pleasant, insects + unpleasant, flowers + unpleasant, and insects + pleasant). Greenwald *et al.* measured effect size in terms of Cohen’s *d*, which is the difference between two means of log-transformed latencies in milliseconds, divided by the standard deviation. Conventional small, medium, and large values of *d* are 0.2, 0.5, and 0.8, respectively. With 32 participants, the IAT comparing flowers and insects resulted in an effect size of 1.35 ($P < 10^{-6}$). Applying our method, we observed the same expected association with an effect size of 1.50 ($P < 10^{-7}$). Similarly, we replicated Greenwald *et al.*’s finding (5) that musical instruments are significantly more pleasant than weapons (see Table 1).

Notice that the word embeddings “know” these properties of flowers, insects, musical instruments, and weapons with no direct experience of the world and no representation of semantics other than the implicit metrics of words’ co-occurrence statistics with other nearby words.

We then used the same technique to demonstrate that machine learning absorbs stereotyped biases as easily as any other. Greenwald *et al.* (5) found extreme effects of race as indicated simply by name. A bundle of names associated with being European American was found to be significantly more easily associated with pleasant than unpleasant terms, compared with a bundle of African-American names.

In replicating this result, we were forced to slightly alter the stimuli because some of the original African-American names did not occur

¹Center for Information Technology Policy, Princeton University, Princeton, NJ, USA. ²Department of Computer Science, University of Bath, Bath BA2 7AY, UK.
*Corresponding author. Email: aylinc@princeton.edu (A.C.); jjb@alum.mit.edu (J.J.B.); arvindn@cs.princeton.edu (A.N.)

Copyright 2017 American Association for the Advancement of Science. All rights reserved.