State of the evidence: a survey of global disparities in clinical trials

Iain James Marshall ¹, Veline L'Esperance,¹ Rachel Marshall,² James Thomas,³ Anna Noel-Storr,⁴ Frank Soboczenski,¹ Benjamin Nye,⁵ Ani Nenkova,⁶ Byron C Wallace⁵

To cite: Marshall IJ, L'Esperance V, Marshall R, *et al*. State of the evidence: a survey of global disparities in clinical trials. *BMJ Global Health* 2021;**6**:e004145. doi:10.1136/bmjgh-2020-004145

Handling editor Seye Abimbola

Received 10 October 2020
Revised 12 November 2020
Accepted 16 November 2020

© Author(s) (or their employer(s)) 2021. Re-use permitted under CC BY. Published by BMJ.

¹School of Population Health and Environmental Sciences, King's College London, London, UK

²Editorial and Methods Department, Cochrane, London, UK

³EPPI-Centre, Department of Social Science, UCL, London, UK

⁴Cochrane Dementia Group, University of Oxford, Oxford, UK

⁵Khoury College of Computer Sciences, Northeastern University, Boston, Massachusetts, USA

⁶Computer and Information Science, University of Pennsylvania, Philadelphia, Pennsylvania, USA

Correspondence to

Dr Iain James Marshall;
iain.marshall@kcl.ac.uk

ABSTRACT

Introduction Ideally, health conditions causing the greatest global disease burden should attract increased research attention. We conducted a comprehensive global study investigating the number of randomised controlled trials (RCTs) published on different health conditions, and how this compares with the global disease burden that they impose.

Methods We use machine learning to monitor PubMed daily, and find and analyse RCT reports. We assessed RCTs investigating the leading causes of morbidity and mortality from the Global Burden of Disease study. Using regression models, we compared numbers of actual RCTs in different health conditions to numbers predicted from their global disease burden (disability-adjusted life years (DALYs)). We investigated whether RCT numbers differed for conditions disproportionately affecting countries with lower socioeconomic development.

Results We estimate 463 000 articles describing RCTs (95% prediction interval 439 000 to 485 000) were published from 1990 to July 2020. RCTs recruited a median of 72 participants (IQR 32–195). 82% of RCTs were conducted by researchers in the top fifth of countries by socio-economic development. As DALYs increased for a particular health condition by 10%, the number of RCTs in the same year increased by 5% (3.2%–6.9%), but the association was weak (adjusted $R^2=0.13$). Conditions disproportionately affecting countries with lower socioeconomic development, including respiratory infections and tuberculosis (7000 RCTs below predicted) and enteric infections (9700 RCTs below predicted), appear relatively under-researched for their disease burden. Each 10% shift in DALYs towards countries with low and middle socioeconomic development was associated with a 4% reduction in RCTs (3.7%–4.9%). These disparities have not changed substantially over time.

Conclusion Research priorities are not well optimised to reduce the global burden of disease. Most RCTs are produced by highly developed countries, and the health needs of these countries have been, on average, favoured.

INTRODUCTION

The Global Burden of Disease (GBD) study has estimated the leading causes of death and disability worldwide.¹ A key strategy to reduce death and morbidity is to conduct

Key questions

What is already known?

- ▶ Prior studies have manually investigated the relationship between published research in different health conditions and the global burden of disease that they impose.
- ▶ However, these analyses have been mostly limited to estimates of research funding from national funders, or smaller scale analysis of older publication records.
- ▶ These studies have highlighted disparities in research relative to burden, but they are not sufficient to enable global targeting of research to optimise improvements in disease burden.

What are the new findings?

- ▶ We automatically process all of PubMed, allowing us to conduct a continually updated, comprehensive analysis of published reports of randomised controlled trials (RCTs), including the number of participants per RCT and the health conditions studied.
- ▶ We found that considerable disparities exist between the relative volume of evidence on some conditions and the global burden of disease that they impose, as calculated by the Global Burden of Disease study.
- ▶ Further, our analysis suggests that there exists a smaller amount of evidence for conditions that impose a comparatively large burden of disease in lower-income countries.

What do the new findings imply?

- ▶ Looking at numbers of RCTs published, and the numbers of participants in these trials, it seems that research priorities are not optimised to reduce the global burden of disease, and that research for conditions affecting higher-income countries has, on average, been favoured.
- ▶ The findings from this study could help research funders to focus research investment in areas where the largest reductions in disease burden could be made.

randomised controlled trials (RCTs) to determine the effects of interventions for improving health.² However, RCTs are highly expensive to conduct; this cost leads to a risk of health inequality, where the health needs of countries which lack resources are

neglected. Health inequalities have been defined as differences in health outcome which are avoidable, remediable and considered unjust.³ From an ethical standpoint (and particularly a concern for justice), health resources should be distributed to help those with the greatest need, and in those who stand to see the greatest benefit. There are very many factors which could affect which health problems are researched, not limited to: the existence of highly effective treatments for certain conditions; national governments prioritising conditions which affect their own country's population and the unpredictable nature of scientific breakthroughs in therapy. Pharmaceutical companies are likely to develop drugs which have the potential to be sold at profit; this can lead to neglect of conditions (eg, infectious diseases) which disproportionately affect low income countries.⁴ Nonetheless, a rational and ethical approach to funding research—seeking to reduce health inequalities and maximise global societal benefit—would take account of the global disease burden.

Studies from the USA and the UK have reported that disparities exist in national health research funding for certain conditions. Some conditions, such as diabetes and cancer, have attracted a comparatively generous amount of research funding in these countries relative to their national and global burden, whereas other conditions, such as stroke and respiratory diseases, have been relatively neglected.^{5–8} In a later study, Evans *et al* mapped the volume of research (RCTs and systematic reviews) published in 2005 to conditions studied in the GBD study. They concluded that global disability-adjusted life years (DALYs; a measure that incorporates the quality of lived years with respect to physical and mental functioning criteria) in 2004 did not explain any of the topic distribution of research articles published in the following year.⁹ By contrast, Atal *et al* examined 117 000 new registrations for clinical trials globally from 2006 to 2015, and reported that trials registered were proportionate to disease burden in high-income countries, but not in non-high income countries.¹⁰

To our knowledge, there has been no comprehensive, global study comparing published research (both in terms of trial and participant numbers) with the relative burden of each condition. We address this gap by evaluating the total volume of evidence (numbers of RCTs and trial participants) published for clinical conditions via an automated analysis of all research articles indexed in PubMed (<http://pubmed.ncbi.nlm.nih.gov>) from 1990 onwards. Our system uses natural language processing to automatically identify text within articles that describes the conditions being studied, and maps these to GBD-defined categories. This permits comparison between the global disease burden and the amount of published evidence that exists for each condition, and could help prioritise research funding to relatively understudied conditions with high burden. Uniquely, as our method is fully automatic, we are able to analyse articles at large scale (classification of 100's of thousands of documents),

and with continuous, live updating—neither of which have been feasible so far with manual methods.

METHODS

We developed a machine learning system which maintains a live database of annotated RCT reports, named Trialstreamer. We have described the computational methods and accuracy of the system components in detail elsewhere,¹¹ and summarise the key points relevant to the current study below.

Monitoring PubMed for RCTs in humans

We first downloaded the full PubMed database (annual baseline data, from 11 December 2018) and then subsequently monitored and added all newly published articles daily. This comprises >30 million articles. We automatically identify the subset of articles that describe RCTs using a machine learning system, which operates over the text of the title and abstract, and the Publication Type database index.¹² Briefly, we ensemble a Support Vector Machine (SVM)¹³ and a Convolutional Neural Network.¹⁴ These machine learning models are *trained*ⁱ on 280 000 abstracts manually labelled as being RCTs or not by Cochrane Crowd, (<https://crowd.cochrane.org>) a collaborative citizen science project. We do not rely on the Publication Type index alone, as we have previously found it to miss a substantial proportion of the 5–7 most recent years of articles¹¹ (due to delay in manual indexing after publication). We next removed any RCTs that were not conducted in humans (eg, animal or agricultural studies) using an SVM model. This model was trained using labelled abstracts from PubMed (labels derived as a function of whether the Medical Subject Headings (MeSH) (MeSH refers to the MeSH vocabulary, maintained by the National Library of Medicine and used for indexing articles in the MEDLINE database) term 'Humans' had been applied or not).

Automatic information extraction

To characterise the trial participants and outcomes, we first automatically extract 'snippets' of text (typically a few consecutive words within the abstract) describing these components. For this, we use a sequence tagging machine learning model (namely a Long Short-Term Memory network¹⁵ followed by a conditional random field.^{16–18} This model was trained using a publicly available dataset called EBM-NLP (Evidence Based Medicine—Natural Language Processing), which comprises 5000 RCT abstracts in which text spans relevant to the respective characteristics (participants, interventions, comparators and outcomes) were manually annotated.

ⁱModel 'training' describes the process of adjusting model parameters to align with (typically large amounts of) data. Here, the models consume text documents (abstracts), with manual labels designating whether these are RCTs or not. Once trained (ie, once the parameters are estimated) the model can be applied to new, unseen abstracts to provide predictions.

These text descriptions are then standardised to MeSH terms. This allows us to infer, for example, that ‘myocardial infarction’ and ‘heart attack’ refer to the same condition, and that both are subtypes of cardiovascular disease. To achieve this step, we use an approach similar to that described by Demner-Fushman *et al*, in their MetaMap software.¹⁹ This approach involves generating a large database of text synonyms for MeSH terms, done by finding alternative descriptions for each term in linked vocabularies. These synonyms are then sought in the extracted population and outcomes text spans.

Linking RCTs with data on global disease burden

In this study, we use the 22 level 2 categories defined by the GBD collaboration. These categories describe broad disease areas such as ‘Cardiovascular diseases’ and ‘Mental disorders’ (the full list is provided in table 1). The GBD collaboration makes available International Classification of Diseases (ICD-10) codes which define the scope of each category. We automatically link these ICD-10 codes to the equivalent MeSH terms using a tool and dataset named Metathesaurus, from the Unified Medical Language System.²⁰ In brief, Metathesaurus provides a database of synonyms between different health science vocabularies, and thus lists of synonymous terms between ICD-10 and MeSH can be retrieved. One clinical author (IJM) manually checked all of the automatically generated mappings from MeSH to GBD category (both removing any errors from the automatic mapping, and adding links where they had been missed). This resulted in a final mapping of 5256 parent MeSH terms to one of the 22 GBD categories.

Descriptive statistics

We report descriptive statistics for RCT counts and participant numbers. We use the country of the first author’s affiliation as a proxy for the country organising or funding the study. In order to capture trials addressing prevention, we also consider RCTs as relevant to the GBD category where rates of the health condition in question are identified as an outcome, even where participants did not have the condition at trial onset. For example, trials that recruited people with diabetes but measured the effects of an intervention on subsequent coronary heart disease were considered as relevant to both the ‘diabetes’ and ‘cardiovascular diseases’ categories. We also present separate estimates based on this distinction, that is, separating instances where a given GBD condition was found in the trial population (‘treatment’ trials) from those where it was inferred to describe an outcome (‘prevention’ trials).

Estimating uncertainty due to prediction error

We randomly sampled articles (titles and abstracts) from the finalised dataset, and two authors independently manually labelled these with respect to study design (RCT in humans vs not; 500 abstracts, done by IJM and BCW), GBD category (250 abstracts, done by IJM and VL), and

number of participants randomised (500 abstracts, done by IJM and BCW). Authors then met to resolve discrepancies by consensus. We then used the bootstrap method to estimate 95% prediction intervals for final counts, using the method described by Fox and colleagues.^{21 22} In short, we use our manually labelled data to calculate 10 000 bootstrap estimates of the precision and recall of each labelling task.²³ These precision and recall estimates are then used to adjust the automated results from the full dataset (generating 10 000 adjusted count estimates for each classification task), which can be used to construct a 95% prediction interval. To evaluate uncertainty around extracted sample sizes we use a similar approach, except we used estimates of the absolute error taken from the manually labelled abstracts. We then simulated the effect of this error on the full dataset using the bootstrap method with 10 000 iterations. These intervals permit assessment of the degree to which model misclassification and bias has affected the final results.

Assessing associations between disease burden and published RCTs

To examine associations between global disease burden and volume of evidence (both numbers of RCTs and participants), we used a series of linear regression models with log-transformed global DALYs as the predictor, and the log-transformed numbers of global RCTs as the dependent variable. We evaluate annual data on DALYs for each clinical condition for the full period available from the GBD study (1990–2017), and adjust for trends over time by incorporating the year as a categorical predictive variable. A detailed description alongside statistical code and data for these analyses is available at the Open Science Framework (DOI: 10.17605/osf.io/3db76).

In our first regression analysis, we examined the association of global DALYs on RCT publications in the same year. We explored incorporating the year (as categorical variable) both as a fixed and a random effect, evaluating using the Hausman test. We compare observed RCT counts for each GBD condition with those predicted by the model.

In our second regression analysis, we examine the effect of lag (given that RCTs will publish reports some time after measures of disease burden are known). Here, we use finite length distributed lag models, and evaluate the association of global DALYs with up to 1, 2, 3, 4 and 5 years lag on published RCTs in the index year. We compare models using the adjusted R² statistic, the Akaike information criterion (AIC) and Bayesian information criterion (BIC).

In our third regression analysis, we add a variable to account for whether a particular condition tends to impose a disproportionate burden in high-income, rather than low-income and middle-income countries.

We examine associations between socioeconomic status and evidence volume, using the Sociodemographic Index (SDI) of the country from the GBD study. As a

Table 1 RCTs in PubMed with estimates of trial population sizes, and comparison with measures of global disease burden

Disease category	Total RCTs, thousands, (95% prediction interval) (rank)	RCTs with the condition as a population, thousands, (95% prediction interval) (rank)	RCTs with the condition as an outcome, Thousands, (95% prediction interval) (rank)	DALYs, thousands (rank)	Death, thousands (rank)	YLD, thousands (rank)	YLLs, thousands (rank)
Mental disorders	55 (49 to 61)(1)	28 (25 to 31)(3)	40 (36 to 44)(1)	2900(11)	<0.1(21)	2900(2)	0.4(21)
Cardiovascular diseases	54 (49 to 60)(2)	37 (33 to 41)(1)	31 (28 to 34)(2)	8700(1)	410(1)	710(11)	8000(1)
Other non-communicable diseases	39 (35 to 43)(3)	24 (21 to 26)(5)	20 (18 to 22)(4)	3500(6)	33(15)	1300(6)	2200(9)
Neurological disorders	39 (35 to 43)(4)	26 (23 to 28)(4)	20 (18 to 23)(3)	2500(12)	58(7)	1700(3)	820(17)
Neoplasms	31 (28 to 34)(5)	29 (26 to 31)(2)	6 (6 to 7)(12)	5400(4)	210(2)	150(20)	5200(4)
Maternal and neonatal disorders	24 (22 to 26)(6)	19 (17 to 21)(6)	13 (11 to 14)(5)	7200(2)	77(5)	630(13)	6500(2)
Diabetes and kidney diseases	21 (19 to 23)(7)	18 (17 to 20)(7)	5 (4 to 5)(13)	2200(16)	52(9)	900(9)	1300(15)
Unintentional injuries	18 (16 to 20)(8)	11 (10 to 12)(11)	9 (8 to 10)(8)	3300(7)	51(10)	800(10)	2500(7)
Musculoskeletal disorders	18 (16 to 20)(9)	13 (12 to 15)(8)	7 (6 to 8)(10)	3100(8)	2(19)	3000(1)	65(19)
Digestive diseases	17 (15 to 19)(10)	13 (12 to 14)(9)	7 (6 to 8)(11)	2100(18)	56(8)	440(14)	1700(14)
Substance use disorders	17 (15 to 19)(11)	8 (7 to 8)(14)	12 (11 to 13)(6)	1000(22)	8(18)	710(12)	310(18)
Skin and subcutaneous diseases	16 (15 to 18)(12)	8 (7 to 9)(13)	12 (11 to 13)(7)	1000(21)	2(20)	990(7)	55(20)
Chronic respiratory diseases	14 (13 to 16)(13)	13 (11 to 14)(10)	4 (4 to 5)(15)	2900(10)	99(4)	990(8)	1900(12)
Sense organ diseases	14 (13 to 16)(14)	8 (7 to 9)(12)	8 (7 to 9)(9)	1400(20)	n/a	1400(4)	n/a
HIV/AIDS and sexually transmitted infections	8 (7 to 9)(15)	7 (6 to 8)(15)	3 (3 to 4)(17)	2500(13)	41(11)	130(21)	2300(8)
Respiratory infections and tuberculosis	8 (7 to 8)(16)	4 (3 to 4)(17)	5 (4 to 5)(14)	6500(3)	120(3)	290(17)	6200(3)
Other infectious diseases	7 (6 to 8)(17)	5 (4 to 5)(16)	3 (2 to 3)(18)	3000(9)	39(12)	110(22)	2900(6)
Nutritional deficiencies	5 (4 to 6)(18)	2 (1 to 2)(19)	4 (3 to 4)(16)	2100(17)	12(17)	1300(5)	840(16)
Neglected tropical diseases and malaria	3 (3 to 4)(19)	2 (2 to 3)(18)	1 (1 to 1)(19)	2300(14)	27(16)	410(15)	1900(10)
Self-harm and interpersonal violence	2 (1 to 2)(20)	0.8 (0.7 to 0.9) (20)	1.0 (0.9 to 1)(20)	2000(19)	37(14)	180(19)	1900(13)
Enteric infections	1 (1 to 2)(21)	0.6 (0.5 to 0.7) (21)	0.8 (0.8 to 0.9)(21)	4100(5)	64(6)	240(18)	3800(5)
Transport injuries	0.4 (0.4 to 0.5)(22)	0.2 (0.2 to 0.2) (22)	0.2 (0.2 to 0.2)(22)	2200(15)	38(13)	300(16)	1900(11)

DALYs, disability-adjusted life years; n/a, not available; RCTs, randomised controlled trials; YLD, years lived with disability; YLLs, years of life lost.

Figure 1 Global clinical trial publications by first author location, 1990–2017.

metric of disparity in disease burden, we used the ratio of DALYs occurring in high SDI countries (the top fifth) to DALYs occurring in low and medium SDI countries (ie, the bottom four fifths).

Finally, we conducted a sensitivity analysis aiming to discover whether our results are sensitive to ‘salami slicing’ (the practice of publishing results from a single trial across multiple publications, which could lead us to overestimate the evidence base). We first estimated the mean number of publications per RCT for each GBD condition, utilising the subset of articles which contained a unique trial registry identifier. Then we use these values to discount our estimate of RCT numbers for each GBD condition. We repeated our regression analyses with these discounted RCT numbers.

RESULTS

We analysed records of 31 367 011 articles indexed in PubMed on 27 July 2020. We estimate that 463 000 articles describing RCTs (95% prediction interval 439 000 to 485 000). We show how these trials are distributed globally (by location of the first author) in figure 1. The number of RCTs published for each of the 22 ‘level 2’ disease categories from the 2017 GBD study is presented in table 1 and figures 2 and 3.

Mental disorders (ranked 11th in terms of total DALYs) had the largest number of RCTs with 55 000, followed by cardiovascular diseases (ranked first) with 54 000. We

Figure 2 Disease categories causing the highest global burden (by DALY), with the number of RCTs published for each). DALY, disability-adjusted life years; RCT, randomised controlled trial.

Figure 3 Trends over time in RCT publications of the top fifteen categories by publication rates from 1990 to 2019. RCT, randomised controlled trials.

separate the results from trials where the clinical condition occurred in the trial population (generally trials of treatment strategies for), and where the condition was a trial outcome (typically trials looking at prevention or diagnosis of the condition) in table 1 and figure 2. In all cases except for mental disorders, the numbers of trials and participants in the ‘treatment’ trials exceeded those in the ‘prevention’ trials.

Overall, we estimate the median number of participants randomised per RCT was 72, with IQR 32–195. We show how the distribution of RCT sample sizes varies for each of the GBD categories in figure 4.

In figure 5 we scatter the volume of evidence estimated for conditions against the global burden that they impose (totals in period 1990–2017) for the 22 ‘level 2’ GBD categories.

Our primary analysis found that, on average, increasing DALYs were associated with a statistically significant increase in RCT publication in the same year, but that variance was not well explained by the model ($\beta=0.52$ for log DALYs, 95% CI 0.33 to 0.70; $R^2=0.13$). This translates to an average increase in the number of RCTs published of 5% (3.2% to 6.9%) for each 10% increase in DALYs in the same year. We report

Figure 4 Number of participants recruited to RCTs pertaining to each of the Global Burden of Disease categories. RCTs, randomised controlled trials.

Figure 5 Total RCTs published globally from 1990 to 2017 against disability-adjusted life years in the same period. RCTs, randomised controlled trials.

differences between the expected vs observed number of RCTs and trial participants based on this model in table 2.

We found that models incorporating between 3 and 5 years of lag had improved fit compared with the equivalent model with no lag effects (demonstrated by reduction in AIC, BIC, and increase in adjusted R^2). The model incorporating lag up to 3 years found an average 7.2% increase in published RCTs for every 10% increase in DALYs (adjusted $R^2=0.29$). The models with lag up to 4, and up to 5 years gave near identical estimates and fit (full analysis results provided in the statistical appendix).

We provide full results for our sensitivity analysis (where we adjust for the differential impact of RCTs producing multiple publications in each disease in our statistical appendix (<https://osf.io/zf97n/>)). We found that RCTs with a registry identifier produced a mean of 1.3 articles, but that there was variation between GBD categories (from 1.2 articles/RCT in nutritional deficiencies to 1.5 articles/RCT in cardiovascular diseases). Compared with our primary results, our sensitivity analysis found a near identical relationship between DALYs and published trials ($\beta=0.48$, 95% CI 0.30 to 0.66; $R^2=0.13$). Adjusting for multiple publications did not affect which conditions

Table 2 Difference between measured number of RCTs, and number predicted based on global disease burden (negative numbers=fewer RCTs were published than predicted)

Disease category	Total (1990–2017)	1990–1993	1994–1997	1998–2001	2002–2005	2006–2009	2010–2013	2014–2017
Cardiovascular diseases	+34 000	+2400	+3000	+3600	+4400	+5100	+6700	+8300
Chronic respiratory diseases	+4200	+370	+500	+670	+720	+710	+620	+630
Diabetes and kidney diseases	+10 000	+490	+710	+800	+1200	+1600	+2400	+3400
Digestive diseases	+7700	+590	+850	+920	+1100	+1100	+1400	+1700
Enteric infections	–9700	–950	–1200	–1200	–1400	–1500	–1700	–1800
HIV/AIDS and sexually transmitted infections	–1100	–120	–180	–190	–400	–280	–51	+100
Maternal and neonatal disorders	+6400	+120	+410	+440	+530	+670	+1500	+2700
Mental disorders	+38 000	+1100	+1900	+2600	+4100	+6100	+9200	+13 000
Musculoskeletal disorders	+5900	+50	+130	+360	+780	+1000	+1600	+2000
Neglected tropical diseases and malaria	–5400	–470	–570	–650	–790	–880	–990	–1100
Neoplasms	+14 000	+980	+1300	+1200	+1700	+2200	+3000	+4000
Neurological disorders	+25 000	+1100	+1700	+2100	+2800	+3800	+5600	+7800
Nutritional deficiencies	–3300	–380	–450	–460	–410	–520	–560	–530
Other infectious diseases	–2500	–390	–400	–340	–350	–260	–350	–410
Other non-communicable diseases	+24 000	+890	+1400	+1800	+2800	+3900	+6100	+7400
Respiratory infections and tuberculosis	–7000	–830	–980	–970	–1100	–1200	–1100	–1000
Self-harm and interpersonal violence	–6800	–620	–840	–890	–920	–1,000	–1,200	–1,300
Sense organ diseases	+6500	+470	+630	+800	+780	+1000	+1300	+1500
Skin and subcutaneous diseases	+9300	+660	+780	+980	+1200	+1500	+1900	+2300
Substance use disorders	+9800	+370	+750	+870	+1200	+1700	+2200	+2700
Transport injuries	–8000	–670	–840	–940	–1100	–1200	–1500	–1700
Unintentional injuries	+6300	+91	+200	+310	+620	+840	+1600	+2600

RCTs, randomised controlled trials.

were found to have relatively high or low associated RCT publications.

In our final model, which examined the effect of the DALY location (ratio of DALYs in the top fifth of countries by SDI vs the bottom four fifths), we found that DALYs, and the DALY location were both associated with increased RCT publications ($\beta_{\text{DALY}}=0.75$, 95% CI 0.56 to 0.91, $\beta_{\text{location}}=0.44$, 95% CI 0.38 to 0.51; adjusted $R^2=0.33$). This may be interpreted that for every 10% shift in DALYs towards high-income countries, the number of RCTs increased by 4% (3.7% to 4.9%), and for every 10% increase in DALYs globally, RCTs increased by 7% (5.7% to 9.1%).

DISCUSSION

We present the results from a natural language processing system that we used to perform a large-scale, comprehensive, automated analysis of all published RCTs indexed in the PubMed database. We find the number of published RCTs covering health conditions correlates only weakly with the burden that they inflict globally. Some conditions (particularly respiratory infections and tuberculosis, enteric infections and transport injuries) attract substantially fewer published RCTs than expected based on the global burden.

Conditions with fewer RCTs and participants than expected tend to disproportionately affect countries with low and medium socioeconomic status. For instance, respiratory infections and tuberculosis, which had the largest disparity between global burden and number of published RCTs disproportionately affects populations in lower income countries.²⁴ By contrast, cancers and diabetes, for which there appears to be a relative abundance of evidence, are comparatively common in high-income countries, although this picture is changing over time.^{25 26} In other words, the top fifth of countries by SDI, who conduct the vast majority of RCTs, are more closely matching RCT conduct with disease burden within their own countries (although still weakly). These findings suggest that a research funding strategy based on individual countries priorities does not optimally meet global needs. However, the picture is complex: global development has meant that the most important diseases affecting high-income countries (ie, cardiovascular disease and cancers) have also become the top health priorities in low-income and middle-income countries a number of years later.

The median number of participants per RCTs was 72 (ie, a typical trial is likely to have around 36 participants per arm, or fewer for trials with >2 arms). Smaller trials can be conducted faster and at lower cost, and thus can be more responsive to need, but have important limitations.²⁷ Small trials are more susceptible to the effects of publication bias, and have a high likelihood of generating false positive, and false negative findings.²⁸ Meta-epidemiological studies have shown that trials with fewer than 100 participants per arm tend to exaggerate

treatment effects, compared with larger ones.^{29 30} Twenty-five years after Yusuf *et al* called for larger, simpler RCTs,³¹ unfortunately we find that a majority of published trials are likely to be too small to provide useful or definitive information.

Prior work has investigated the relationship between burden of disease and National Institutes of Health funding.^{5 8} These efforts relied on manual compilation of data, which necessarily limited the scope of analysis. Additionally, these efforts could only analyse conditions for which funding data was available. The authors of the landmark 1999 New England Journal of Medicine study noted that data were not available for conditions making up 38% of the total DALY estimate at that time.⁵ Conducting this analysis manually additionally precludes the realisation of a 'living' view of the evidence base. By contrast, we have developed a fully automated approach that facilitates comprehensive, real-time analysis of the published RCT literature. For illustration, at the time of writing (September 2020), our system has indexed 158 RCTs examining treatments for COVID-19. This contrasts with PubMed, which has indexed (largely manually) 58 with the equivalent MeSH terms.

Because our coverage of diseases is broader than prior work on the relationship between funding and burden of disease, the results are not directly comparable. However, some findings do qualitatively align with findings from prior analyses of research funding. For example, we also find that diabetes generated a relatively generous number of RCTs, as do cancers (in general). Future analyses using the methods and data could examine specific conditions within the categories presented here, investigating, for example, whether stroke is adequately researched compared with coronary artery disease, or the distribution of trials relevant to different subtypes of cancer.

Our analysis has some limitations. First, our primary analysis examines publications and not RCTs themselves. Our sensitivity analysis confirms previous findings that RCTs may generate more than one publication,³² and so our estimates of publication numbers are likely to be an overestimate of the number of trials conducted. However, we found adjusting for estimates of duplicate publication made little difference to estimates of research done as a function of global burden. Second, the automated data extraction methods used here are only able to count explicit mentions of a disease or disorder, and explicit descriptions of participant numbers from the abstract, but abstracts do not always contain all the necessary information (in this study, for example, 34% of articles publications did not report the number of study participants). There is evidence of improvement over time after the publication of the Consolidated Standards of Reporting Trials recommendations in 1996.³³

Last, we restrict our enquiry to PubMed and therefore have not counted RCTs which have produced publications in journals not indexed in PubMed. PubMed tends to under-represent articles published in non-English languages (though includes some where an

English-language translations abstracts exist). PubMed also does not index conference presentations or pharmaceutical company data repositories. By focusing on published articles, this analysis will by construction miss RCTs whose results are never published (although unpublished trial results are arguably not effectively contributing to reducing the global disease burden). Nonetheless, PubMed is reasonably comprehensive; and those RCT publications that it misses have been found to comprise a relatively small fraction of the available health research.^{34 35}

Why and how do some conditions attract relatively high numbers of RCTs and trial participants? We have described above how the priorities of research funders (nation states, pharmaceutical companies and disease-specific charities) may not align with the global need. But disparities in funding may not be the sole explanation: organisational and cultural differences within medical subdisciplines could plausibly play a role. National strategies exist for increasing participation in oncology research (eg, 2008 guidance from the American Society of Clinical Oncology which set a target of >10% of patients participating in clinical trials³⁶); a large proportion of research looking at how best to increase trial recruitment has arisen from cancer RCTs.³⁷ Further investigation of how the better performing medical disciplines achieve their high recruitment rates might uncover strategies which are more widely applicable.

Clinical trials are one small part of tackling the global disease burden and the UN sustainable development goals, and this wider landscape needs to be considered when determining research priorities. For example, the provision of clean water is likely to be more important for preventing enteric infections than any specific health intervention, and probably does not require an RCT to prove it.³⁸ There may be other conditions (eg, HIV³⁹) where established treatments can enable those with the condition to have a normal lifespan without significant disability. This contrasts with some types of cancer (eg, pancreatic cancer⁴⁰), where we still lack reliably effective treatments. For pharmacological research, the ability to run new trials depends on the availability of new candidate drugs to test.⁴¹ Such new candidate drugs depend on unpredictable scientific breakthroughs, which will not neatly tally with global disease burden rankings.

Nonetheless, these data on the relative attention attracted by different health conditions may provide a useful metric to consider alongside other indicators of research need. Given that the vast majority of trials are being conducted by the wealthiest 20% of countries, targeting national health priorities is not sufficient, and a global approach is needed. We publish daily updates of the figures in this paper at www.robotreviewer.net/stateoftheevidence, and make the source code and dataset freely available via the Open Science Framework (DOI: <https://dx.doi.org/10.17605/osf.io/3DB76>).

Twitter Iain James Marshall @ijmarshall

Acknowledgements We are enormously grateful to the volunteers from the Cochrane Crowd, who labelled hundreds of thousands of articles manually, enabling us to develop the machine learning system. Our grateful thanks to Danielle Wenner from the Department of Philosophy, Carnegie Mellon University, who reviewed early drafts and provided us with valuable feedback on the ethics of research funding distribution. We thank the Global Burden of Disease collaboration for making their data freely available for academic use.

Contributors Initial study conceived by IJM and BCW, with intellectual input from all authors. Automatic information extraction system created by: IJM, BCW, BN, AN and FS. Curation of data for training machine learning models: JT, AN-S, BCW, BN, AN and IJM. Statistical analysis designed and conducted by IJM, VL and BCW. Initial draft of manuscript by IJM and BCW. Critical revisions for important intellectual content: all; Approval of final manuscript: all.

Funding This work is funded by the UK Medical Research Council (MRC), through its Skills Development Fellowship programme (IJM), MR/N015185/1, and the US National Institutes of Health (NIH) under the National Library of Medicine, 2R01-LM012086-05 (IJM, FS and BCW). VL is funded by a National Institute for Health Research (NIHR), (Doctoral Research Fellow- DRF-2017-10-13). JT received grants from Cochrane during the conduct of the study.

Disclaimer The views expressed are those of the authors and not necessarily those of the NHS, the NIHR or the Department of Health and Social Care.

Map disclaimer The depiction of boundaries on this map does not imply the expression of any opinion whatsoever on the part of BMJ (or any member of its group) concerning the legal status of any country, territory, jurisdiction or area or of its authorities. This map is provided without any warranty of any kind, either express or implied.

Competing interests None declared.

Patient and public involvement Patients and/or the public were not involved in the design, or conduct, or reporting, or dissemination plans of this research.

Patient consent for publication Not required.

Provenance and peer review Not commissioned; externally peer reviewed.

Data availability statement Data are available in a public, open access repository. All data are available in a public, open access repository (DOI 10.17605/osf.io/3db76).

Open access This is an open access article distributed in accordance with the Creative Commons Attribution 4.0 Unported (CC BY 4.0) license, which permits others to copy, redistribute, remix, transform and build upon this work for any purpose, provided the original work is properly cited, a link to the licence is given, and indication of whether changes were made. See: <https://creativecommons.org/licenses/by/4.0/>.

ORCID iD

Iain James Marshall <http://orcid.org/0000-0003-2594-2654>

REFERENCES

- 1 Kyu HH, Abate D, Abate KH, *et al*. Global, regional, and national disability-adjusted life-years (DALYs) for 359 diseases and injuries and healthy life expectancy (HALE) for 195 countries and territories, 1990-2017: a systematic analysis for the global burden of disease study 2017. *Lancet* 2018;392:1859-922.
- 2 Chalmers I. The Cochrane collaboration: preparing, maintaining, and disseminating systematic reviews of the effects of health care. *Ann N Y Acad Sci* 1993;703:156-65.
- 3 Welch V, Petkovic J, Jull J, *et al*. Chapter 16: equity and specific populations. version 6.0. Cochrane Handbook for systematic reviews of interventions. version 6.0, 2019. Available: <http://www.training.cochrane.org/handbook>
- 4 Trouiller P, Olliaro P, Torreele E, *et al*. Drug development for neglected diseases: a deficient market and a public-health policy failure. *The Lancet* 2002;359:2188-94.
- 5 Gross CP, Anderson GF, Powe NR. The relation between funding by the National Institutes of Health and the burden of disease. *N Engl J Med* 1999;340:1881-7.
- 6 Luengo-Fernandez R, Leal J, Gray AM. UK research expenditure on dementia, heart disease, stroke and cancer: are levels of spending related to disease burden? *Eur J Neurol* 2012;19:149-54.
- 7 Moses H, Matheson DHM, Cairns-Smith S, *et al*. The anatomy of medical research: US and international comparisons. *JAMA* 2015;313:174-89.

- 8 Gillum LA, Gouveia C, Dorsey ER, *et al*. Nih disease funding levels and burden of disease. *PLoS One* 2011;6:e16837.
- 9 Evans JA, Shim J-M, Ioannidis JPA. Attention to local health burden and the global disparity of health research. *PLoS One* 2014;9:e90147.
- 10 Atal I, Trinquart L, Ravaud P, *et al*. A mapping of 115,000 randomized trials revealed a mismatch between research effort and health needs in non-high-income regions. *J Clin Epidemiol* 2018;98:123–32.
- 11 Marshall IJ, Nye B, Kuiper J, *et al*. Trialstreamer: a living, automatically updated database of clinical trial reports. *J Am Med Inform Assoc* 2020;ocaa163.
- 12 Marshall IJ, Noel-Storr A, Kuiper J, *et al*. Machine learning for identifying randomized controlled trials: an evaluation and practitioner's guide. *Res Synth Methods* 2018;9:602–14.
- 13 Joachims T. Text categorization with Support Vector Machines: Learning with many relevant features. In: Nédellec C, Rouveiol C, eds. *Machine learning: ECML-98*. Berlin, Heidelberg: Springer, 1998: 137–42.
- 14 Kim Y. Convolutional neural networks for sentence classification. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Doha, Qatar, 2014. p.:1746–51.
- 15 Hochreiter S, Schmidhuber J. Long short-term memory. *Neural Comput* 1997;9:1735–80.
- 16 Sutton C, McCallum A. An introduction to conditional random fields. now publishers, 2012. Available: <https://books.google.co.uk/books?id=Ve6HMAEACAAJ>
- 17 Lample G, Ballesteros M, Subramanian S, *et al*. Neural architectures for named entity recognition. *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2016. p.:260–70.
- 18 Nye B, Jessy Li J, Patel R, *et al*. A corpus with multi-level annotations of patients, interventions and outcomes to support language processing for medical literature. *Proc Conf Assoc Comput Linguist Meet* 2018;2018:197–207.
- 19 Demner-Fushman D, Rogers WJ, Aronson AR. MetaMap Lite: an evaluation of a new Java implementation of MetaMap. *J Am Med Inform Assoc* 2017;24:841–4.
- 20 Aronson AR, Lang F-M. An overview of MetaMap: historical perspective and recent advances. *J Am Med Inform Assoc* 2010;17:229–36.
- 21 Fox MP, Lash TL, Greenland S. A method to automate probabilistic sensitivity analyses of misclassified binary variables. *Int J Epidemiol* 2005;34:1370–6.
- 22 Lash TL, Fink AK. Semi-Automated sensitivity analysis to assess systematic errors in observational data. *Epidemiology* 2003;14:451–8.
- 23 Steyerberg EW, Harrell FE, Borsboom GJ, *et al*. Internal validation of predictive models: efficiency of some procedures for logistic regression analysis. *J Clin Epidemiol* 2001;54:774–81.
- 24 Troeger C, Blacker B, Khalil IA, *et al*. Estimates of the global, regional, and national morbidity, mortality, and aetiologies of lower respiratory infections in 195 countries, 1990–2016: a systematic analysis for the global burden of disease study 2016. *Lancet Infect Dis* 2018;18:1191–210.
- 25 Torre LA, Bray F, Siegel RL, *et al*. Global cancer statistics, 2012. *CA Cancer J Clin* 2015;65:87–108.
- 26 King H, Aubert RE, Herman WH. Global burden of diabetes, 1995–2025: prevalence, numerical estimates, and projections. *Diabetes Care* 1998;21:1414–31.
- 27 Hackshaw A. Small studies: strengths and limitations. *Eur Respir J* 2008;32:1141–3.
- 28 Ioannidis JPA. Why most published research findings are false. *PLoS Med* 2005;2:e124.
- 29 Zhang Z, Xu X, Ni H. Small studies may overestimate the effect sizes in critical care meta-analyses: a meta-epidemiological study. *Crit Care* 2013;17:R2.
- 30 Nüesch E, Trelle S, Reichenbach S, *et al*. Small study effects in meta-analyses of osteoarthritis trials: meta-epidemiological study. *BMJ* 2010;341:c3515.
- 31 Yusuf S, Collins R, Peto R. Why do we need some large, simple randomized trials? *Stat Med* 1984;3:409–20.
- 32 Ebrahim S, Montoya L, Kamal el Din M, *et al*. Randomized trials are frequently fragmented in multiple secondary publications. *J Clin Epidemiol* 2016;79:130–9.
- 33 Moher D, Jones A, Lepage L, *et al*. Use of the CONSORT statement and quality of reports of randomized trials: a comparative before-and-after evaluation. *JAMA* 2001;285:1992–5.
- 34 Halladay CW, Trikalinos TA, Schmid IT, *et al*. Using data sources beyond PubMed has a modest impact on the results of systematic reviews of therapeutic interventions. *J Clin Epidemiol* 2015;68:1076–84.
- 35 Marshall IJ, Marshall R, Wallace BC, *et al*. Rapid reviews may produce different results to systematic reviews: a meta-epidemiological study. *J Clin Epidemiol* 2019;109:30–41.
- 36 Zon R, Meropol NJ, Catalano RB, *et al*. American Society of clinical oncology statement on minimum standards and exemplary attributes of clinical trial sites. *J Clin Oncol* 2008;26:2562–7.
- 37 Treweek S, Pitkethly M, Cook J, *et al*. Strategies to improve recruitment to randomised trials. *Cochrane Database Syst Rev* 2018;2:MR000013.
- 38 Anand S, Hanson K. Disability-adjusted life years: a critical review. *J Health Econ* 1997;16:685–702.
- 39 HIV-CAUSAL Collaboration, Ray M, Logan R, *et al*. The effect of combined antiretroviral therapy on the overall mortality of HIV-infected individuals. *AIDS* 2010;24:123–37.
- 40 Hidalgo M. Pancreatic cancer. *N Engl J Med* 2010;362:1605–17.
- 41 Waring MJ, Arrowsmith J, Leach AR, *et al*. An analysis of the attrition of drug candidates from four major pharmaceutical companies. *Nat Rev Drug Discov* 2015;14:475–86.