Etica della ricerca scientifica

Etica della ricerca sull’uomo

Etica della sperimentazione sugli animali

Etica della ricerca

Le responsabilità della scienza

Etica deontologica:

la correttezza etica di un comportamento è assoluta e inderogabile

Etica normativa basata sulle regole

Etica utilitarista:

La migliore azione è quella che massimizza l'utilità (p.es. il bene degli esseri umani, oppure di tutti gli esseri senzienti)

Etica basata sul peso delle conseguenze

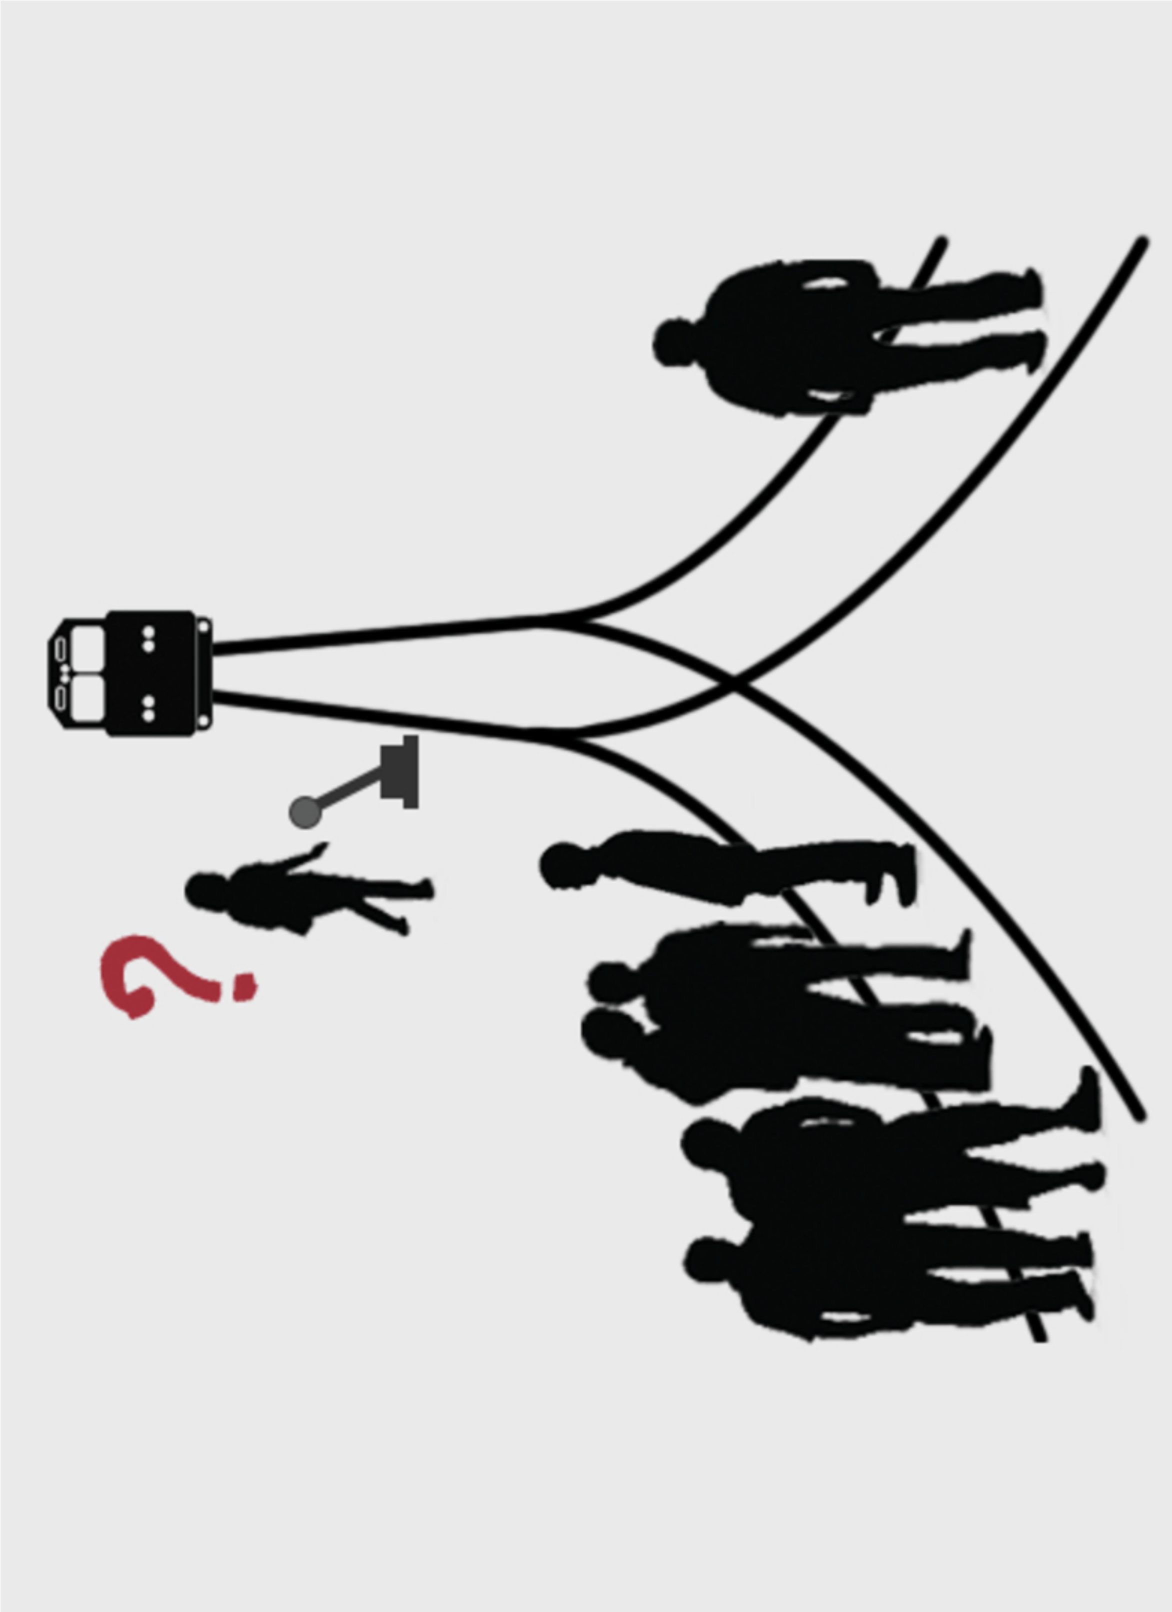

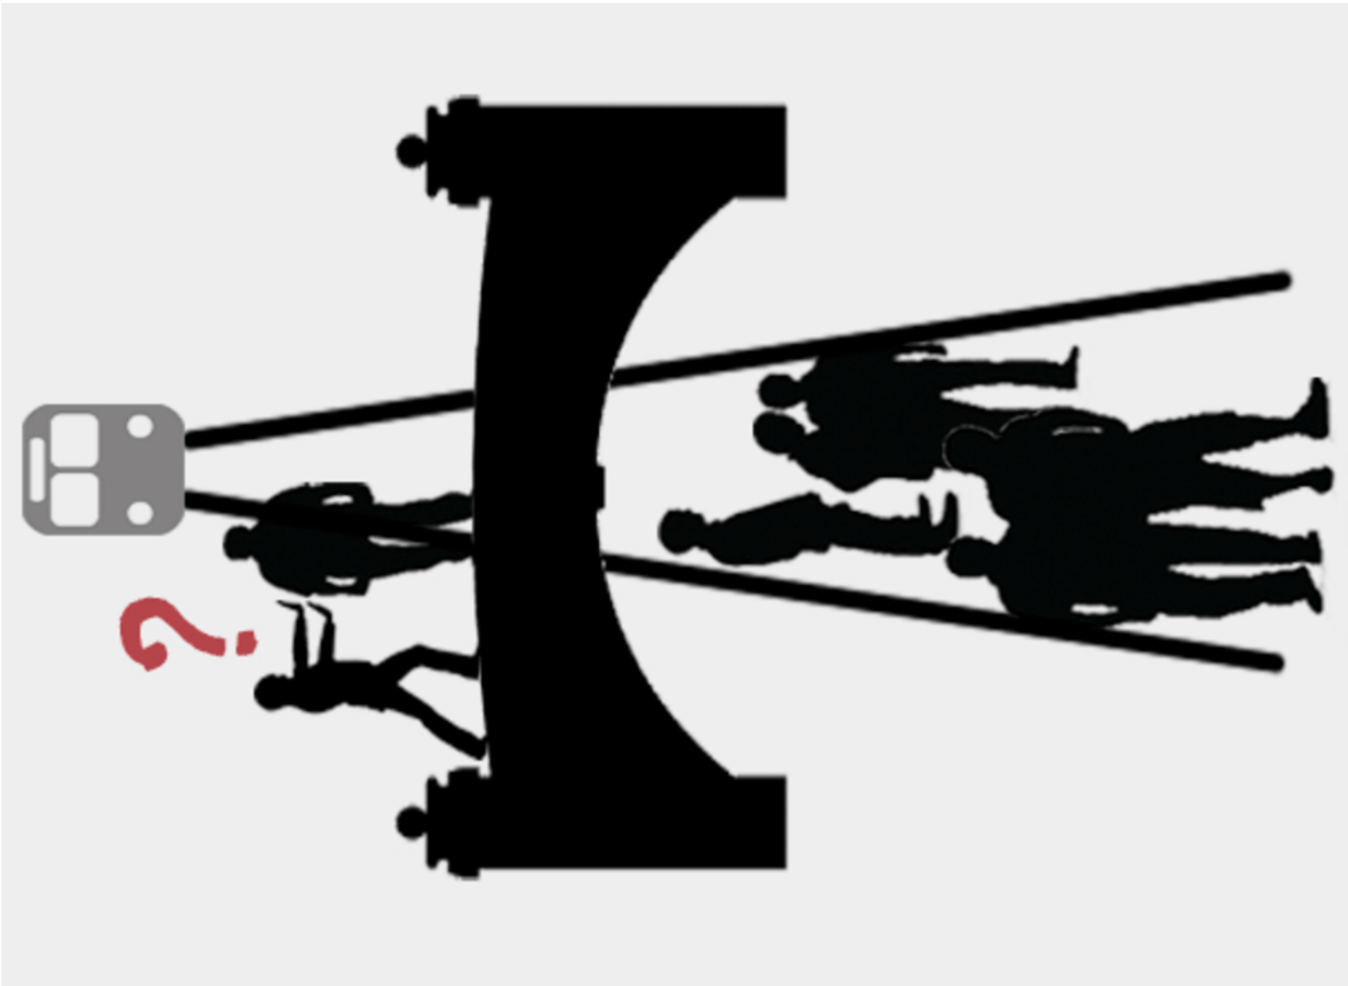

Etica della ricerca sull'uomo


LPDC

Ricerca sull’uomo

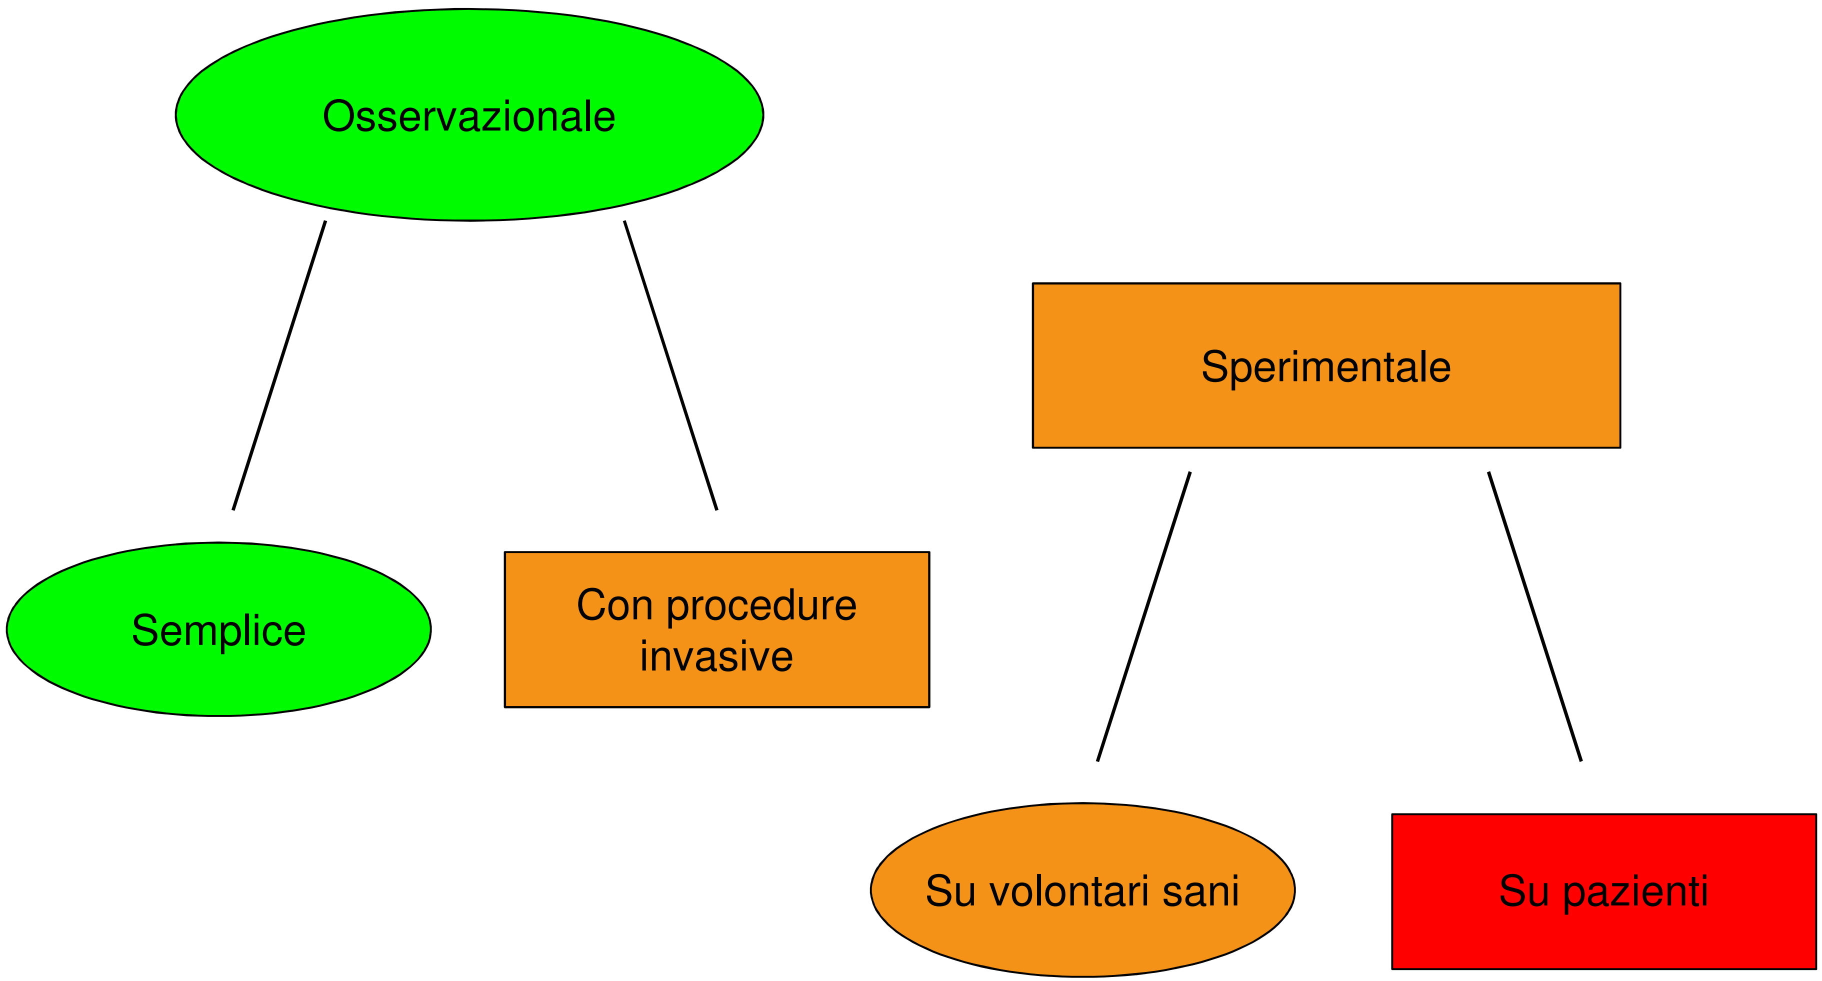

Etica della sperimentazione sugli animali

Decreto legislativo n. 26 del 2014

Art. 1

Oggetto e ambito di applicazione

1. Il presente decreto stabilisce misure relative alla protezione degli animali utilizzati ai fini scientifici o educativi, a tal fine, sono disciplinati i seguenti aspetti:

a) la sostituzione, la riduzione dell'uso di animali nelle procedure e il perfezionamento delle tecniche di allevamento, di alloggiamento, di cura e di impiego degli animali nelle procedure;

b) la provenienza, l'allevamento, l'identificazione, la cura, l'alloggiamento e la soppressione degli animali;

c) le attivita' degli allevatori, dei fornitori e degli utilizzatori;

d) la valutazione e l'autorizzazione dei progetti che prevedono l'uso degli animali nelle procedure.

2. E' consentito l'utilizzo degli animali ai fini scientifici o educativi soltanto quando, per ottenere il risultato ricercato, non sia possibile utilizzare altro metodo o una strategia di sperimentazione scientificamente valida, ragionevolmente e praticamente applicabile che non implichi l'impiego di animali vivi.

3. Il presente decreto si applica ai seguenti animali:

a) animali vertebrati vivi non umani, comprese:

1) forme larvali capaci di alimentarsi autonomamente;

2) forme fetali di mammiferi a partire dall'ultimo terzo del loro normale sviluppo;

b) cefalopodi vivi.

4. Il presente decreto si applica agli animali:

a) utilizzati o destinati a essere utilizzati nelle procedure, o appositamente allevati affinche' i loro organi o tessuti possano essere usati ai fini scientifici, anche se si trovano in una fase di sviluppo precedente a quella di cui al comma 3, lettera a), e se l'animale viene fatto vivere oltre detta fase di sviluppo ed e' probabile che, a seguito delle procedure effettuate, provi dolore,

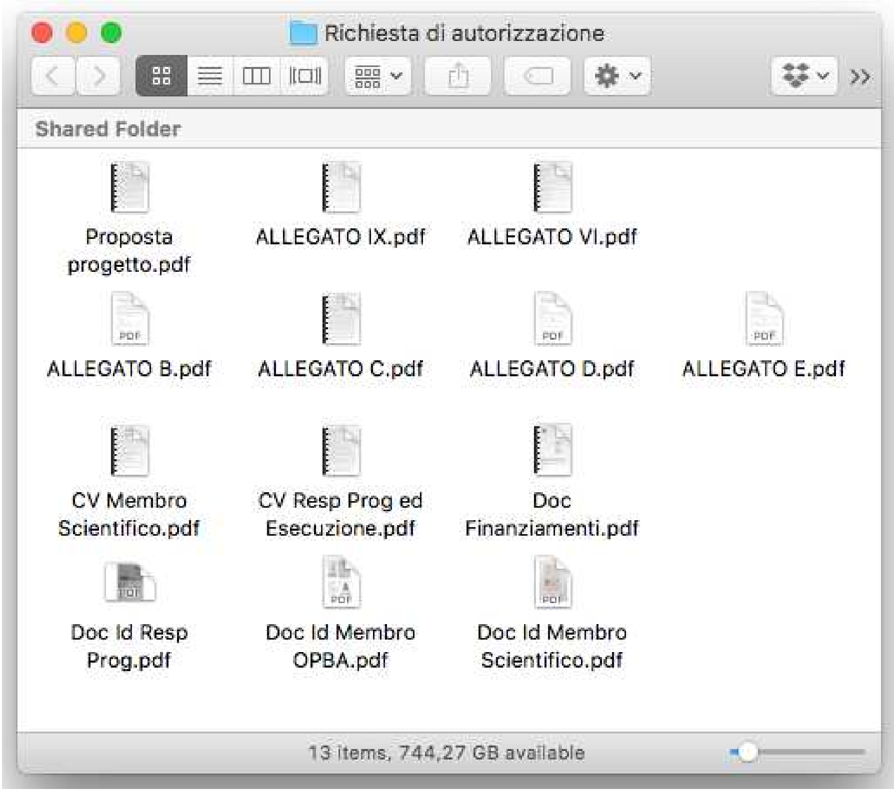

Etica della ricerca

Gli scienziati lavorano per l’intero genere umano

Gli scienziati lavorano per l’intero genere umano,

perché le loro scoperte si diffondono in tutto il mondo

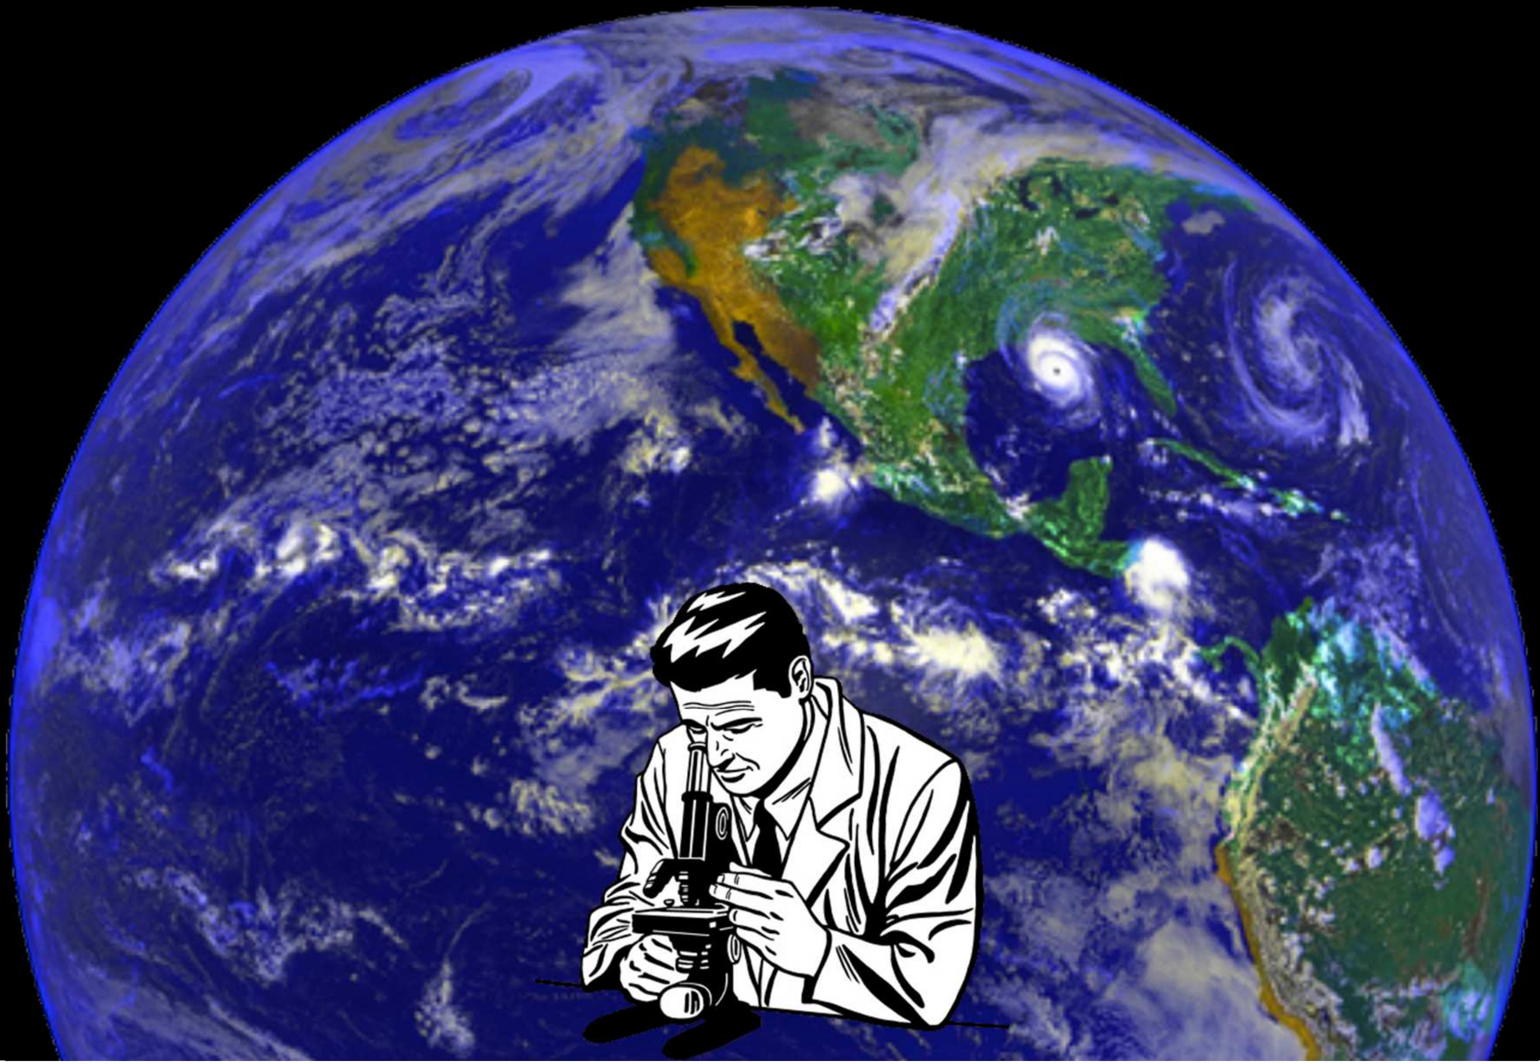

In nessun caso è permesso loro di alterare dati e risultati

- Tendenziosità
- Soppressione parziale di dati
- Alterazione arbitraria di dati
- Fabbricazione di dati

Gli scienziati devono cooperare con i loro colleghi

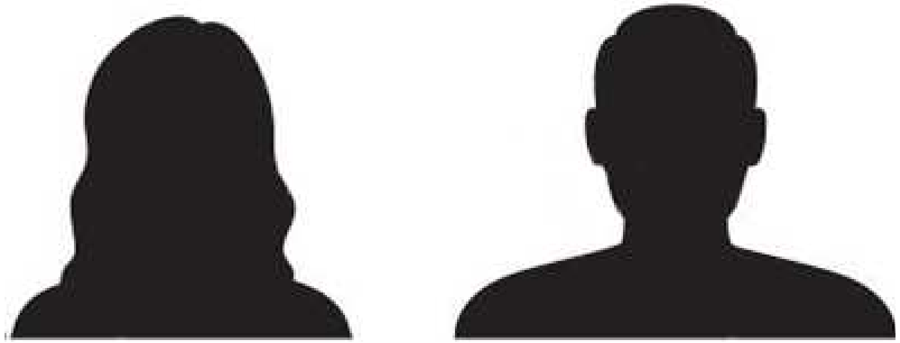

Referees

Fare ricerca è un lavoro

Come tutti i lavori, deve essere remunerato

Gli scienziati devono formare i giovani

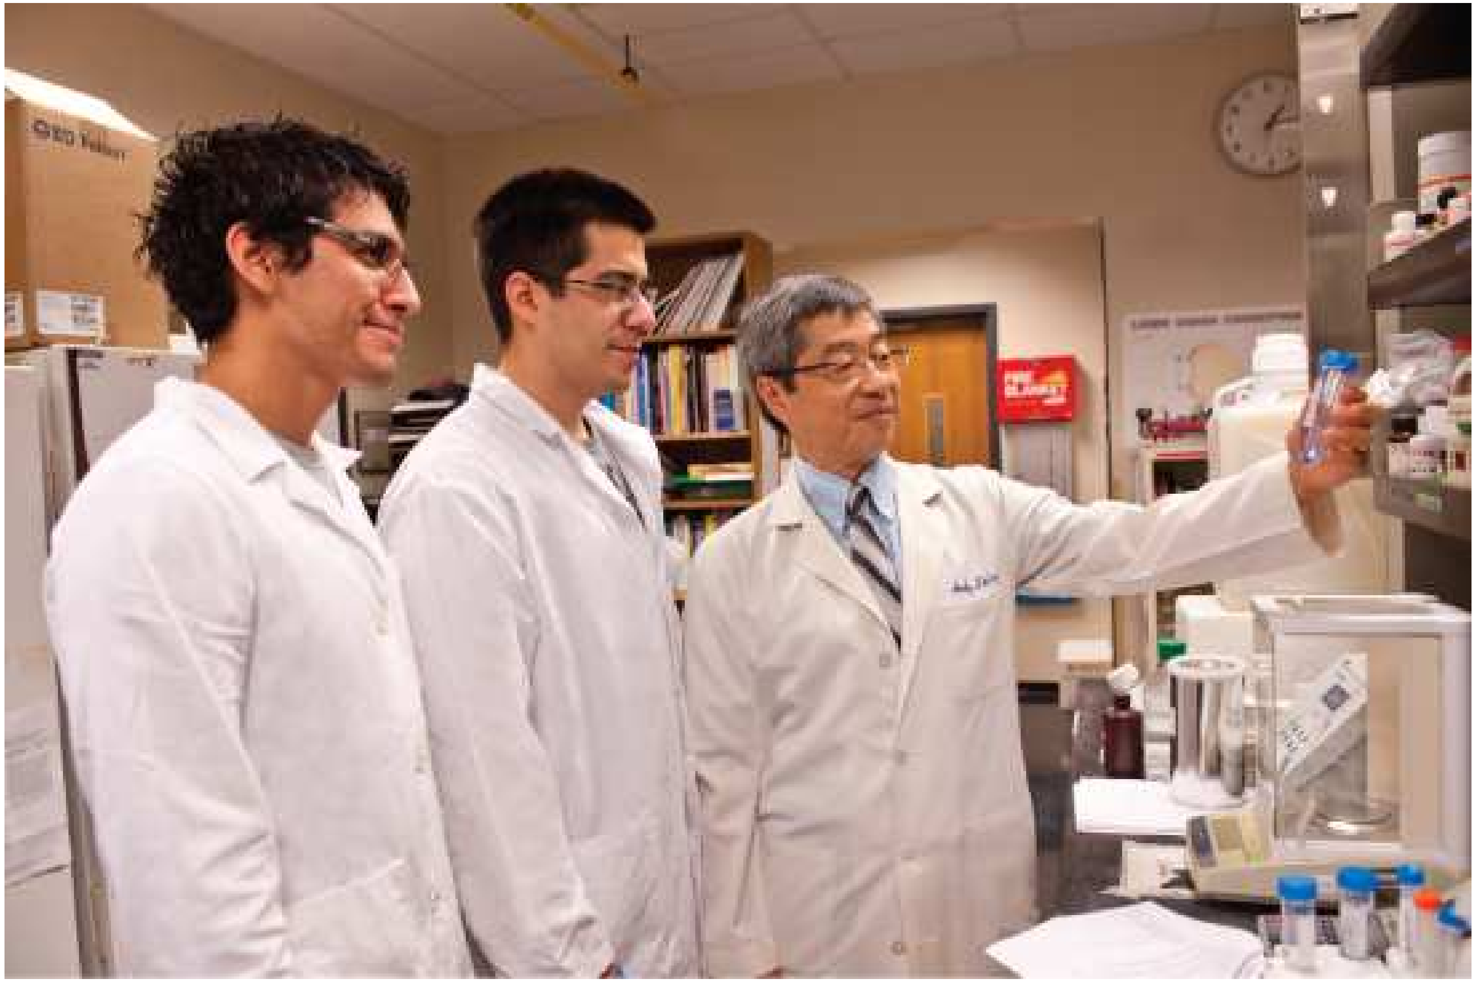

La responsabilità della scienza

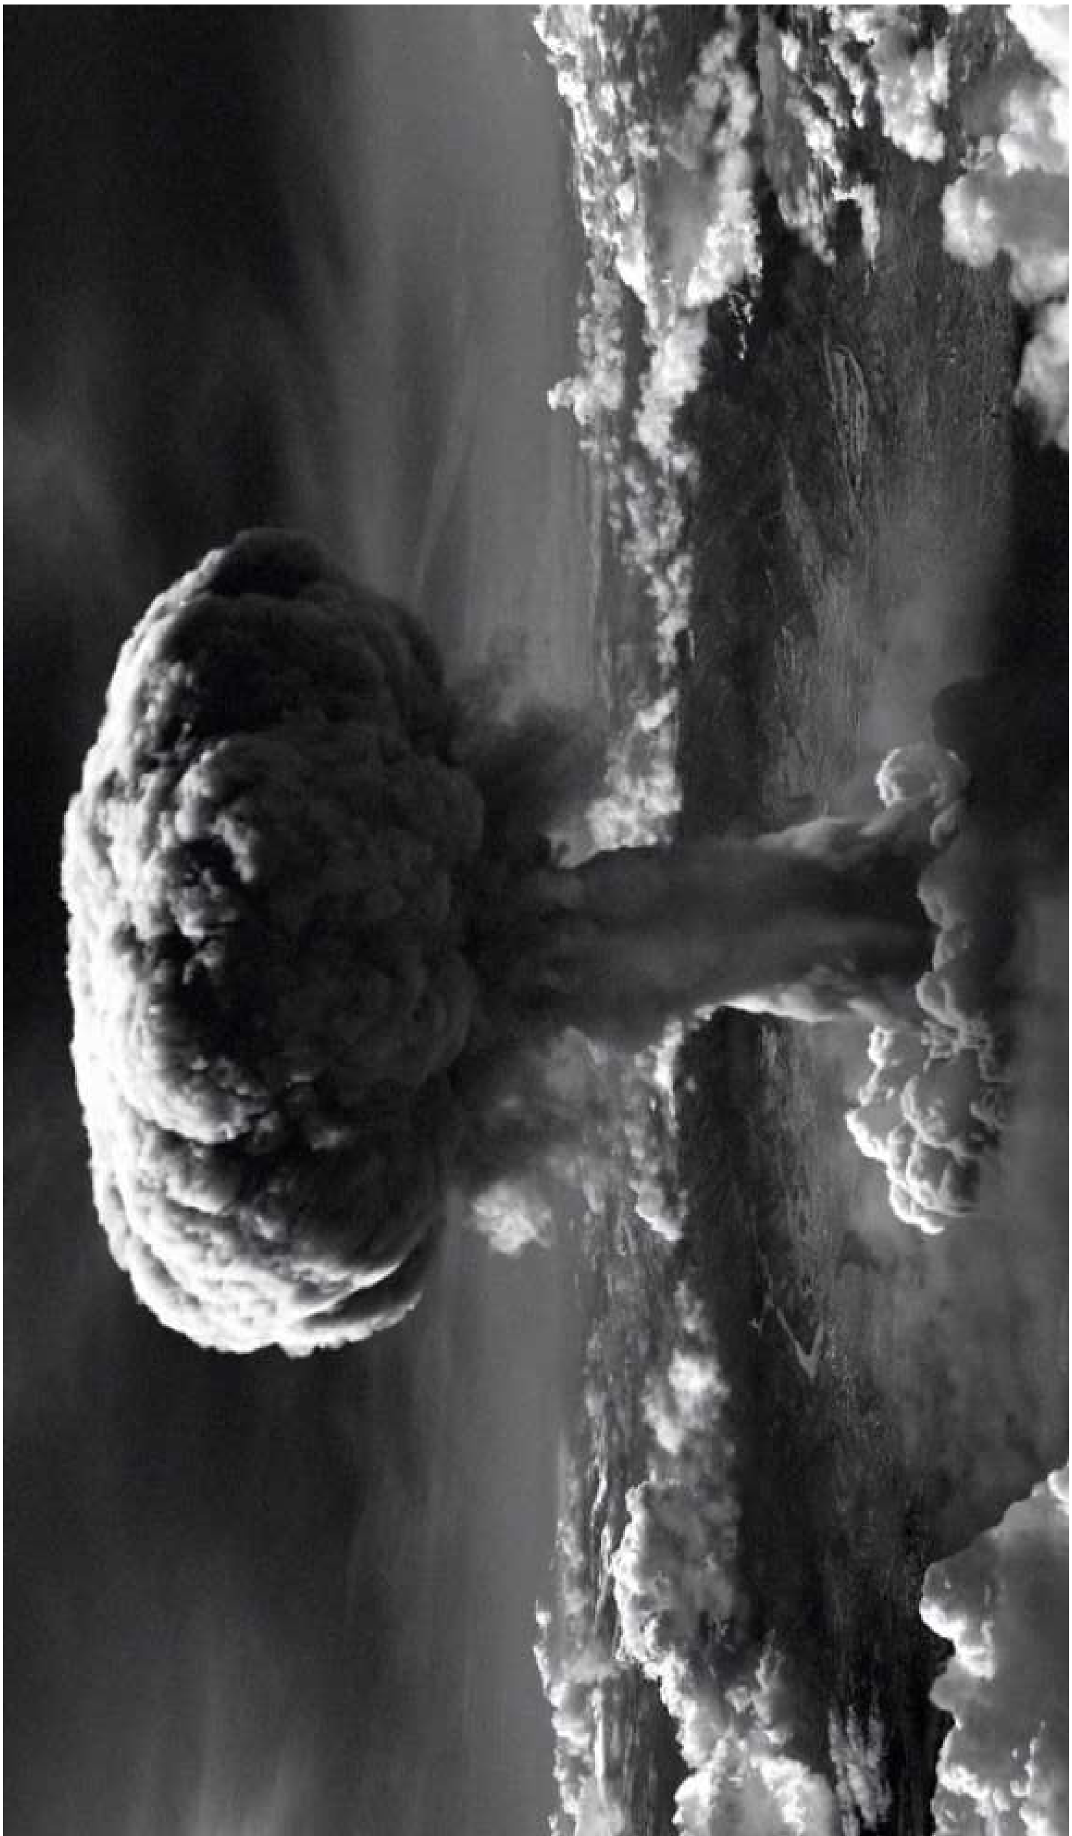

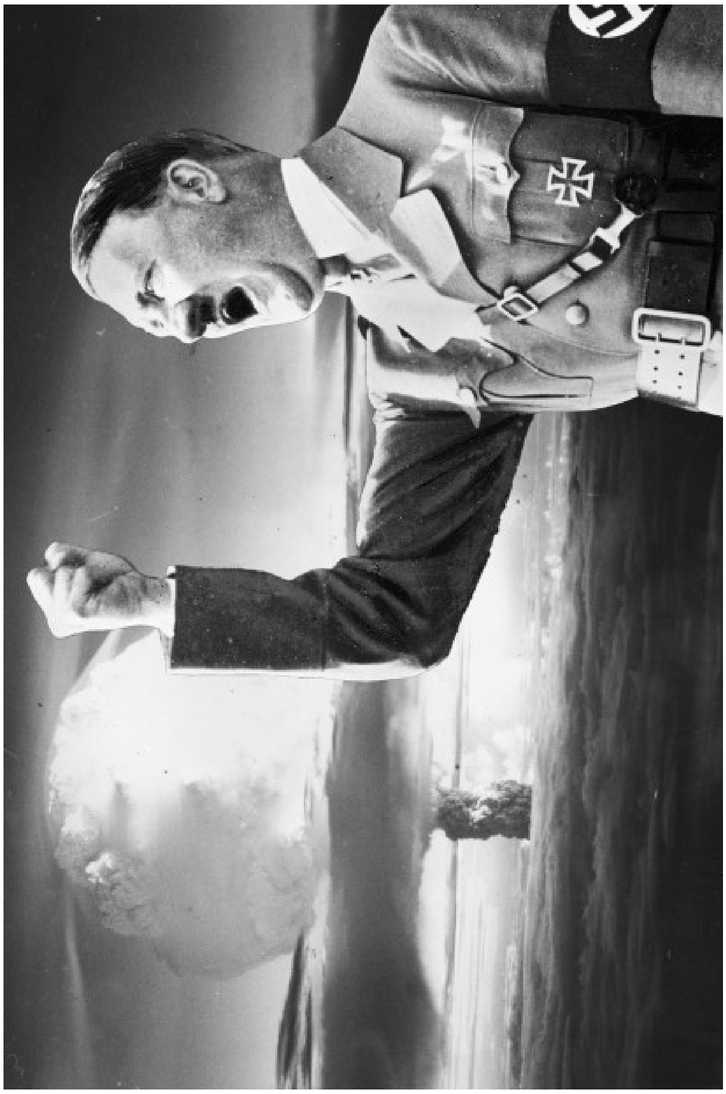

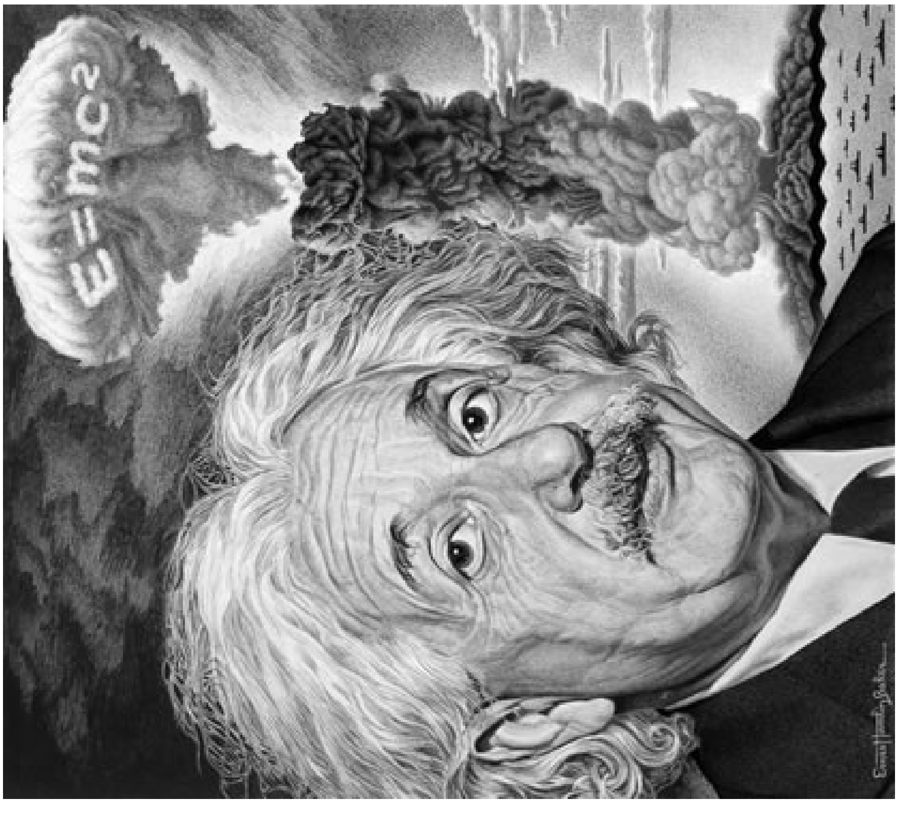
$E=mc^2$

Costituzione della Repubblica italiana

Art. 33.

L’arte e la scienza sono libere e libero ne è l’insegnamento.

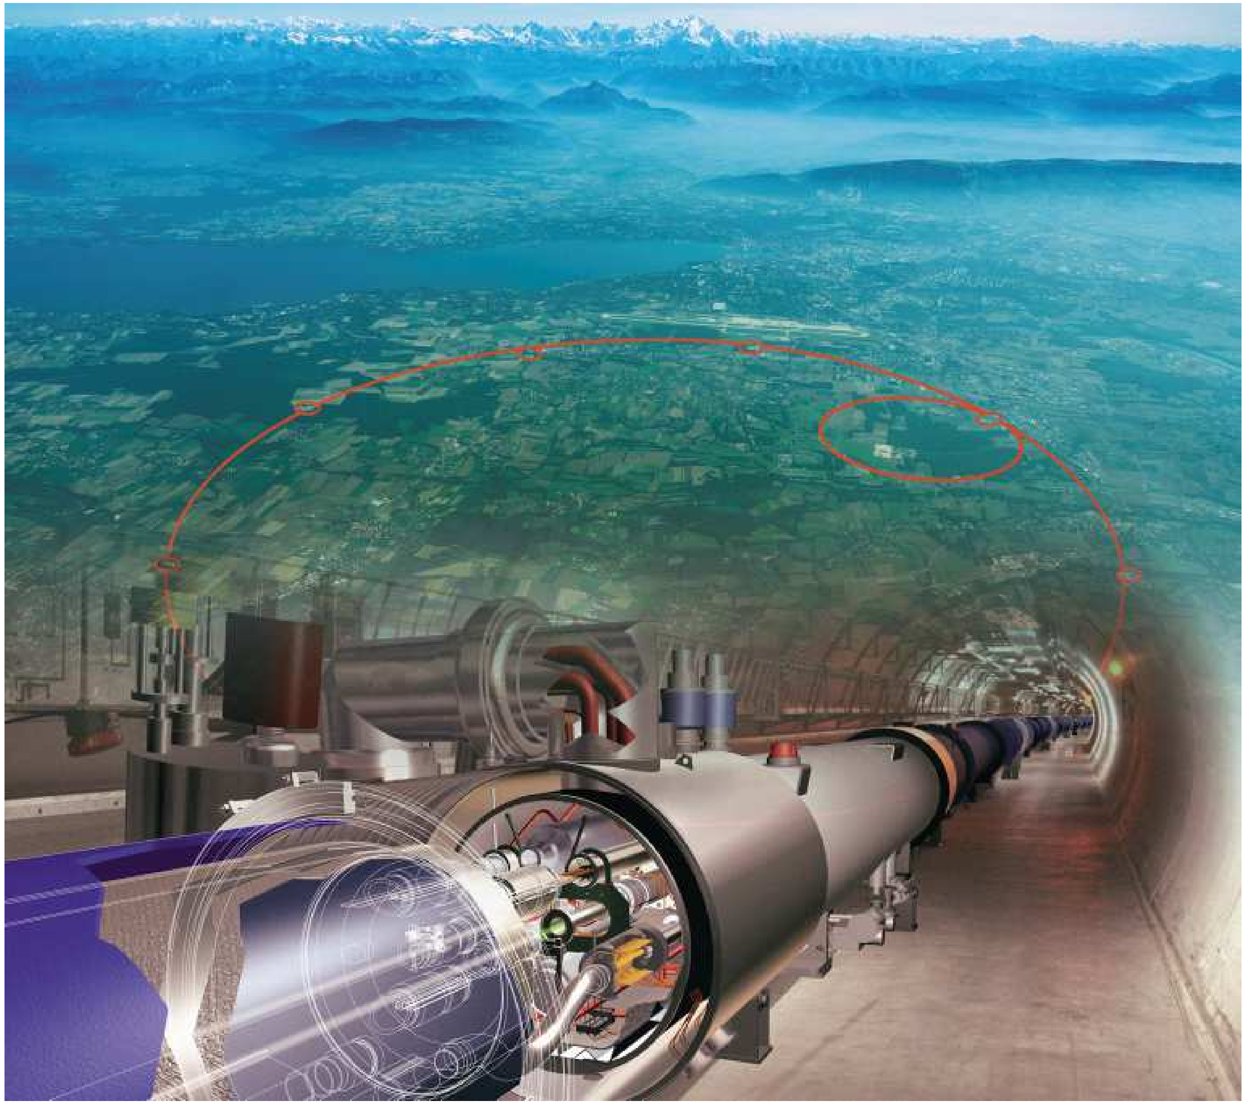

1999

Intelligenza artificiale

COMPUTER SCIENCE

A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play

David Silver[1,2*†], Thomas Hubert[1*], Julian Schrittwieser[1*], Ioannis Antonoglou[1], Matthew Lai[1], Arthur Guez[1], Marc Lanctot[1], Laurent Sifre[1], Dharshan Kumaran[1], Thore Graepel[1], Timothy Lillicrap[1], Karen Simonyan[1], Demis Hassabis[1†]

The game of chess is the longest-studied domain in the history of artificial intelligence. The strongest programs are based on a combination of sophisticated search techniques, domain-specific adaptations, and handcrafted evaluation functions that have been refined by human experts over several decades. By contrast, the AlphaGo Zero program recently achieved superhuman performance in the game of Go by reinforcement learning from self-play. In this paper, we generalize this approach into a single AlphaZero algorithm that can achieve superhuman performance in many challenging games. Starting from random play and given no domain knowledge except the game rules, AlphaZero convincingly defeated a world champion program in the games of chess and shogi (Japanese chess), as well as Go.

The study of computer chess is as old as computer science itself. Charles Babbage, Alan Turing, Claude Shannon, and John von Neumann devised hardware, algorithms, and theory to analyze and play the game of chess. Chess subsequently became a grand challenge task for a generation of artificial intelligence researchers, culminating in high-performance computer chess programs that play at a superhuman level (*1*, *2*). However, these systems are highly tuned to their domain and cannot be generalized to other games without substantial human effort, whereas general game-playing systems (*3*, *4*) remain comparatively weak.

A long-standing ambition of artificial intelligence has been to create programs that can instead learn for themselves from first principles (*5*, *6*). Recently, the AlphaGo Zero algorithm achieved superhuman performance in the game of Go by representing Go knowledge with the use of deep convolutional neural networks (*7*, *8*), trained solely by reinforcement learning from games of self-play (*9*). In this paper, we introduce AlphaZero, a more generic version of the AlphaGo Zero algorithm that accommodates, without special casing, a broader class of game rules. We apply AlphaZero to the games of chess and shogi, as well as Go, by using the same algorithm and network architecture for all three games. Our results demonstrate that a general-purpose reinforcement learning algorithm can learn, tabula rasa—without domain-specific human knowledge or data, as evidenced by the same algorithm succeeding in multiple domains—superhuman performance across multiple challenging games.

A landmark for artificial intelligence was achieved in 1997 when Deep Blue defeated the human world chess champion (*1*). Computer chess programs continued to progress steadily beyond human level in the following two decades. These programs evaluate positions by using handcrafted features and carefully tuned weights, constructed by strong human players and programmers, combined with a high-performance alpha-beta search that expands a vast search tree by using a large number of clever heuristics and domain-specific adaptations. In (*10*) we describe these augmentations, focusing on the 2016 Top Chess Engine Championship (TCEC) season 9 world champion Stockfish (*11*); other strong chess programs, including Deep Blue, use very similar architectures (*1*, *12*).

In terms of game tree complexity, shogi is a substantially harder game than chess (*13*, *14*): It is played on a larger board with a wider variety of pieces; any captured opponent piece switches sides and may subsequently be dropped anywhere on the board. The strongest shogi programs, such as the 2017 Computer Shogi Association (CSA) world champion Elmo, have only recently defeated human champions (*15*). These programs use an algorithm similar to those used by computer chess programs, again based on a highly optimized alpha-beta search engine with many domain-specific adaptations.

AlphaZero replaces the handcrafted knowledge and domain-specific augmentations used in traditional game-playing programs with deep neural networks, a general-purpose reinforcement learning algorithm, and a general-purpose tree search algorithm.

Instead of a handcrafted evaluation function and move-ordering heuristics, AlphaZero uses a deep neural network $(\mathbf{p}, v) = f_\theta(s)$ with parameters θ. This neural network $f_\theta(s)$ takes the board position s as an input and outputs a vector of move probabilities $\mathbf{p}$ with components $p_a = \Pr(a|s)$ for each action a and a scalar value v estimating the expected outcome z of the game from position s, $v \approx \mathbb{E}[z|s]$. AlphaZero learns these move probabilities and value estimates entirely from self-play; these are then used to guide its search in future games.

Instead of an alpha-beta search with domain-specific enhancements, AlphaZero uses a general-purpose Monte Carlo tree search (MCTS) algorithm. Each search consists of a series of simulated games of self-play that traverse a tree from root state s_{root} until a leaf state is reached. Each simulation proceeds by selecting in each state s a move a with low visit count (not previously frequently explored), high move probability, and high value (averaged over the leaf states of

[1]DeepMind, 6 Pancras Square, London N1C 4AG, UK. [2]University College London, Gower Street, London WC1E 6BT, UK.
*These authors contributed equally to this work.
†Corresponding author. Email: davidsilver@google.com (D.S.); dhcontact@google.com (D.H.)

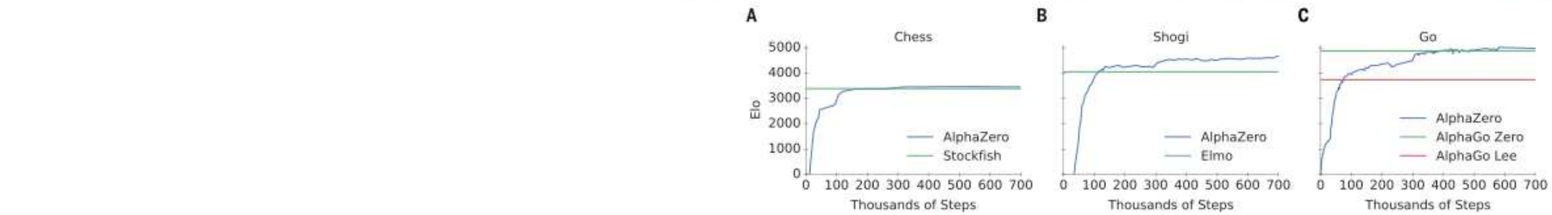

Fig. 1. Training AlphaZero for 700,000 steps. Elo ratings were computed from games between different players where each player was given 1 s per move. (**A**) Performance of AlphaZero in chess compared with the 2016 TCEC world champion program Stockfish. (**B**) Performance of AlphaZero in shogi compared with the 2017 CSA world champion program Elmo. (**C**) Performance of AlphaZero in Go compared with AlphaGo Lee and AlphaGo Zero (20 blocks over 3 days).

A graph placement methodology for fast chip design

https://doi.org/10.1038/s41586-021-03544-w

Received: 3 November 2020

Accepted: 13 April 2021

Published online: 9 June 2021

Check for updates

Azalia Mirhoseini[1,4✉], Anna Goldie[1,3,4✉], Mustafa Yazgan[2], Joe Wenjie Jiang[1], Ebrahim Songhori[1], Shen Wang[1], Young-Joon Lee[2], Eric Johnson[1], Omkar Pathak[2], Azade Nazi[1], Jiwoo Pak[2], Andy Tong[2], Kavya Srinivasa[2], William Hang[3], Emre Tuncer[2], Quoc V. Le[1], James Laudon[1], Richard Ho[2], Roger Carpenter[2] & Jeff Dean[1]

Chip floorplanning is the engineering task of designing the physical layout of a computer chip. Despite five decades of research[1], chip floorplanning has defied automation, requiring months of intense effort by physical design engineers to produce manufacturable layouts. Here we present a deep reinforcement learning approach to chip floorplanning. In under six hours, our method automatically generates chip floorplans that are superior or comparable to those produced by humans in all key metrics, including power consumption, performance and chip area. To achieve this, we pose chip floorplanning as a reinforcement learning problem, and develop an edge-based graph convolutional neural network architecture capable of learning rich and transferable representations of the chip. As a result, our method utilizes past experience to become better and faster at solving new instances of the problem, allowing chip design to be performed by artificial agents with more experience than any human designer. Our method was used to design the next generation of Google's artificial intelligence (AI) accelerators, and has the potential to save thousands of hours of human effort for each new generation. Finally, we believe that more powerful AI-designed hardware will fuel advances in AI, creating a symbiotic relationship between the two fields.

In this work, we propose a new graph placement method based on reinforcement learning (RL), and demonstrate state-of-the-art results on chip floorplanning, a challenging problem[2] that has long defied automation, despite five decades of research. Our method generates manufacturable chip floorplans in under 6 h, compared to the strongest baseline, which requires months of intense effort by human experts.

A computer chip is divided into dozens of blocks, each of which is an individual module, such as a memory subsystem, compute unit or control logic system. These blocks can be described by a netlist, a hypergraph of circuit components, such as macros (memory components) and standard cells (logic gates such as NAND, NOR and XOR), all of which are connected by wires. Chip floorplanning involves placing netlists onto chip canvases (two-dimensional grids) so that performance metrics (for example, power consumption, timing, area and wirelength) are optimized, while adhering to hard constraints on density and routing congestion.

Since the 1960s, many approaches to chip floorplanning have been proposed, falling into three broad categories: partitioning-based methods[3–5], stochastic/hill-climbing approaches[6–8] and analytic solvers[9–14]. However, none of these approaches could achieve human-level performance, and the exponential growth in chip complexity has rendered these techniques largely unusable on modern chips.

The limitations of these prior approaches are varied. For example, partitioning-based methods sacrifice the quality of the global solution in order to scale to larger netlists, and a poor early partition may result in an unsalvageable final result. Hill-climbing approaches have low convergence rates and do not scale to modern chip netlists, which have millions or billions of nodes[13]. Prior to this work, analytic solvers were the leading approach, but they can only optimize for differentiable loss functions, meaning that they cannot effectively optimize for critical metrics, such as routing congestion or timing violations. Our method, on the other hand, can scale to netlists with millions of nodes, and optimizes directly for any mixture of differentiable or non-differentiable cost functions. Furthermore, our method improves in both speed and quality of result because it is exposed to more instances of the chip placement problem.

Owing to the limitations of these prior methods, human physical designers must iterate for months with commercial electronic design automation (EDA) tools, taking as input a register transfer level (RTL) description of the chip netlist, generating a manual placement of that netlist onto the chip canvas, and waiting up to 72 h for EDA tools to evaluate that placement. On the basis of this feedback, the human designer either concludes that the design criteria have been achieved, generates an updated floorplan for evaluation, or provides feedback to upstream RTL designers, who then modify the low-level code to make the placement task easier (for example, resolve timing violations).

To address the chip floorplanning problem, we developed an RL method capable of generalizing across chips—meaning that it can learn from experience to become both better and faster at placing

[1]Google Research, Brain Team, Google, Mountain View, CA, USA. [2]Google Chip Implementation and Infrastructure (CI2) Team, Google, Sunnyvale, CA, USA. [3]Computer Science Department, Stanford University, Stanford, CA, USA. [4]These authors contributed equally: Azalia Mirhoseini, Anna Goldie. ✉e-mail: azalia@google.com; agoldie@google.com

H5N1

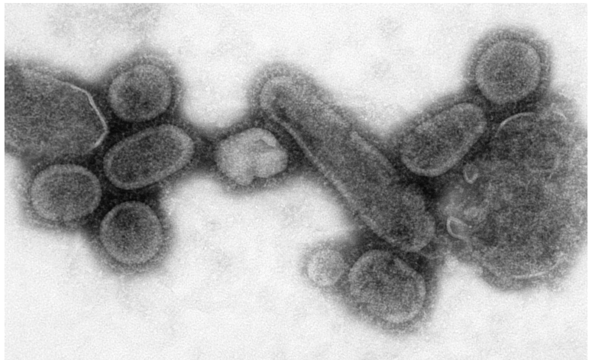

REPORT

Airborne Transmission of Influenza A/H5N1 Virus Between Ferrets

Sander Herfst,[1] Eefje J. A. Schrauwen,[1] Martin Linster,[1] Salin Chutinimitkul,[1] Emmie de Wit,[1]* Vincent J. Munster,[1]* Erin M. Sorrell,[1] Theo M. Bestebroer,[1] David F. Burke,[2] Derek J. Smith,[1,2,3] Guus F. Rimmelzwaan,[1] Albert D. M. E. Osterhaus,[1] Ron A. M. Fouchier[1]†

Highly pathogenic avian influenza A/H5N1 virus can cause morbidity and mortality in humans but thus far has not acquired the ability to be transmitted by aerosol or respiratory droplet ("airborne transmission") between humans. To address the concern that the virus could acquire this ability under natural conditions, we genetically modified A/H5N1 virus by site-directed mutagenesis and subsequent serial passage in ferrets. The genetically modified A/H5N1 virus acquired mutations during passage in ferrets, ultimately becoming airborne transmissible in ferrets. None of the recipient ferrets died after airborne infection with the mutant A/H5N1 viruses. Four amino acid substitutions in the host receptor-binding protein hemagglutinin, and one in the polymerase complex protein basic polymerase 2, were consistently present in airborne-transmitted viruses. The transmissible viruses were sensitive to the antiviral drug oseltamivir and reacted well with antisera raised against H5 influenza vaccine strains. Thus, avian A/H5N1 influenza viruses can acquire the capacity for airborne transmission between mammals without recombination in an intermediate host and therefore constitute a risk for human pandemic influenza.

Science 336:1534-1541, 2012

Modifiche alla linea germinale umana

Biology's last taboo

P. PLAILLY/EURELIOS/SPL

Will gene therapy ever extend to inducing changes in humans that can be inherited down through generations? Maybe so, if the concerns over safety can be ironed out. Jonathan Knight considers the technical challenges and the ethical arguments.

For a few days in January, a young rhesus monkey called ANDi became one of the most famous animals on the planet, his endearing features plastered across countless newspapers. Engineered by a team at the Oregon Regional Primate Research Center near Portland so that each of his cells carries a gene for a glowing green protein, ANDi is the world's first transgenic primate[1]. His name derives from 'inserted DNA' spelt backwards.

Some reports suggested that ANDi is just one step removed from 'designer' human babies. But if anything, he shows that transgenic people are still far from a practical reality. ANDi gives out not even the faintest green glow — the introduced gene was active in two other fetuses that aborted spontaneously, but in ANDi it seems to lie silent. Nor do ANDi's creators hope to develop a method for use in people. Rather, they are interested in creating genetically modified monkeys for use in medical research. "We discourage any extrapolations to people," says Gerald Schatten, the team's leader, now at the University of Pittsburgh.

But when asked to look a decade or more into the future, many experts believe that 'germline' gene therapy, which aims to produce babies carrying altered genes that will be inherited down the generations, might become part of mainstream medicine. Given advances in basic research, they argue, the safety concerns that currently put germline gene therapy off-limits might be overcome.

Objections to human germline manipulation run deeper than the safety issues. For some people, tampering with our genetic inheritance in this way is fundamentally wrong. Others fear that what starts as an attempt to cure inherited diseases will soon mutate into a programme of genetic 'improvement'. But previous experience suggests that such concerns may not restrain those intent on pushing back the boundaries of reproductive and genetic medicine (see 'The unregulated frontier', overleaf). "Altering the genome is essentially the endpoint of the whole genomic revolution," claims Gregory Stock of the University of California, Los Angeles, an expert on the societal impact of emerging genetic technologies.

Tackling the taboo

In March 1998, Stock organized a conference in Los Angeles that threw a spotlight on human germline manipulation, then considered a taboo subject. At the meeting, leading scientists spoke in favour of the idea, arguing that excessive regulation could hamper research of medical value[2]. Last year, a report from the American Association for the Advancement of Science dampened the enthusiasm, calling for a moratorium on attempts to change a person's genes in ways that could affect future generations[3]. But it left the door open: with safe methods, proper supervision and adequate public discussion, the report concluded that germline gene therapy might one day be acceptable.

Any future attempt at human germline gene therapy will draw on experience with transgenic animals. The most basic procedure, injecting DNA fragments into one of the unfused nuclei in a newly fertilized egg, was pioneered in mice more than two decades ago[4]. The

NEWSMAKERS

Monkey business: ANDi, the first genetically modified primate, was born containing a jellyfish gene.

…la storia dell'umanita' insegna che una capacita' tecnologica acquisita non rimane a lungo inutilizzata. Al contrario qualcuno, in qualche luogo, la metterà in pratica con uno scopo od un altro, legalmente o contro la legge e il volere comune. L'uomo non si è lasciato fermare dal timore di produrre, con le sue invenzioni e la sua tecnologia, disastri ecologici o armi di sterminio. Né è riuscito realmente ad astenersi, come pure si era solennemente ripromesso, dal colonizzare l'Antartide o lo spazio extraterrestre a fini di potere. C'è sempre un interesse forte, una ragion di Stato, una nobile causa o una forza sociale inarrestabile che porta a scoperchiare il vaso di Pandora. È bene quindi che cominciamo a dibattere l'ingegnerizzazione genetica dell'uomo, al fine di sceverarne tutte le implicazioni, per poterla meglio controllare e indirizzare o, se così decidessimo, per cercare di metterla al bando con le massime probabilità di successo. Nascondere la testa nella sabbia per non vedere il futuro che ci sta venendo incontro non allontanerà i pericoli: ci farà solo sorprendere impreparati.

2003

Protein Cell 2015, 6(5):363–372
DOI 10.1007/s13238-015-0153-5

CrossMark Protein & Cell

RESEARCH ARTICLE

CRISPR/Cas9-mediated gene editing in human tripronuclear zygotes

Puping Liang, Yanwen Xu, Xiya Zhang, Chenhui Ding, Rui Huang, Zhen Zhang, Jie Lv, Xiaowei Xie, Yuxi Chen, Yujing Li, Ying Sun, Yaofu Bai, Zhou Songyang, Wenbin Ma, Canquan Zhou✉, Junjiu Huang✉

Guangdong Province Key Laboratory of Reproductive Medicine, the First Affiliated Hospital, and Key Laboratory of Gene Engineering of the Ministry of Education, School of Life Sciences, Sun Yat-sen University, Guangzhou 510275, China

✉ Correspondence: hjunjiu@mail.sysu.edu.cn (J. Huang), zhoucanquan@hotmail.com (C. Zhou)

Received March 30, 2015 Accepted April 1, 2015

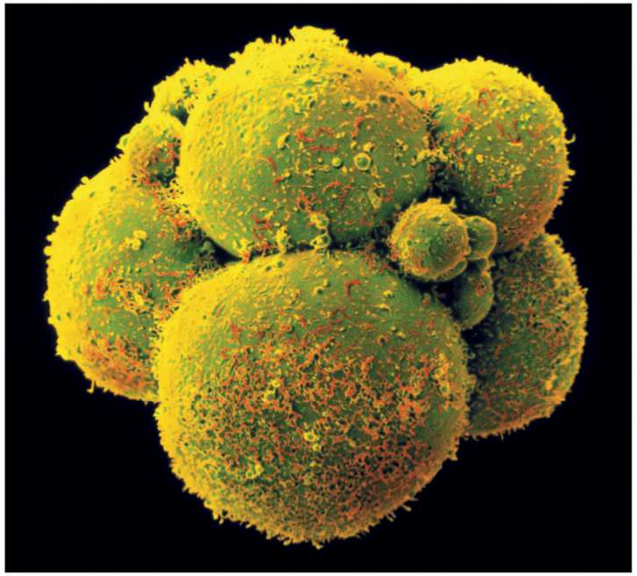

YORGOS NIKAS/SPL

Human embryos are at the centre of a debate over the ethics of gene editing.

GENETIC MODIFICATION

Embryo editing sparks epic debate

In wake of paper describing genetic modification of human embryos, scientists disagree about ethics.

BY DAVID CYRANOSKI AND SARA REARDON

In a world first, Chinese scientists have reported that they have used powerful gene-editing techniques to modify human embryos. Their paper[1], published in the Beijing-based journal *Protein & Cell* on 18 April, came as no surprise to the scientific community, but it has ignited a wide-ranging debate about what types of gene-editing research are ethical. The publication also raises questions about the appropriate way to publish such work (see 'Publishing policy').

In the paper, researchers led by Junjiu Huang, a gene-function researcher at Sun Yat-sen University in Guangzhou, describe how they used a system of molecules called CRISPR/Cas9, known for its ease of use, to cut DNA in human embryos and then attempted to repair it by introducing new DNA.

In a deliberate attempt to head off ethical concerns, the team used non-viable embryos obtained from fertility clinics, in which eggs had been fertilized by two sperm and so could not result in a live birth.

Gene-editing techniques such as those that rely on CRISPR/Cas9 had previously been used to modify DNA in adult human cells and animal embryos. Earlier this year, rumours were circulating that the methods were being applied in human embryos too, but the Huang paper is the first published report of this.

The team used CRISPR/Cas9 to modify a gene that can cause a potentially fatal blood disorder called β-thalassaemia when it is mutated. Some researchers have suggested that such a procedure, if done in a viable embryo, could eradicate devastating genetic diseases before a baby is born. Others say that such work crosses an ethical line: in response to the rumours that the work was being carried out, researchers warned in *Nature*[2] and *Science*[3] in March that because the genetic changes to embryos — a procedure known as germline modification — are heritable, they could have an unpredictable effect on future generations.

Researchers have also expressed concerns that any gene-editing research in human embryos could be a slippery slope towards unsafe, unethical or non-medical uses of the technique.

SERIOUS OBSTACLES

Huang's team says that its results reveal serious obstacles to using the method in a clinical setting. The team injected 86 embryos with CRISPR/Cas9, along with other molecules designed to add in new DNA. The researchers then waited 48 hours, by which time the embryos would have grown to about eight cells each. Of the 71 embryos that survived, 54 were genetically tested. This revealed that just 28 were successfully spliced, and that only 4 of those contained the genetic material designed to repair the cuts. "That's why we stopped," says Huang. "We still think it's too immature."

His team also found a surprising number of 'off-target' mutations assumed to be introduced by the CRISPR/Cas9 complex acting on other parts of the genome. The effect is one of the main safety concerns surrounding ▸


2018
SECOND INTERNATIONAL SUMMIT ON
HUMAN GENOME EDITING
Convened by
NATIONAL ACADEMY OF SCIENCES
NATIONAL ACADEMY OF MEDICINE
THE ROYAL SOCIETY

The international journal of science / 13 May 2021

nature

Give research into solar geoengineering a chance

There is no substitute for aggressive cuts in greenhouse-gas emissions. But the risks and benefits of technologies that could mitigate global warming need to be evaluated.

By at least one measure, US President Joe Biden's online climate summit last month was a success: several governments, including that of the United States, made fresh pledges to reduce greenhouse-gas emissions. Combined with earlier announcements from other countries and the European Union, these pledges would reduce emissions in 2030 by the equivalent of more than 3 billion tonnes of carbon dioxide, more than the current annual carbon emissions of India. But even this reduction – if achieved – would not be enough for the world to remain on a plausible path to limit warming to 1.5 °C relative to pre-industrial times.

World leaders must look for ways to close that gap at the United Nations climate convention in Glasgow, UK, in November, and then implement their commitments. Clearly, there is a long and difficult road ahead. So governments and scientists must continue to evaluate carbon capture and other climate strategies that can be used to remove carbon dioxide from the atmosphere. They should also explore solar geoengineering, which involves altering clouds or adding reflective particles to the stratosphere to reflect sunlight back into space and cool the planet. The effect of such stratospheric injections would be similar to the cooling that happens after volcanic eruptions.

Some studies suggest that solar geoengineering could provide much-needed short-term relief if global warming becomes unbearable (P. Irvine *et al. Nature Clim. Change* **9**, 295–299; 2019). But technical, environmental and ethical questions remain, including how to ensure that the cooling works as desired – and who decides the setting for the thermostat. And then there are the potential knock-on effects, which could vary across regions and sectors of society (J. Proctor *et al. Nature* **560**, 480–483; 2018). More research is needed to understand these issues.

Some scientists are vociferously opposed to solar geoengineering, which could go awry in unpredictable ways and, once started, could be difficult to safely shut down. There are also concerns that even a move to research solar geoengineering creates 'moral hazard', leading to misplaced confidence and detracting from efforts to rein in greenhouse-gas emissions. Some who oppose it fear that once research begins, the roll-out of the technology will be unstoppable, no matter what the findings.

> **"Governments need to step up climate efforts and evaluate all possible options for action."**

Researchers who study geoengineering counter that the science needs to be understood, and that the world must consider the thorny questions of international governance that would arise if a country moves forward with an ill-conceived programme. But researchers have struggled to raise funding, conduct experiments and address legitimate concerns about their work.

In March, scientists working on the Stratospheric Controlled Perturbation Experiment (SCoPEx) had to cancel a balloon flight scheduled for June in Sweden. The flight, designed to test equipment that would be used in future experiments involving the release of particles, faced opposition from environmentalists and representatives of northern Scandinavia's Sami communities, who argued that geoengineering research is a distraction from other work on climate change. An advisory committee recommended delaying the flight to allow for further public engagement.

Until now, public bodies have not prioritized reaching such a consensus. But two welcome developments suggest there is potential for this to happen.

The US National Academies of Sciences, Engineering, and Medicine recommended in March that the US government establish a coordinated federal research programme to investigate solar geoengineering. It is the most explicit call yet from a major scientific body for a government research programme, and comes at the right time.

Costing US$100 million to $200 million over 5 years, the multi-agency programme would explore the core environmental science of altering clouds or releasing particles on a large scale, as well as the ethics and public perception of this technology. The proposal includes a code of conduct, and a public registry for research proposals and results. The academies also called for broad international engagement, information sharing and discussions about global governance. This is crucial: solar geoengineering would affect the entire planet, and the United States must not go it alone.

There is also progress on the international front. The Carnegie Climate Governance Initiative (C2G), an advocacy group based in New York City, has been working to engage the UN on solar-geoengineering research and governance issues. In March 2019, environment ministers debated a resolution calling on the UN Environment Programme to assess geoengineering science and technology. That measure failed, in part because of opposition from the administration of former US president Donald Trump. Switzerland and other countries are preparing to bring forward a fresh resolution next year. And C2G is working to have solar geoengineering discussed at the UN General Assembly in 2023.

Solar-geoengineering research brings risks, and there are other, more promising ways to address global warming. But the world remains on a path to dangerous climate change. Governments need to step up climate efforts and evaluate all possible options for action. If solar geoengineering is harmful, leaders will need evidence so that they can rule out the technology.

If solar geoengineering is harmful, leaders will need evidence so that they can rule out the technology.

Punti notevoli

- Condurre la ricerca sull'uomo in modo etico è responsabilità morale e legale dello scienziato
- La ricerca sugli animali è un male necessario: chi li impiega si assume l'obbligo di rispettarli e proteggerli
- La ricerca scientifica è al servizio dell'umanità
- Per il ricercatore il rigore scientifico deve essere una forma mentis inderogabile
- Lo scienziato deve cooperare coi colleghi nell'interesse della scienza
- Il ricercatore ha precisi doveri verso i giovani che impiega presso di sé
- Dovremmo imporre limiti alla conoscenza? Se sì, in che modo?