Academic writing

Frontiers and Careers workshop

Andreas Trabesinger

Reinschrift
SCIENCE COMMUNICATION

How to prepare a manuscript

Overview and some practical tools

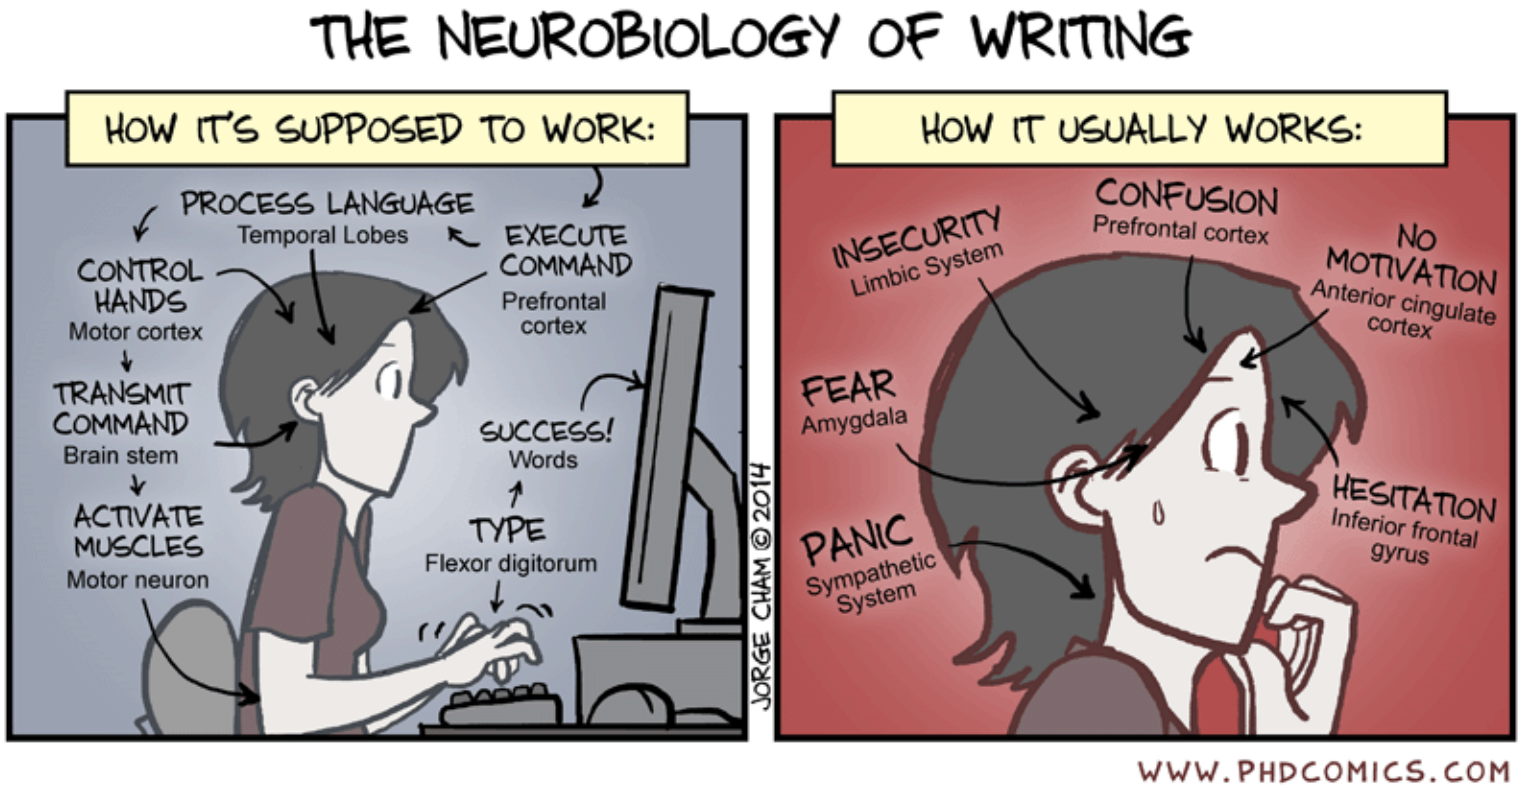

https://phdcomics.com/comics.php?f=1734

Key points covered in this talk

- Scientific storytelling
- The writing process as a multiscale problem
- Perspectives on AI-assisted writing

My background

- PhD in physics from ETH Zurich (*in vivo* NMR and MRI)
- Five years of postdoctoral experience at UC Berkeley and ETH Zurich ('ultralow-field' NMR with SQUIDs and atomic magnetometers)
- Joined in 2005 the editorial team of *Nature Physics*, involved in the launch and running of the journal
- Left *Nature Physics* in February 2012 to found my own company to assist scientists with scientific writing, editing, translations, teaching and consulting

Reinschrift
SCIENCE COMMUNICATION

The writing process

General points

- Clear writing is clear thinking.
- Good writing is rewriting.

> *"Look, writing is hard. It's like some vile, incurable disease: There are bad days, and there are worse days. And writing well is a two-stage process: (1) write not so well; (2) fix it."* **Jim Holt**

- What you write is most likely to be your legacy.
- Clarity rules.

https://www.nytimes.com/2017/05/17/books/review/do-i-make-myself-clear-harold-evans.html

The two modes of writing

1) Creative

→ Find your ‘story’

2) (Semi-)mechanical

→ Construct your text

Scientific storytelling

Writing a paper is different from keeping a lab journal.

Lab journal

- Chronological
- Individual findings
- Adding pieces to the puzzle

➢ Describes what you are doing.

Paper

- Coherent narrative
- Findings added to reach novel conclusions
- Conclusions supported by data

➢ Reports what you have done.

Key elements of a good story

“The White Rabbit put on his spectacles. ‘Where shall I begin, please your Majesty?’ he asked. ‘Begin at the beginning,’ the King said, gravely, ‘and go on till you come to the end: then stop.’”

Lewis Carroll, Alice's Adventures in Wonderland

Reinschrift
SCIENCE COMMUNICATION

From: Faber, J. How to write the introduction of a scientific article. *JWFO* **1**, E133 (2012).

Key elements of a good story

Background, open questions

→ Why is the work interesting?
→ What are major gaps or roadblocks?

Specific question addressed

→ What is your contribution?
→ Which gap have you filled?

Main results and their significance

→ What can you or we do now that your or we could not do before?
→ What do you or we know now that you or we didn't know before?

Reinschrift
SCIENCE COMMUNICATION

Key elements of a good story

Background, open questions

→ Why?

Specific question addressed

→ What?

Main results and their significance

→ So what?

Reinschrift
SCIENCE COMMUNICATION

The ‘hourglass structure’

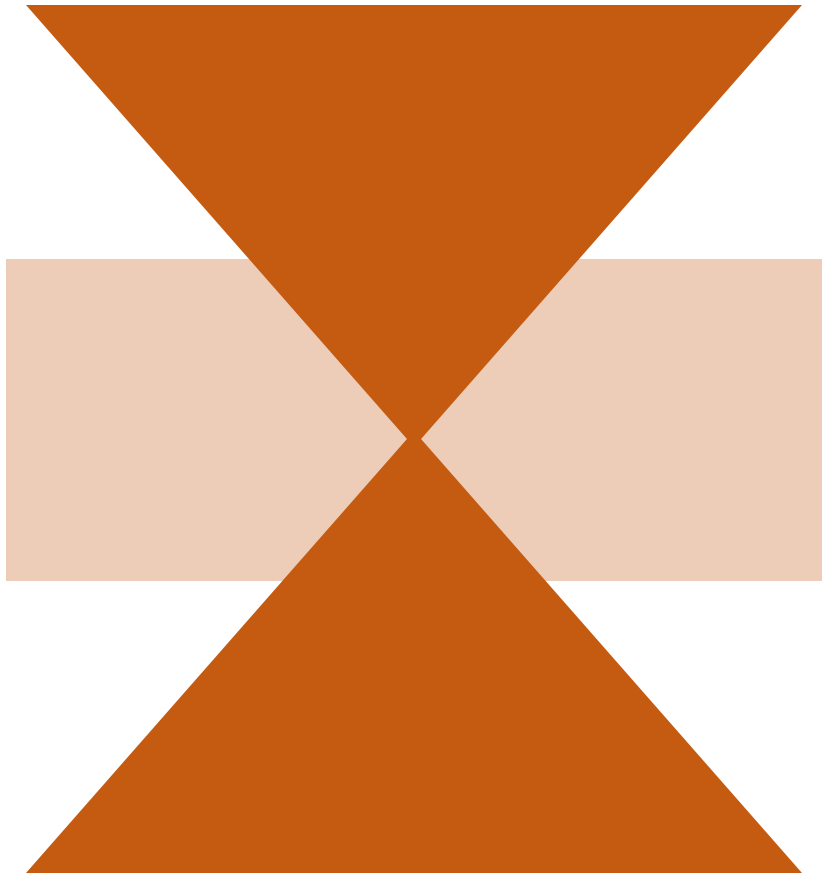

Beyond your story:
Background

Your story

Beyond your story:
Perspectives

See also:
Wu, J. Improving the writing of research papers: IMRAD and beyond. *Landscape Ecology* **26**, 1345 (2011).

How to construct a *Nature* summary paragraph

Annotated example taken from *Nature* **435,** 114–118 (5 May 2005).

One or two sentences providing a **basic introduction** to the field, comprehensible to a scientist in any discipline.

Two to three sentences of **more detailed background**, comprehensible to scientists in related disciplines.

One sentence clearly stating the **general problem** being addressed by this particular study.

One sentence summarizing the main result (with the words "**here we show**" or their equivalent).

Two or three sentences explaining what the **main result** reveals in direct comparison to what was thought to be the case previously, or how the main result adds to previous knowledge.

One or two sentences to put the results into a more **general context**.

Two or three sentences to provide a **broader perspective**, readily comprehensible to a scientist in any discipline, may be included in the first paragraph if the editor considers that the accessibility of the paper is significantly enhanced by their inclusion. Under these circumstances, the length of the paragraph can be up to 300 words. (This example is 190 words without the final section, and 250 words with it).

During cell division, mitotic spindles are assembled by microtubule-based motor proteins[1,2]. The bipolar organization of spindles is essential for proper segregation of chromosomes, and requires plus-end-directed homotetrameric motor proteins of the widely conserved kinesin-5 (BimC) family[3]. Hypotheses for bipolar spindle formation include the 'push–pull mitotic muscle' model, in which kinesin-5 and opposing motor proteins act between overlapping microtubules[2,4,5]. However, the precise roles of kinesin-5 during this process are unknown. Here we show that the vertebrate kinesin-5 Eg5 drives the sliding of microtubules depending on their relative orientation. We found in controlled *in vitro* assays that Eg5 has the remarkable capability of simultaneously moving at ~20 nm s^{-1} towards the plus-ends of each of the two microtubules it crosslinks. For anti-parallel microtubules, this results in relative sliding at ~40 nm s^{-1}, comparable to spindle pole separation rates *in vivo*[6]. Furthermore, we found that Eg5 can tether microtubule plus-ends, suggesting an additional microtubule-binding mode for Eg5. Our results demonstrate how members of the kinesin-5 family are likely to function in mitosis, pushing apart interpolar microtubules as well as recruiting microtubules into bundles that are subsequently polarized by relative sliding. We anticipate our assay to be a starting point for more sophisticated *in vitro* models of mitotic spindles. For example, the individual and combined action of multiple mitotic motors could be tested, including minus-end-directed motors opposing Eg5 motility. Furthermore, Eg5 inhibition is a major target of anti-cancer drug development, and a well-defined and quantitative assay for motor function will be relevant for such developments.

https://www.nature.com/documents/nature-summary-paragraph.pdf

An example

The Nobel Prize in Chemistry 2024

Ill. Niklas Elmehed © Nobel Prize Outreach

David Baker

Prize share: 1/2

Ill. Niklas Elmehed © Nobel Prize Outreach

Demis Hassabis

Prize share: 1/4

Ill. Niklas Elmehed © Nobel Prize Outreach

John M. Jumper

Prize share: 1/4

The Nobel Prize in Chemistry 2024 was divided, one half awarded to David Baker "for computational protein design", the other half jointly to Demis Hassabis and John M. Jumper "for protein structure prediction"

https://www.nobelprize.org/prizes/chemistry/2024/summary/

An example

Article | Open access | Published: 15 July 2021

Highly accurate protein structure prediction with AlphaFold

John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, Alex Bridgland, Clemens Meyer, Simon A. A. Kohl, Andrew J. Ballard, Andrew Cowie, Bernardino Romera-Paredes, Stanislav Nikolov, Rishub Jain, Jonas Adler, Trevor Back, Stig Petersen, David Reiman, Ellen Clancy, Michal Zielinski, Martin Steinegger, Michalina Pacholska, Tamas Berghammer, Sebastian Bodenstein, David Silver, Oriol Vinyals, Andrew W. Senior, Koray Kavukcuoglu, Pushmeet Kohli & Demis Hassabis

— Show fewer authors

Nature **596**, 583–589 (2021) | Cite this article

1.72m Accesses | **3808** Altmetric | Metrics

An example

Proteins are essential to life, and understanding their structure can facilitate a mechanistic understanding of their function. Through an enormous experimental effort[1,2,3,4], the structures of around 100,000 unique proteins have been determined[5], but this represents a small fraction of the billions of known protein sequences[6,7]. Structural coverage is bottlenecked by the months to years of painstaking effort required to determine a single protein structure. Accurate computational approaches are needed to address this gap and to enable large-scale structural bioinformatics. Predicting the three-dimensional structure that a protein will adopt based solely on its amino acid sequence—the structure prediction component of the 'protein folding problem'[8]—has been an important open research problem for more than 50 years[9]. Despite recent progress[10,11,12,13,14], existing methods fall far short of atomic accuracy, especially when no homologous structure is available. Here we provide the first computational method that can regularly predict protein structures with atomic accuracy even in cases in which no similar structure is known. We validated an entirely redesigned version of our neural network-based model, AlphaFold, in the challenging 14th Critical Assessment of protein Structure Prediction (CASP14)[15], demonstrating accuracy competitive with experimental structures in a majority of cases and greatly outperforming other methods. Underpinning the latest version of AlphaFold is a novel machine learning approach that incorporates physical and biological knowledge about protein structure, leveraging multi-sequence alignments, into the design of the deep learning algorithm.

One or two sentences providing a **basic introduction** to the field, comprehensible to a scientist in any discipline.

Two to three sentences of **more detailed background**, comprehensible to scientists in related disciplines.

One sentence clearly stating the **general problem** being addressed by this particular study.

One sentence summarizing the main result (with the words "**here we show**" or their equivalent).

Two or three sentences explaining what the **main result** reveals in direct comparison to what was thought to be the case previously, or how the main result adds to previous knowledge.

One or two sentences to put the results into a more **general context**.

Two or three sentences to provide a **broader perspective**, readily comprehensible to a scientist in any discipline, may be included in the first paragraph if the editor considers that the accessibility of the paper is significantly enhanced by their inclusion. Under these circumstances, the length of the paragraph can be up to 300 words. (This example is 190 words without the final section, and 250 words with it).

Nature **596**, 583–589 (2021).

Reinschrift
SCIENCE COMMUNICATION

Storytelling

- When is a good moment to publish?
- What to include in the paper?

→ Have a story

Reinschrift
SCIENCE COMMUNICATION

Storytelling

Don’t forget the ‘so what’.

The Ratio of Proton and Electron Masses

FRIEDRICH LENZ
Düsseldorf, Germany
(Received April 5, 1951)

THE most exact value at present[1] for the ratio of proton to electron mass is 1836.12 ± 0.05. It may be of interest to note that this number coincides with $6\pi^5 = 1836.12$.

[1] Sommer, Thomas, and Hipple, Phys. Rev. **80**, 487 (1950).

Physical Review **82**, 554 (1951).

Storytelling — Posters

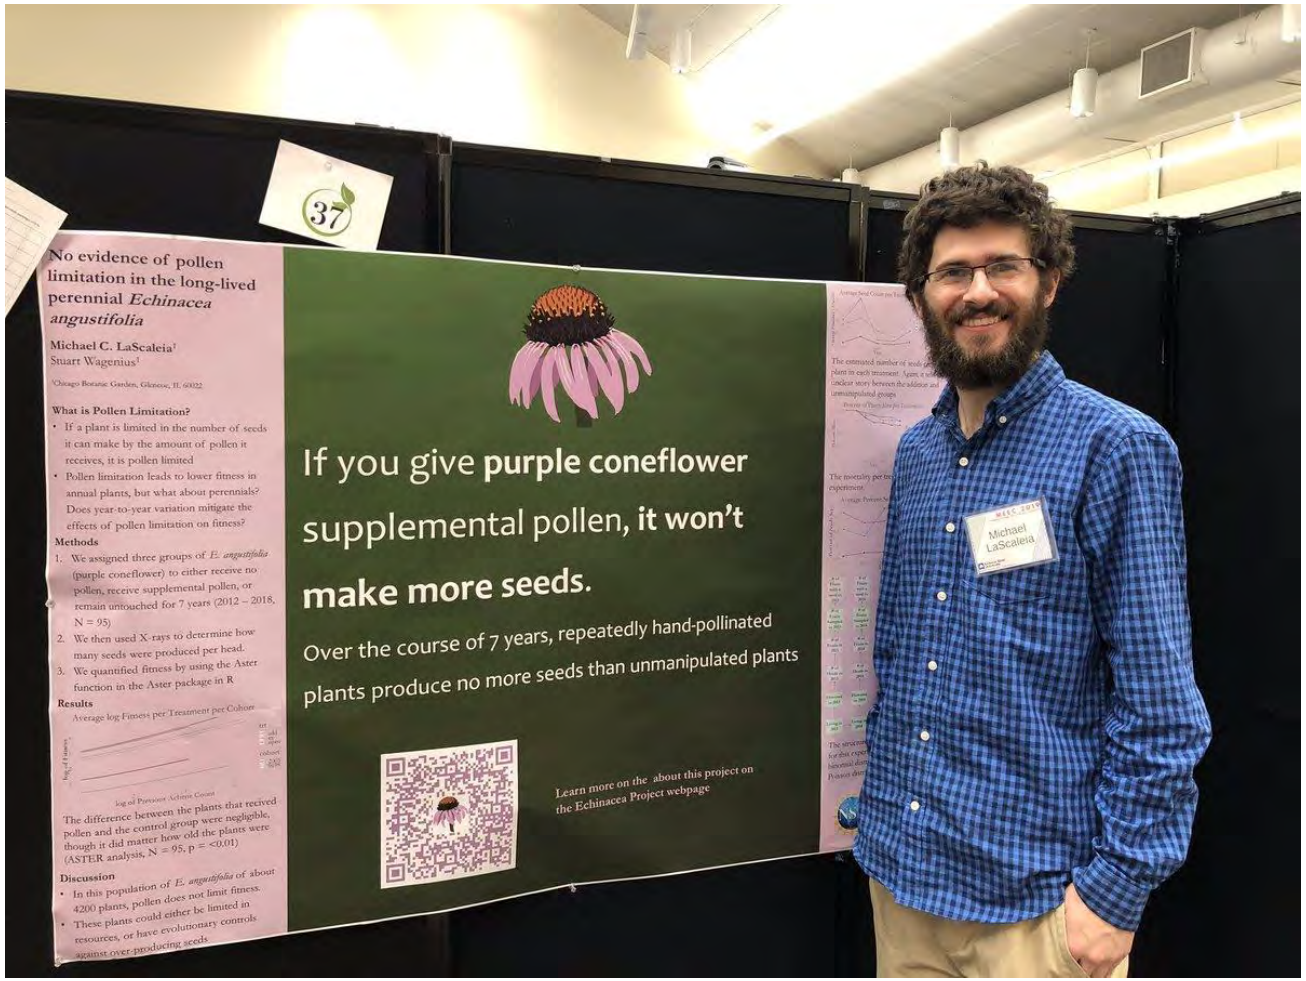

Michael C LaScaleia @MCLaScience
https://x.com/MCLaScience/status/1122505413386231810

Reinschrift
SCIENCE COMMUNICATION

Storytelling — Science communication

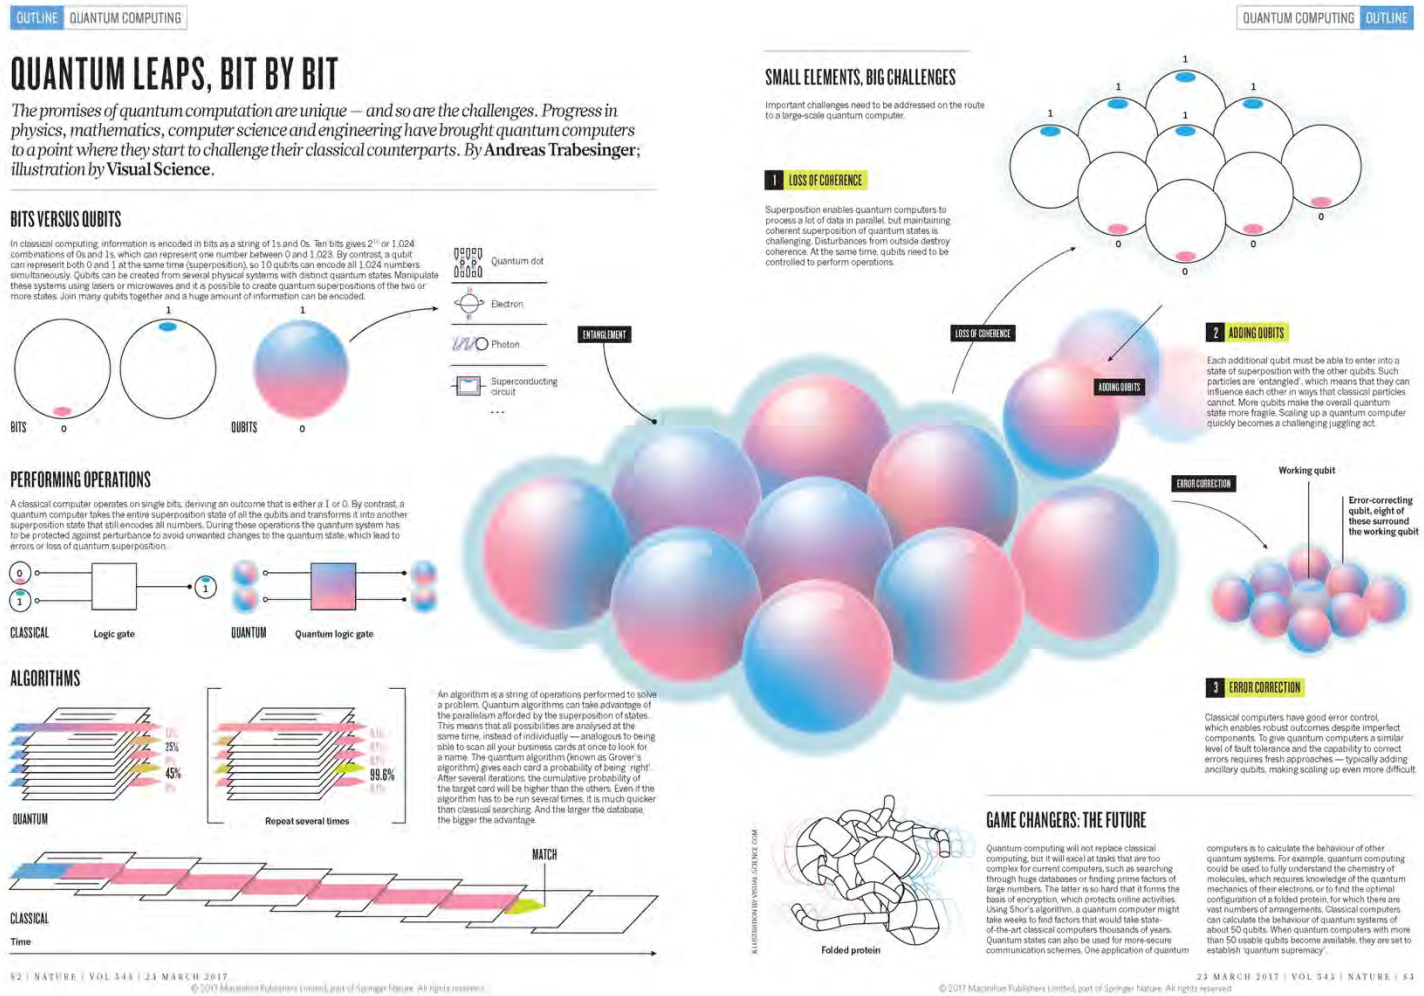

OUTLINE QUANTUM COMPUTING

QUANTUM LEAPS, BIT BY BIT

The promises of quantum computation are unique — and so are the challenges. Progress in physics, mathematics, computer science and engineering have brought quantum computers to a point where they start to challenge their classical counterparts. By **Andreas Trabesinger**; *illustration by* **Visual Science**.

BITS VERSUS QUBITS

In classical computing, information is encoded in bits as a string of 1s and 0s. Ten bits gives 2^{10} or 1,024 combinations of 0s and 1s, which can represent one number between 0 and 1,023. By contrast, a qubit can represent both 0 and 1 at the same time (superposition), so 10 qubits can encode all 1,024 numbers simultaneously. Qubits can be created from several physical systems with distinct quantum states. Manipulate these systems using lasers or microwaves and it is possible to create quantum superpositions of the two or more states. Join many qubits together and a huge amount of information can be encoded.

Quantum dot
Electron
Photon
Superconducting circuit

BITS 0 1 QUBITS 0 1

ENTANGLEMENT

PERFORMING OPERATIONS

A classical computer operates on single bits, deriving an outcome that is either a 1 or 0. By contrast, a quantum computer takes the entire superposition state of all the qubits and transforms it into another superposition state that still encodes all numbers. During these operations the quantum system has to be protected against perturbance to avoid unwanted changes to the quantum state, which lead to errors or loss of quantum superposition.

CLASSICAL Logic gate QUANTUM Quantum logic gate

ALGORITHMS

An algorithm is a string of operations performed to solve a problem. Quantum algorithms can take advantage of the parallelism afforded by the superposition of states. This means that all possibilities are analysed at the same time, instead of individually — analogous to being able to scan all your business cards at once to look for a name. The quantum algorithm (known as Grover's algorithm) gives each card a probability of being 'right'. After several iterations, the cumulative probability of the target card will be higher than the others. Even if the algorithm has to be run several times, it is much quicker than classical searching. And the larger the database, the bigger the advantage.

QUANTUM 12.5% 25% 6.3% 45% 6.3% Repeat several times 3.1% 6.3% 12.5% 99.6% 6.3%

CLASSICAL MATCH Time

QUANTUM COMPUTING OUTLINE

SMALL ELEMENTS, BIG CHALLENGES

Important challenges need to be addressed on the route to a large-scale quantum computer.

1 LOSS OF COHERENCE

Superposition enables quantum computers to process a lot of data in parallel, but maintaining coherent superposition of quantum states is challenging. Disturbances from outside destroy coherence. At the same time, qubits need to be controlled to perform operations.

LOSS OF COHERENCE ADDING QUBITS

2 ADDING QUBITS

Each additional qubit must be able to enter into a state of superposition with the other qubits. Such particles are 'entangled', which means that they can influence each other in ways that classical particles cannot. More qubits make the overall quantum state more fragile. Scaling up a quantum computer quickly becomes a challenging juggling act.

ERROR CORRECTION Working qubit Error-correcting qubit, eight of these surround the working qubit

3 ERROR CORRECTION

Classical computers have good error control, which enables robust outcomes despite imperfect components. To give quantum computers a similar level of fault tolerance and the capability to correct errors requires fresh approaches — typically adding ancillary qubits, making scaling up even more difficult.

GAME CHANGERS: THE FUTURE

Quantum computing will not replace classical computing, but it will excel at tasks that are too complex for current computers, such as searching through huge databases or finding prime factors of large numbers. The latter is so hard that it forms the basis of encryption, which protects online activities. Using Shor's algorithm, a quantum computer might take weeks to find factors that would take state-of-the-art classical computers thousands of years. Quantum states can also be used for more-secure communication schemes. One application of quantum computers is to calculate the behaviour of other quantum systems. For example, quantum computing could be used to fully understand the chemistry of molecules, which requires knowledge of the quantum mechanics of their electrons, or to find the optimal configuration of a folded protein, for which there are vast numbers of arrangements. Classical computers can calculate the behaviour of quantum systems of about 50 qubits. When quantum computers with more than 50 usable qubits become available, they are set to establish 'quantum supremacy'.

Folded protein

ILLUSTRATION BY VISUAL SCIENCE.COM

A. Trabesinger *Nature* **543**, S2–S3 (2017)
https://www.nature.com/articles/543S2a

Reinschrift
SCIENCE COMMUNICATION

Proposal writing

Experimental Searches for Oscillating and Transient effects from the Dark Sector

Fact Sheet | Reporting | Results | News & Multimedia

Objective

The objective of the proposed project is to pioneer a magnetometry-based experimental framework for the detection of time-varying signatures of the 'dark sector'. This novel approach will enable systematic searches for particles contributing to the dark matter and for dark-energy components.

The nature of dark matter and that of dark energy are among the central open problems in modern physics. There are only few experimental bounds and so far no conclusive observations of dark-sector particles or fields. Experiments enabling a direct coupling to the dark sector and thus a systematic search for and study of the contributing particles and fields would open up new vistas for areas ranging from particle physics to astrophysics and cosmology, and would in particular provide insights into the physics beyond the Standard Model.

Here, we propose a framework for such experimental searches based on high-precision magnetometers, and networks thereof. Our approach is distinct from existing efforts in two ways. First, it will enable searches for so-far unexplored couplings to ultra-light bosonic particles present in the Universe that could be components of dark matter and/or dark energy, in particular axions and axion-like particles (ALPs). Second, we will develop and use devices and methods tailored to search for oscillating and transient, rather than time-independent, effects. Specifically, we will use nuclear magnetic resonance (NMR) techniques for detecting spin precession caused by background axion and ALP dark matter, and geographically separated magnetometers for identify transient effects, such as crossing domain walls of ALP fields, which have been proposed as a possible dark-energy component.

The devices and methods developed in the framework of this project will provide the essential components for unique searches for a broad class of dark-matter and dark-energy candidates and might enable the key experiments to understanding the dark sector.

Project Information

Dark-OsT
Grant agreement ID: 695405

DOI
10.3030/695405

Closed project

Start date
1 August 2016

End date
31 July 2021

Funded under
EXCELLENT SCIENCE - European Research Council (ERC)

Total cost
€ 2 474 875

EU contribution
€ 2 474 875

Hosted by
JOHANNES GUTENBERG-UNIVERSITAT MAINZ
Germany

https://cordis.europa.eu/project/id/695405

Reinschrift
SCIENCE COMMUNICATION

The objective of the proposed project is to pioneer a magnetometry-based experimental framework for the detection of time-varying signatures of the ‘dark sector’. This novel approach will enable systematic searches for particles contributing to the dark matter and for dark-energy components.

In a nutshell

The nature of dark matter and that of dark energy are among the central open problems in modern physics. There are only few experimental bounds and so far no conclusive observations of dark-sector particles or fields. Experiments enabling a direct coupling to the dark sector and thus a systematic search for and study of the contributing particles and fields would open up new vistas for areas ranging from particle physics to astrophysics and cosmology, and would in particular provide insights into the physics beyond the Standard Model.

Background, open problems

Here, we propose a framework for such experimental searches based on high-precision magnetometers, and networks thereof. Our approach is distinct from existing efforts in two ways. First, it will enable searches for so-far unexplored couplings to ultra-light bosonic particles present in the Universe that could be components of dark matter and/or dark energy, in particular axions and axion-like particles (ALPs). Second, we will develop and use devices and methods tailored to search for oscillating and transient, rather than time-independent, effects. Specifically, we will use nuclear magnetic resonance (NMR) techniques for detecting spin precession caused by background axion and ALP dark matter, and geographically separated magnetometers for identify transient effects, such as crossing domain walls of ALP fields, which have been proposed as a possible dark-energy component.

Question addressed

The devices and methods developed in the framework of this project will provide the essential components for unique searches for a broad class of dark-matter and dark-energy candidates and might enable the key experiments to understanding the dark sector.

Main results and their significance

An ‘hourglass plus’

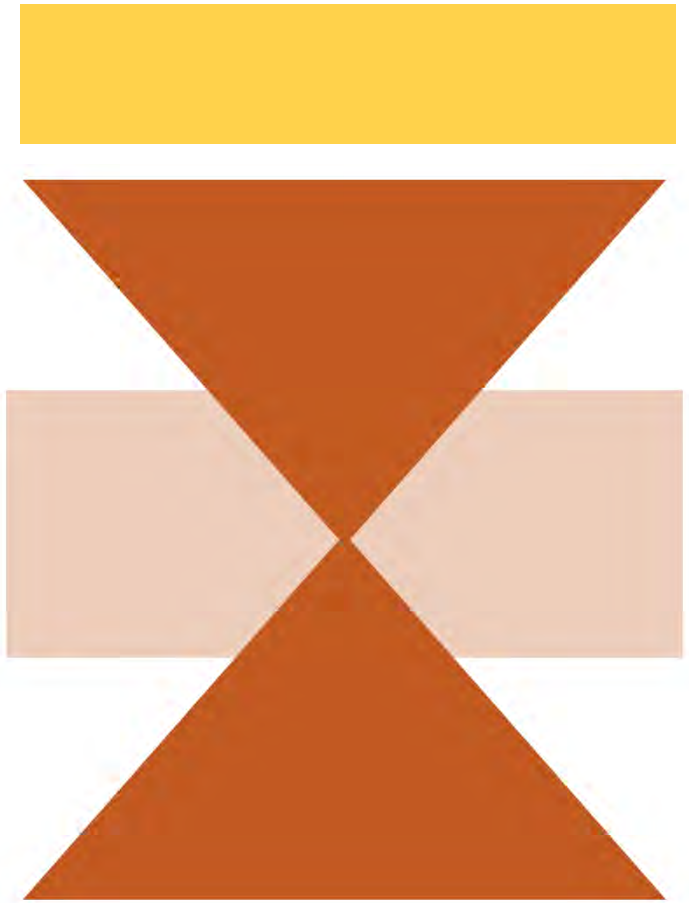

In a nutshell

Beyond your story:
Background

Your story

Beyond your story:
Perspectives

Reinschrift
SCIENCE COMMUNICATION

See also:
Wu, J. Improving the writing of research papers: IMRAD and beyond. *Landscape Ecology* **26**, 1345 (2011).

Summaries in interdisciplinary journals

nature physics

Article

https://doi.org/10.1038/s41567-024-02566-1

Empowering deep neural quantum states through efficient optimization

Received: 20 March 2023

Accepted: 28 May 2024

Published online: 1 July 2024

Check for updates

Ao Chen & Markus Heyl

Computing the ground state of interacting quantum matter is a long-standing challenge, especially for complex two-dimensional systems. Recent developments have highlighted the potential of neural quantum states to solve the quantum many-body problem by encoding the many-body wavefunction into artificial neural networks. However, this method has faced the critical limitation that existing optimization algorithms are not suitable for training modern large-scale deep network architectures. Here, we introduce a minimum-step stochastic-reconfiguration optimization algorithm, which allows us to train deep neural quantum states with up to 10^6 parameters. We demonstrate our method for paradigmatic frustrated spin-1/2 models on square and triangular lattices, for which our trained deep networks approach machine precision and yield improved variational energies compared to existing results. Equipped with our optimization algorithm, we find numerical evidence for gapless quantum-spin-liquid phases in the considered models, an open question to date. We present a method that captures the emergent complexity in quantum many-body problems through the expressive power of large-scale artificial neural networks.

It has been an ever-persisting quest in condensed-matter and quantum many-body physics to capture the essence of quantum many-body systems that is covered behind their exponential complexity. Although many numerical methods have been developed to access the quantum many-body problem with strong interactions, it still remains an extraordinary challenge to obtain accurate ground-state solutions, especially for complex and large two-dimensional systems. The respective challenges depend on the method utilized, such as the 'curse of dimensionality' in exact diagonalization[1], the notorious sign problem[2] in quantum Monte Carlo approaches[3] or the growth of entanglement and matrix contraction complexity in tensor network methods[4]. One of the paradigmatic instances of such complex two-dimensional quantum matter is the putative quantum-spin-liquid (QSL) phase in frustrated magnets[5]. Although a large variety of different numerical methods have been applied, the nature of many of the presumed QSLs still remains debated, such as the prototypical frustrated Heisenberg J_1–J_2 magnets on square[6–15] or triangular lattices[13–22].

Recently, neural quantum states (NQSs) have been introduced as a promising alternative for solving the quantum many-body problem by means of artificial neural networks[23]. This approach has already seen tremendous progress for QSLs[24–26]. However, this method also faces an outstanding challenge critically limiting its capabilities and its potential to date. Due to the rugged quantum landscape[27] with many saddle points, it is typically necessary to utilize stochastic reconfiguration (SR)[28] in the optimization. SR is a quantum generalization of natural gradient descent[29] and has a $\mathcal{O}(N_p^3)$ complexity for a network with N_p parameters, which impedes the training of deep networks. Consequently, the current applications of NQS mainly focus on shallow networks, such as a restricted Boltzmann machine (RBM)[23,30] or shallow convolutional neural networks (CNNs)[25,31] with no more than ten layers and around 10^3 parameters. Many efforts have been made to overcome the optimization difficulty in deep NQS based on either iterative solvers[23], approximate optimizers[32–36] or large-scale supercomputers[37,38]. However, the cost of SR still represents the key

Center for Electronic Correlations and Magnetism, University of Augsburg, Augsburg, Germany. e-mail: chenao.phys@gmail.com; markus.heyl@uni-a.de

Research briefing

Check for updates

Efficient optimization of deep neural quantum states

An improved optimization algorithm enables the training of large-scale neural quantum states in which the enormous number of neuron connections capture the intricate complexity of quantum many-body wavefunctions. This advance leads to unprecedented accuracy in paradigmatic quantum models, opening up new avenues for simulating and understanding complex quantum phenomena.

The problem

The quantum many-body problem has remained an outstanding challenge when it comes to the regime of strongly interacting particles in two dimensions — a regime that hosts important exotic phases (such as quantum spin liquids) and represents an experimental frontier (for example, in Rydberg atom arrays). Recently, the neural quantum state (NQS) has been introduced as a computational approach to numerically solve the quantum many-body problem[1]. In this approach, the quantum state is encoded into an artificial neural network (as illustrated in Fig. 1a) and then the power of machine learning is brought to bear.

However, a drawback in the NQS approach has been that the training of the underlying networks is not efficient. This has imposed critical limitations on network sizes and therefore on the capability of the NQS to solve the quantum many-body problem on a large scale in challenging regimes[2].

The solution

The training of an NQS, to determine the ground state of a quantum many-body system, involves the minimization of the system's energy. This is typically achieved through so-called imaginary-time evolution, which involves the inversion of a matrix of size $N_p \times N_p$ as its key step, where N_p denotes the number of parameters in the underlying artificial neural network. The computational effort of this inversion grows as N_p^3 and becomes intractable for large networks. The key contribution of our work is a reformulation of the fundamental equation in imaginary-time evolution, which curtails the growth of computational effort, making it linear with N_p (Fig. 1b). As a result, it is now possible to train unprecedentedly large NQSs, with up to one million parameters, and to make full use of the power of artificial neural networks.

This new approach has enabled us to target strongly interacting quantum many-body systems with a new level of accuracy, outperforming traditional methods. We have demonstrated this concretely for frustrated quantum magnets in two dimensions, which are considered to be among the most difficult quantum problems[3]: for so-called J_1–J_2 Heisenberg models, we have shown that ground states can be obtained with unprecedented accuracy, compared with traditional methods, and we have found strong numerical evidence that the quantum spin liquid phases in these systems are gapless — something that has been the subject of long-standing debate.

The implications

In quantum many-body physics, there is the famous philosophy of 'more is different', highlighting the emergence of new phenomena at different levels of complexity. For machine learning, there is a philosophy of connectionism, meaning that new levels of machine intelligence can emerge where there is an enormous number of artificial-neuron connections. Our work brings these two ideas together and shows that the emergent machine intelligence in an NQS is able to capture the complex emergent phenomena in strongly-correlated systems — opening up a new perspective on quantum many-body physics.

Although we have now demonstrated that quantum magnets can be solved very efficiently using NQSs, other challenges remain. For instance, clearly, the goal eventually would be to address systems of fermions or electrons, such as in the context of the paradigmatic Hubbard model. However, owing to the antisymmetric nature of fermions, the NQS cannot be directly applied to fermionic systems. With the training of NQSs becoming more efficient, it is now pressing to design neural networks that are better suited to the underlying fermionic structures[4].

Based on the advances made using NQSs in the ground-state search for quantum many-body systems, it is also possible to envisage using NQSs for quantum dynamics[5] — which would be very welcome in view of the tremendous experimental developments in quantum simulators, such as Rydberg atom arrays.

Ao Chen & Markus Heyl
University of Augsburg, Augsburg, Germany.

This is a summary of:
Chen, A. & Heyl, M. Empowering deep neural quantum states through efficient optimization. *Nat. Phys.* https://doi.org/10.1038/s41567-024-02566-1 (2024).

Published online: 9 July 2024

Reinschrift
SCIENCE COMMUNICATION

Summaries in interdisciplinary journals

Research briefing

Check for updates

Efficient optimization of deep neural quantum states

An improved optimization algorithm enables the training of large-scale neural quantum states in which the enormous number of neuron connections capture the intricate complexity of quantum many-body wavefunctions. This advance leads to unprecedented accuracy in paradigmatic quantum models, opening up new avenues for simulating and understanding complex quantum phenomena.

The problem

The quantum many-body problem has remained an outstanding challenge when it comes to the regime of strongly interacting particles in two dimensions — a regime that hosts important exotic phases (such as quantum spin liquids) and represents an experimental frontier (for example, in Rydberg atom arrays). Recently, the neural quantum state (NQS) has been introduced as a computational approach to numerically solve the quantum many-body problem[1]. In this approach, the quantum state is encoded into an artificial neural network (as illustrated in Fig. 1a) and then the power of machine learning is brought to bear.

However, a drawback in the NQS approach has been that the training of the underlying networks is not efficient. This has imposed critical limitations on network sizes and therefore on the capability of the NQS to solve the quantum many-body problem on a large scale in challenging regimes[2].

The solution

The training of an NQS, to determine the ground state of a quantum many-body system, involves the minimization of the system's energy. This is typically achieved through so-called imaginary-time evolution, which involves the inversion of a matrix of size $N_p \times N_p$ as its key step, where N_p denotes the number of parameters in the underlying artificial neural network. The computational effort of this inversion grows as N_p^3 and becomes intractable for large networks. The key contribution of our work is a reformulation of the fundamental equation in imaginary-time evolution, which curtails the growth of computational effort, making it linear with N_p (Fig. 1b). As a result, it is now possible to train unprecedentedly large NQSs, with up to one million parameters, and to make full use of the power of artificial neural networks.

This new approach has enabled us to target strongly interacting quantum many-body systems with a new level of accuracy, outperforming traditional methods. We have demonstrated this concretely for frustrated quantum magnets in two dimensions, which are considered to be among the most difficult quantum problems[3]: for so-called J_1–J_2 Heisenberg models, we have shown that ground states can be obtained with unprecedented accuracy, compared with traditional methods, and we have found strong numerical evidence that the quantum spin liquid phases in these systems are gapless — something that has been the subject of long-standing debate.

The implications

In quantum many-body physics, there is the famous philosophy 'more is different', highlighting the emergence of new phenomena at different levels of complexity. For machine learning, there is a philosophy of connectionism, meaning that new levels of machine intelligence can emerge where there is an enormous number of artificial-neuron connections. Our work brings these two ideas together and shows that the emergent machine intelligence in an NQS is able to capture the complex emergent phenomena in strongly-correlated systems — opening up a new perspective on quantum many-body physics.

Although we have now demonstrated that quantum magnets can be solved very efficiently using NQSs, other challenges remain. For instance, clearly, the goal eventually would be to address systems of fermions or electrons, such as in the context of the paradigmatic Hubbard model. However, owing to the antisymmetric nature of fermions, the NQS cannot be directly applied to fermionic systems. With the training of NQSs becoming more efficient, it is now pressing to design neural networks that are better suited to the underlying fermionic structures[4].

Based on the advances made using NQSs in the ground-state search for quantum many-body systems, it is also possible to envisage using NQSs for quantum dynamics[5] — which would be very welcome in view of the tremendous experimental developments in quantum simulators, such as Rydberg atom arrays.

Ao Chen & Markus Heyl
University of Augsburg, Augsburg, Germany.

This is a summary of:
Chen, A. & Heyl, M. Empowering deep neural quantum states through efficient optimization. *Nat. Phys.* https://doi.org/10.1038/s41567-024-02566-1 (2024).

Published online: 9 July 2024

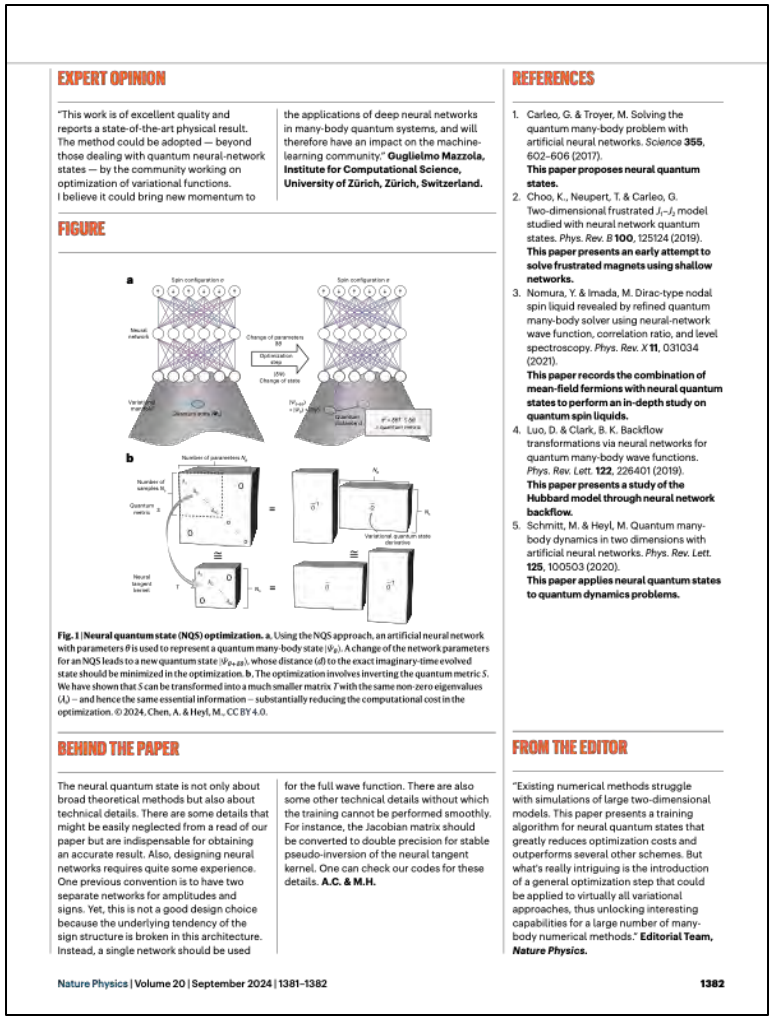

EXPERT OPINION

"This work is of excellent quality and reports a state-of-the-art physical result. The method could be adopted — beyond those dealing with quantum neural-network states — by the community working on optimization of variational functions. I believe it could bring new momentum to the applications of deep neural networks in many-body quantum systems, and will therefore have an impact on the machine-learning community." **Guglielmo Mazzola, Institute for Computational Science, University of Zürich, Zürich, Switzerland.**

FIGURE

Fig. 1 | Neural quantum state (NQS) optimization. a. Using the NQS approach, an artificial neural network with parameters θ is used to represent a quantum many-body state $|\Psi_\theta\rangle$. A change of the network parameters for an NQS leads to a new quantum state $|\Psi_{\theta+d\theta}\rangle$, whose distance ($d\theta$) to the exact imaginary-time evolved state should be minimized in the optimization. **b.** The optimization involves inverting the quantum metric S. We have shown that S can be transformed into a much smaller matrix T with the same non-zero eigenvalues (λ_i) — and hence the same essential information — substantially reducing the computational cost in the optimization.

BEHIND THE PAPER

The neural quantum state is not only about broad theoretical methods but also about technical details. There are some details that might be easily neglected from a read of our paper but are indispensable for obtaining an accurate result. Also, designing neural networks requires quite some experience. One previous convention is to have two separate networks for amplitudes and signs. Yet, this is not a good design choice because the underlying tendency of the sign structure is broken in this architecture. Instead, a single network should be used for the full wave function. There are also some other technical details without which the training cannot be performed smoothly. For instance, the Jacobian matrix should be converted to double precision for stable pseudo-inversion of the neural tangent kernel. One can check our codes for these details. **A.C. & M.H.**

REFERENCES

1. Carleo, G. & Troyer, M. Solving the quantum many-body problem with artificial neural networks. *Science* **355**, 602–606 (2017).
This paper proposes neural quantum states.
2. Choo, K., Neupert, T. & Carleo, G. Two-dimensional frustrated J_1–J_2 model studied with neural network quantum states. *Phys. Rev. B* **100**, 125124 (2019).
This paper presents an early attempt to solve frustrated magnets using shallow networks.
3. Nomura, Y. & Imada, M. Dirac-type nodal spin liquid revealed by refined quantum many-body solver using neural-network wave function, correlation ratio, and level spectroscopy. *Phys. Rev. X* **11**, 031034 (2021).
This paper records the combination of mean-field fermions with neural quantum states to perform an in-depth study on quantum spin liquids.
4. Luo, D. & Clark, B. K. Backflow transformations via neural networks for quantum many-body wave functions. *Phys. Rev. Lett.* **122**, 226401 (2019).
This paper presents a study of the Hubbard model through neural network backflow.
5. Schmitt, M. & Heyl, M. Quantum many-body dynamics in two dimensions with artificial neural networks. *Phys. Rev. Lett.* **125**, 100503 (2020).
This paper applies neural quantum states to quantum dynamics problems.

FROM THE EDITOR

"Existing numerical methods struggle with simulations of large two-dimensional models. This paper presents a training algorithm for neural quantum states that greatly reduces optimization costs and outperforms several other schemes. But what's really intriguing is the introduction of a general optimization step that could be applied to virtually all variational approaches, thus unlocking interesting capabilities for a large number of many-body numerical methods." **Editorial Team, *Nature Physics*.**

Reinschrift
SCIENCE COMMUNICATION

FINER research questions

Feasible

- Feasibility depends on resources, time and finances.
- Access to people, materials and data is crucial.
- Data must be able to be collected within available resources.

Interesting

- Research should be based on genuine interest.

Novel

- The question should not simply copy existing research.
- It should offer new insights or extend existing insights.

Ethical

- Ethics is the primary requirement and requires approval.
- Minimising risks to participants, privacy and confidentiality are critical.

Relevant

- The question should arouse academic and intellectual interest.

Reinschrift
SCIENCE COMMUNICATION

Hulley, S.B. *et al.* Designing clinical research. 3rd ed. Philadelphia (PA): Lippincott Williams and Wilkins; 2007.

Getting to the heart of the research project

The two modes of writing

1) Creative

→ Find your ‘story’

2) (Semi-)mechanical

→ Construct your text

Constructing your text

- Manuscript level
- Section level
- Paragraph level
- Sentence level
- Word level

Manuscript level

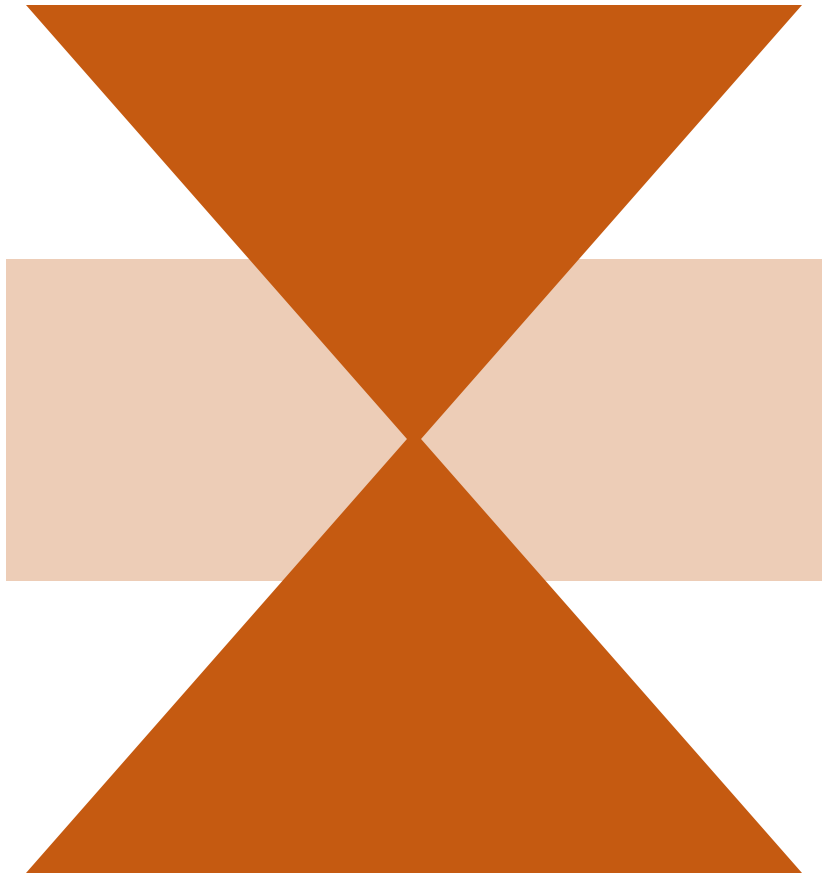

Beyond your story:

Background

Your story

Beyond your story:

Perspectives

See also:
Wu, J. Improving the writing of research papers: IMRAD and beyond. *Landscape Ecology* **26**, 1345 (2011).

Constructing your text

- Manuscript level
- Section level
- Paragraph level
- Sentence level
- Word level

Typical structure of research papers

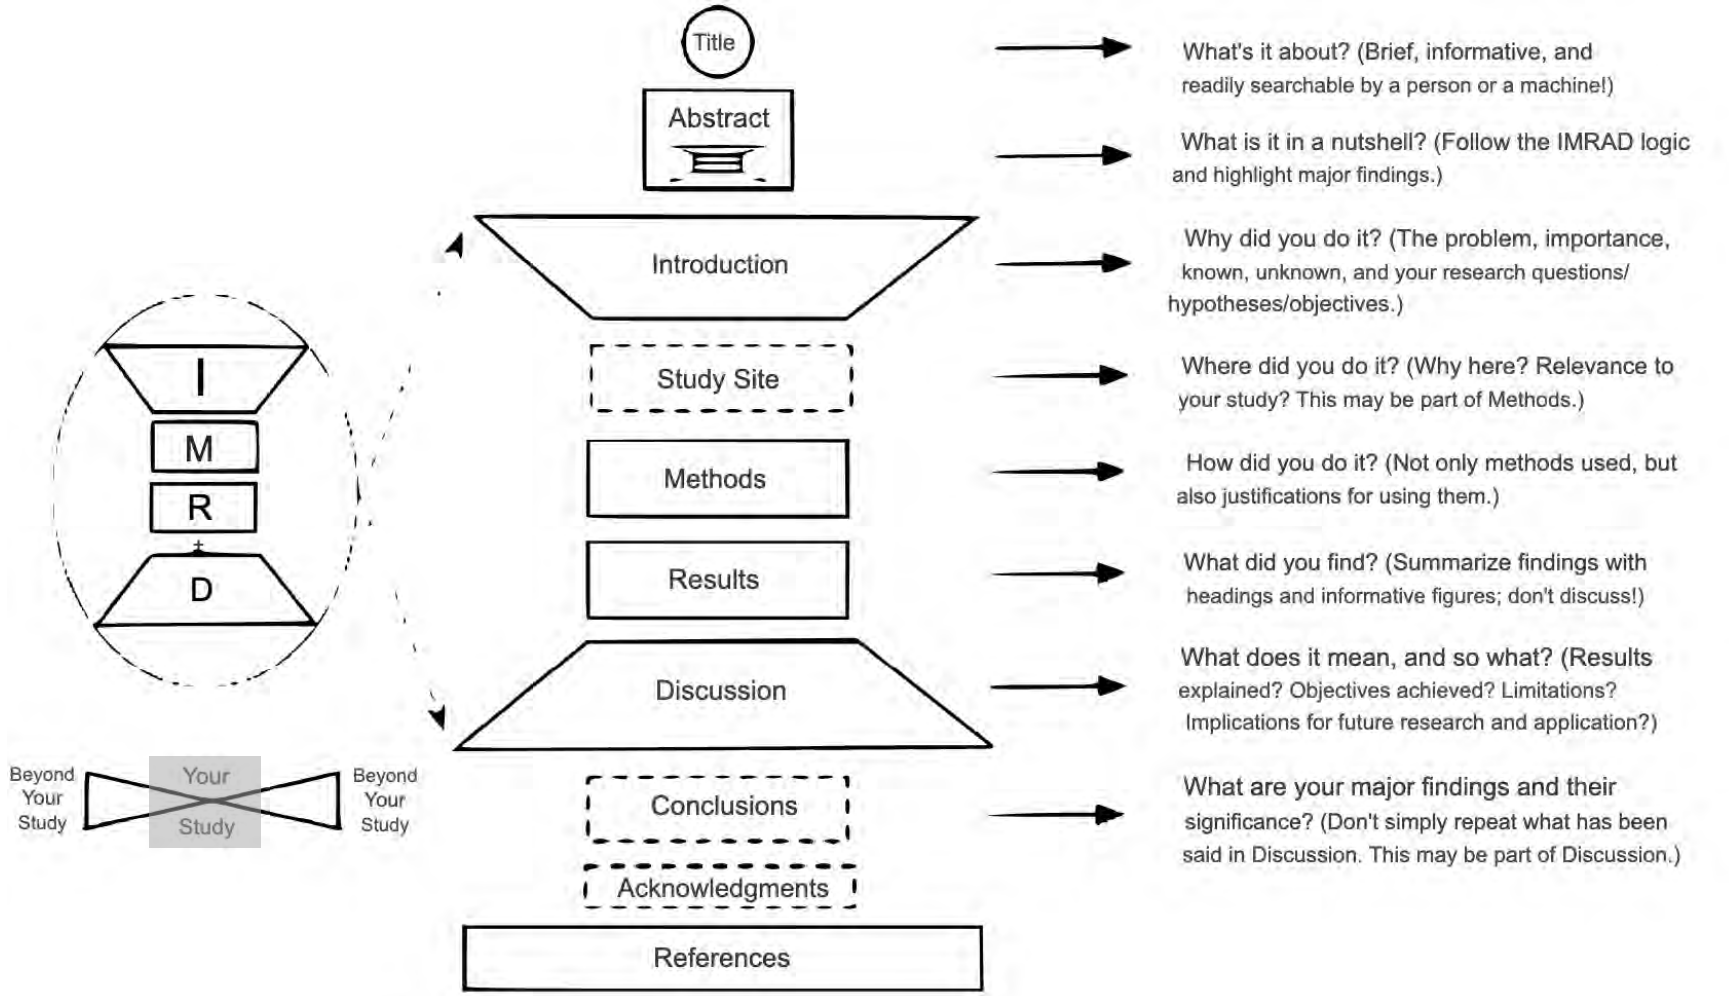

Reinschrift
SCIENCE COMMUNICATION

Wu, J. Improving the writing of research papers: IMRAD and beyond. *Landscape Ecology* **26**, 1345 (2011).

Typical structure of research papers

- Title
- Abstract
- Introduction
- Methods and results
- Discussion
- Conclusions
- Detailed methods
- Supplementary Information/Material

Reinschrift
SCIENCE COMMUNICATION

Typical structure of research papers

- **Title**
- Abstract
- Introduction
- Methods and results
- Discussion
- Conclusions
- Detailed methods
- Supplementary Information/Material

Reinschrift
SCIENCE COMMUNICATION

Paper title

Should describe what you have *found*, not what you have *done*

- "Observational evidence of nocturnal behaviour in *Mesocricetus auratus*"
- *not* "Observing the circadian cycle of *Mesocricetus auratus*"
- *not* "Do hamsters sleep at night?"

Concise, enticing, specific, accurate

Should include main keywords

Avoid

→ Acronyms
→ Questions
→ Clichés

Reinschrift
SCIENCE COMMUNICATION

Bad titles

Avoid clichés like the plague

- Holy Grail
- Silver bullet / Magic bullet
- Shedding light
- Missing link
- Paradigm shift
- Rosetta Stone

See also: https://www.nature.com/articles/nphys724

Clichés

Formalin-fixed paraffin-embedded tissue: The holy grail of clinical proteomics?

Specific subsets of mesenchymal stroma cells to treat lung disorders – Finding the Holy Grail

In what sense a neutron star-black hole binary is the holy grail for testing gravity?[1]

Network Pharmacology: A Rosetta Stone for Traditional Chinese Medicine

Silyl Imine Electrophiles in Enantioselective Catalysis: A Rosetta Stone for Peptide Homologation, Enabling Diverse *N*-Protected Aryl Glycines from Aldehydes in Three Steps

No rules without exceptions

IOP PUBLISHING JOURNAL OF PHYSICS A: MATHEMATICAL AND THEORETICAL

J. Phys. A: Math. Theor. **44** (2011) 492001 (5pp) doi:10.1088/1751-8113/44/49/492001

FAST TRACK COMMUNICATION

Can apparent superluminal neutrino speeds be explained as a quantum weak measurement?

M V Berry[1], **N Brunner**[1], **S Popescu**[1] **and P Shukla**[2]

[1] H H Wills Physics Laboratory, Tyndall Avenue, Bristol BS8 1TL, UK
[2] Department of Physics, Indian Institute of Technology, Kharagpur, India

Received 12 October 2011, in final form 27 October 2011
Published 11 November 2011
Online at stacks.iop.org/JPhysA/44/492001

Abstract
Probably not.

Typical structure of research papers

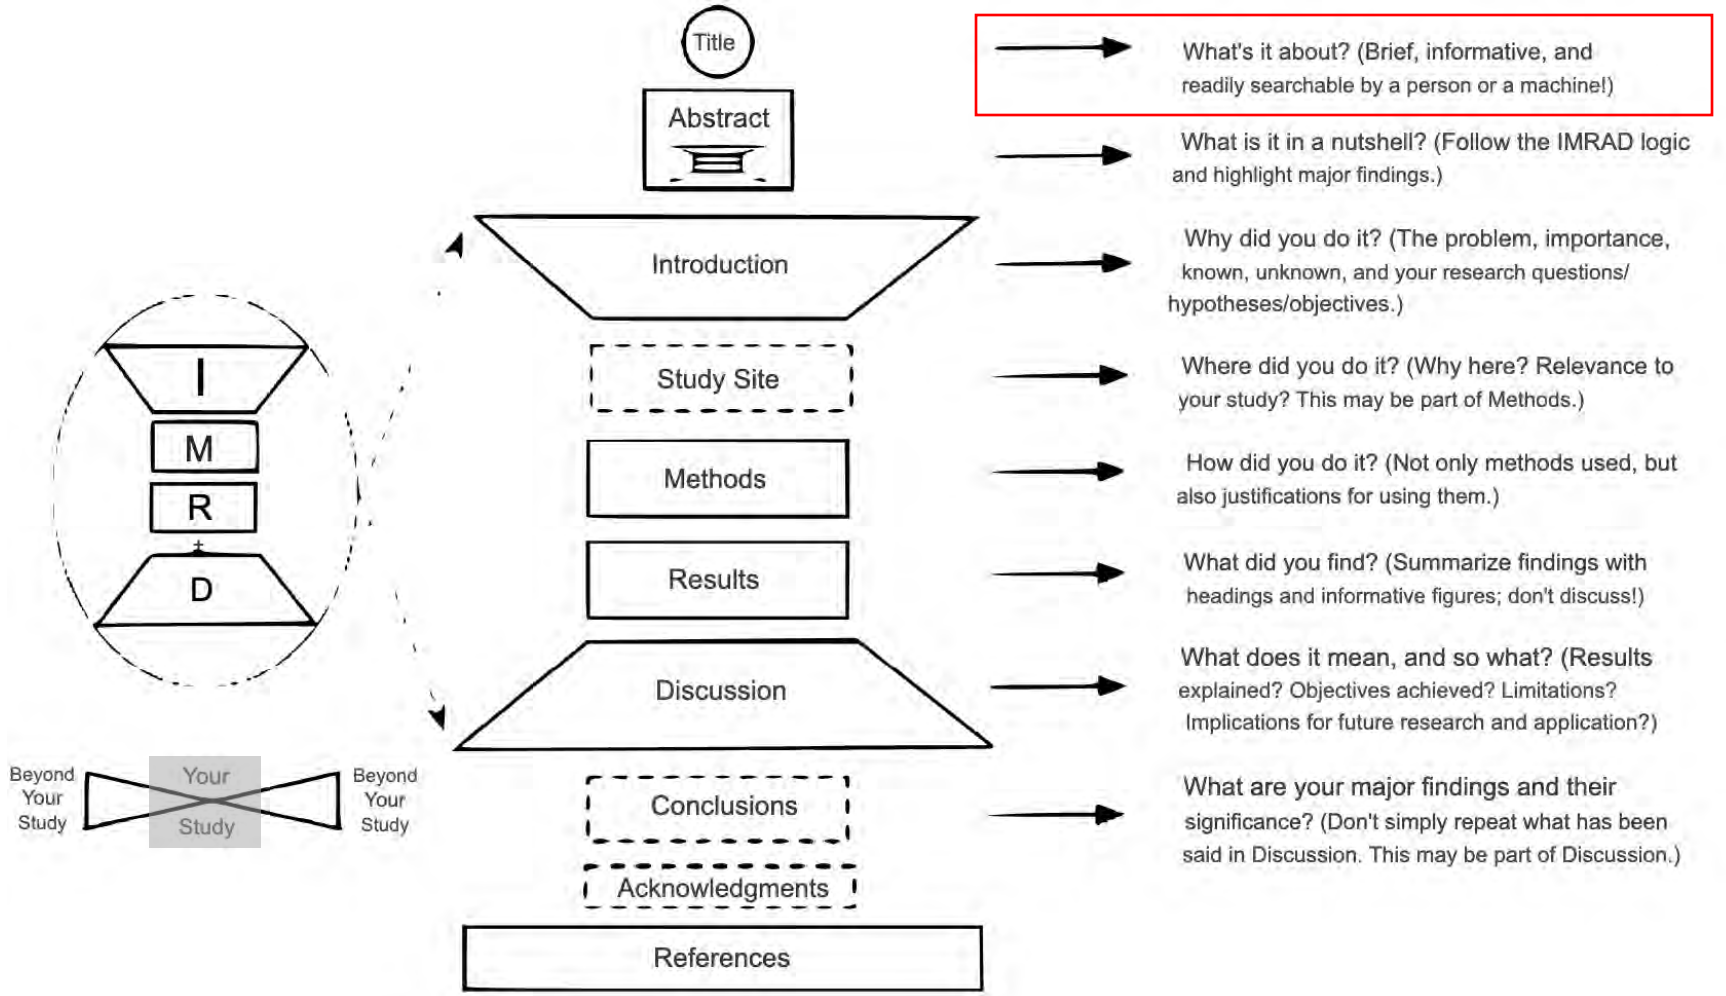

Reinschrift
SCIENCE COMMUNICATION

Wu, J. Improving the writing of research papers: IMRAD and beyond. *Landscape Ecology* **26**, 1345 (2011).

Typical structure of research papers

- Title
- **Abstract**
- Introduction
- Methods and results
- Discussion
- Conclusions
- Detailed methods
- Supplementary Information/Material

Reinschrift
SCIENCE COMMUNICATION

Abstract

- Your story ‘in a nutshell’
- Compact, self-contained summary
  - Background, open problems
    → Why?
  - Specific question addressed
    → What?
  - Main results and their (demonstrated) significance
    → So what?

Typical structure of research papers

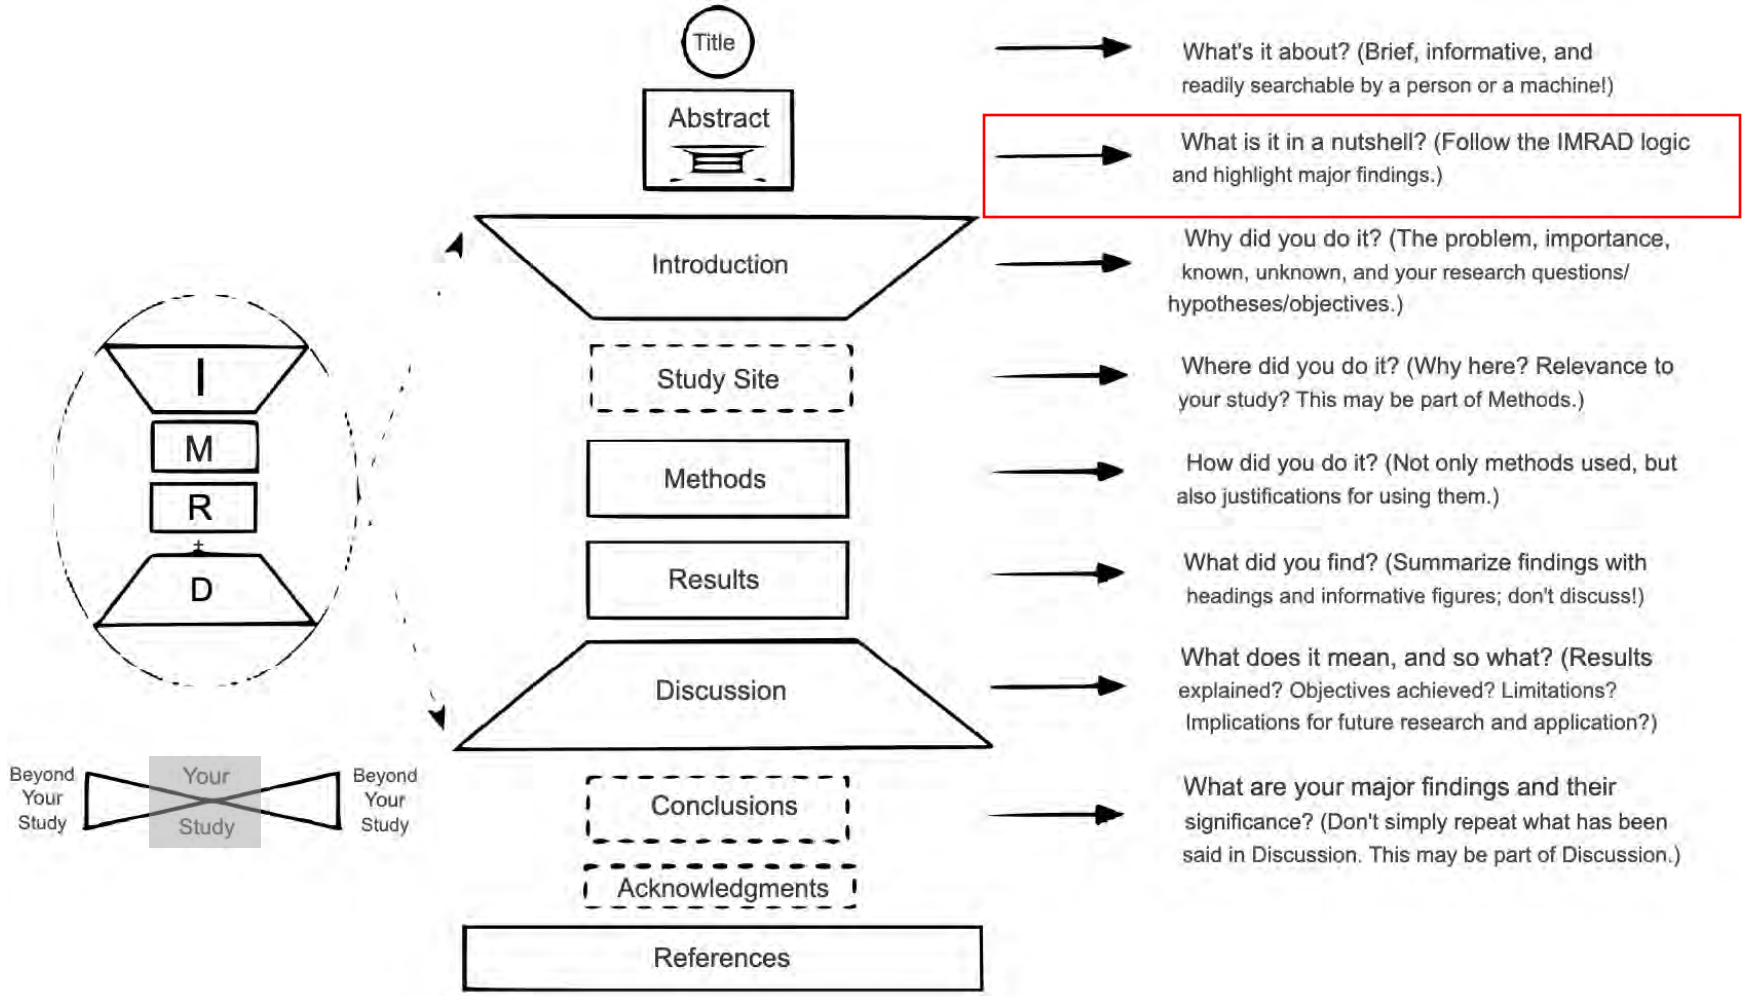

Reinschrift
SCIENCE COMMUNICATION

Wu, J. Improving the writing of research papers: IMRAD and beyond. *Landscape Ecology* **26**, 1345 (2011).

Introduction

- ‘Set the scene’ — for *your* paper and for your target audience

  → Keep in mind what your audience might (or, should) know already about the field

- Explain the (relevant) background to the story to be told *here*
- How does your study relate to (and possibly differ from) precedent work?
- In how far is your approach original?
- Why are your specific results and findings significant?

Introduction

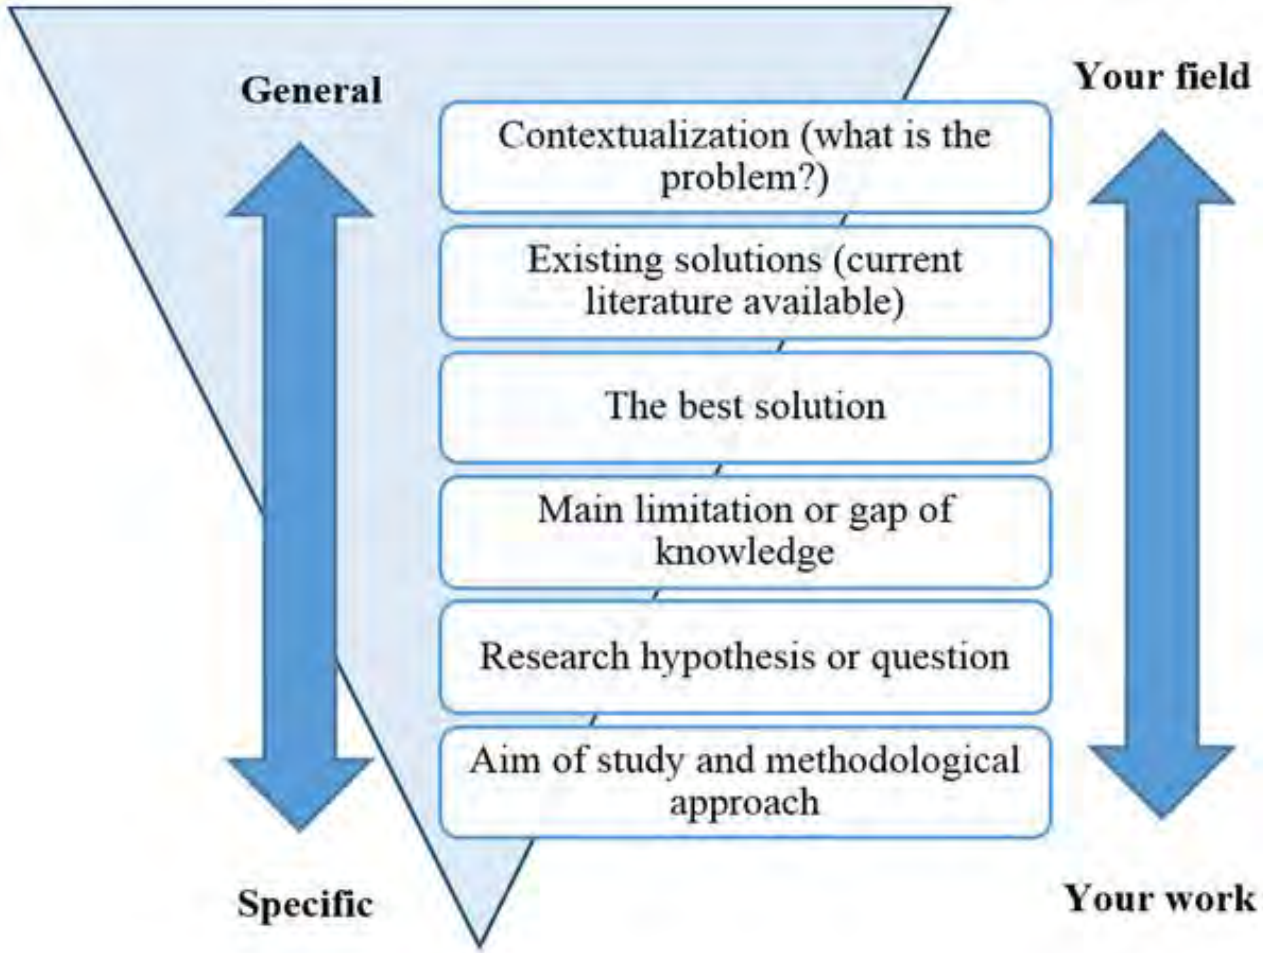

Reinschrift
SCIENCE COMMUNICATION

Bahadoran Z , Jeddi S, Mirmiran P, Ghasemi A.
The Principles of Biomedical Scientific Writing: Introduction.
Int. J. Endocrinol. Metab. **16**, e84795 (2018).

Typical structure of research papers

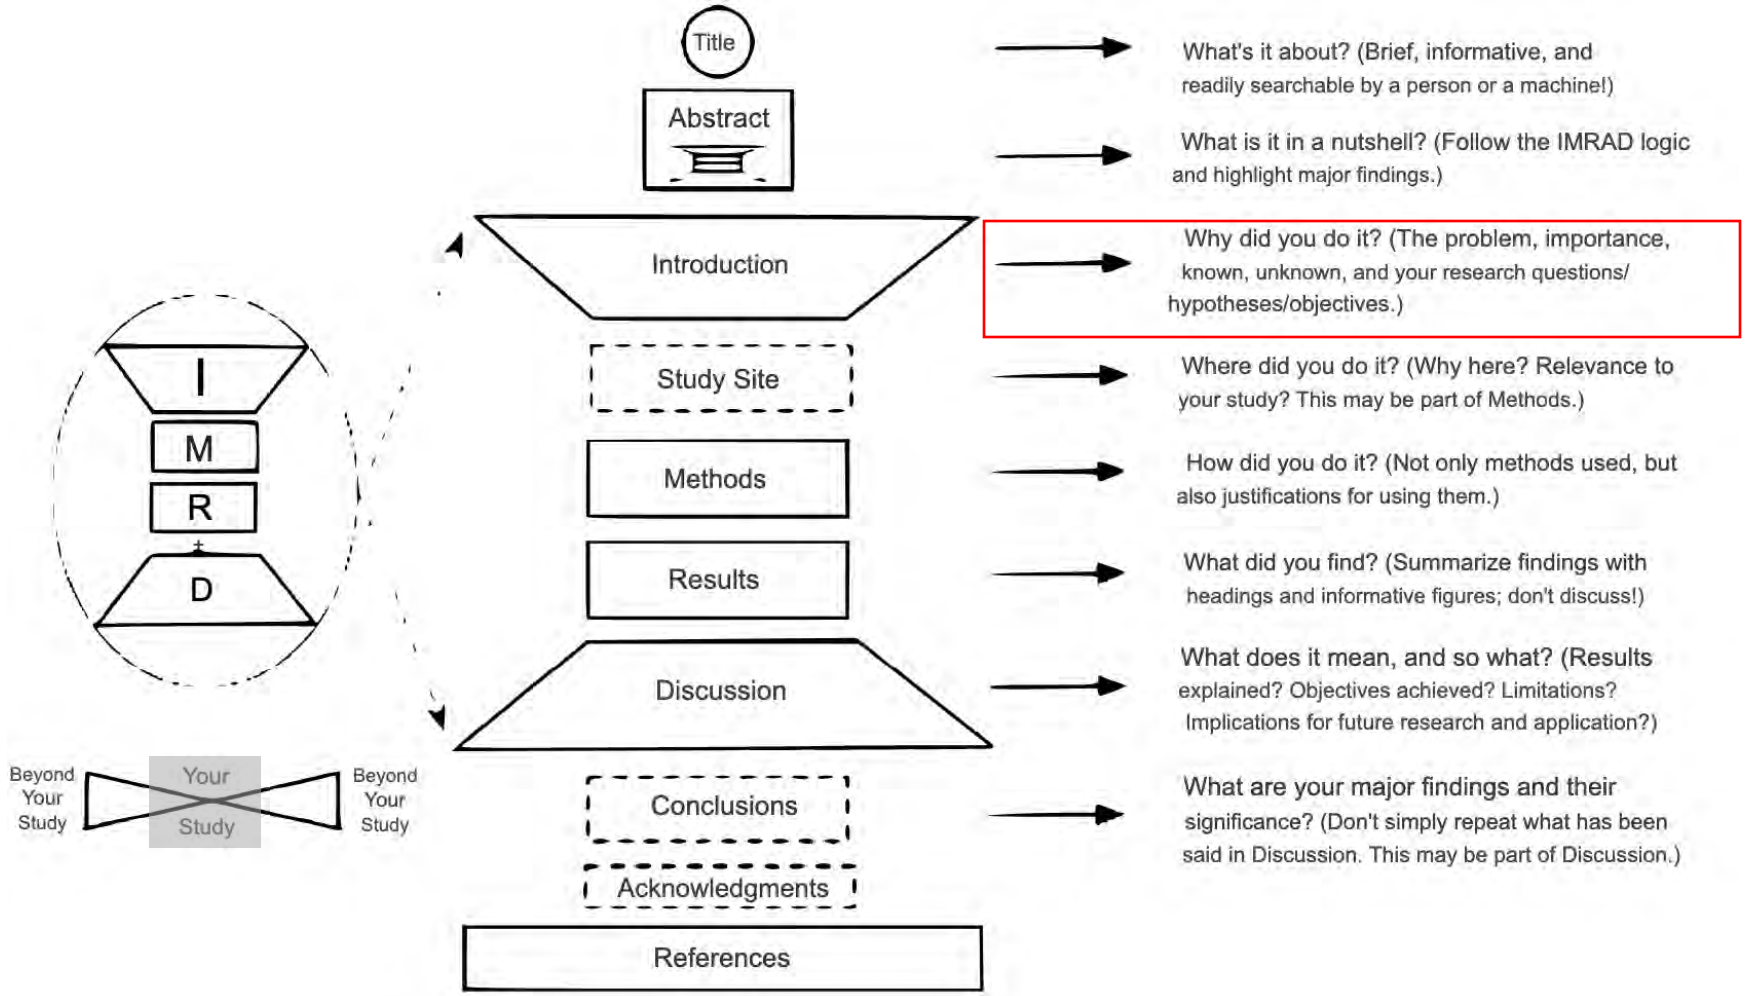

Reinschrift
SCIENCE COMMUNICATION

Wu, J. Improving the writing of research papers: IMRAD and beyond. *Landscape Ecology* **26**, 1345 (2011).

Typical structure of research papers

- Title
- Abstract
- Introduction
- **Methods and results**
- **Discussion**
- **Conclusions**
- Detailed methods
- Supplementary Information/Material

Reinschrift
SCIENCE COMMUNICATION

Results / Discussion / Conclusions

Results

Present the data of your research

→ Not a data dump, but coherent selection to support your conclusions

Discussion

Analyse the data and interpret them

→ Based on your data, make your point

Conclusions

Discuss the potential implications and limitations of the work and give outlook

→ Try to avoid repetition or concluding paragraph that merely summarizes what has been said before

Typical structure of research papers

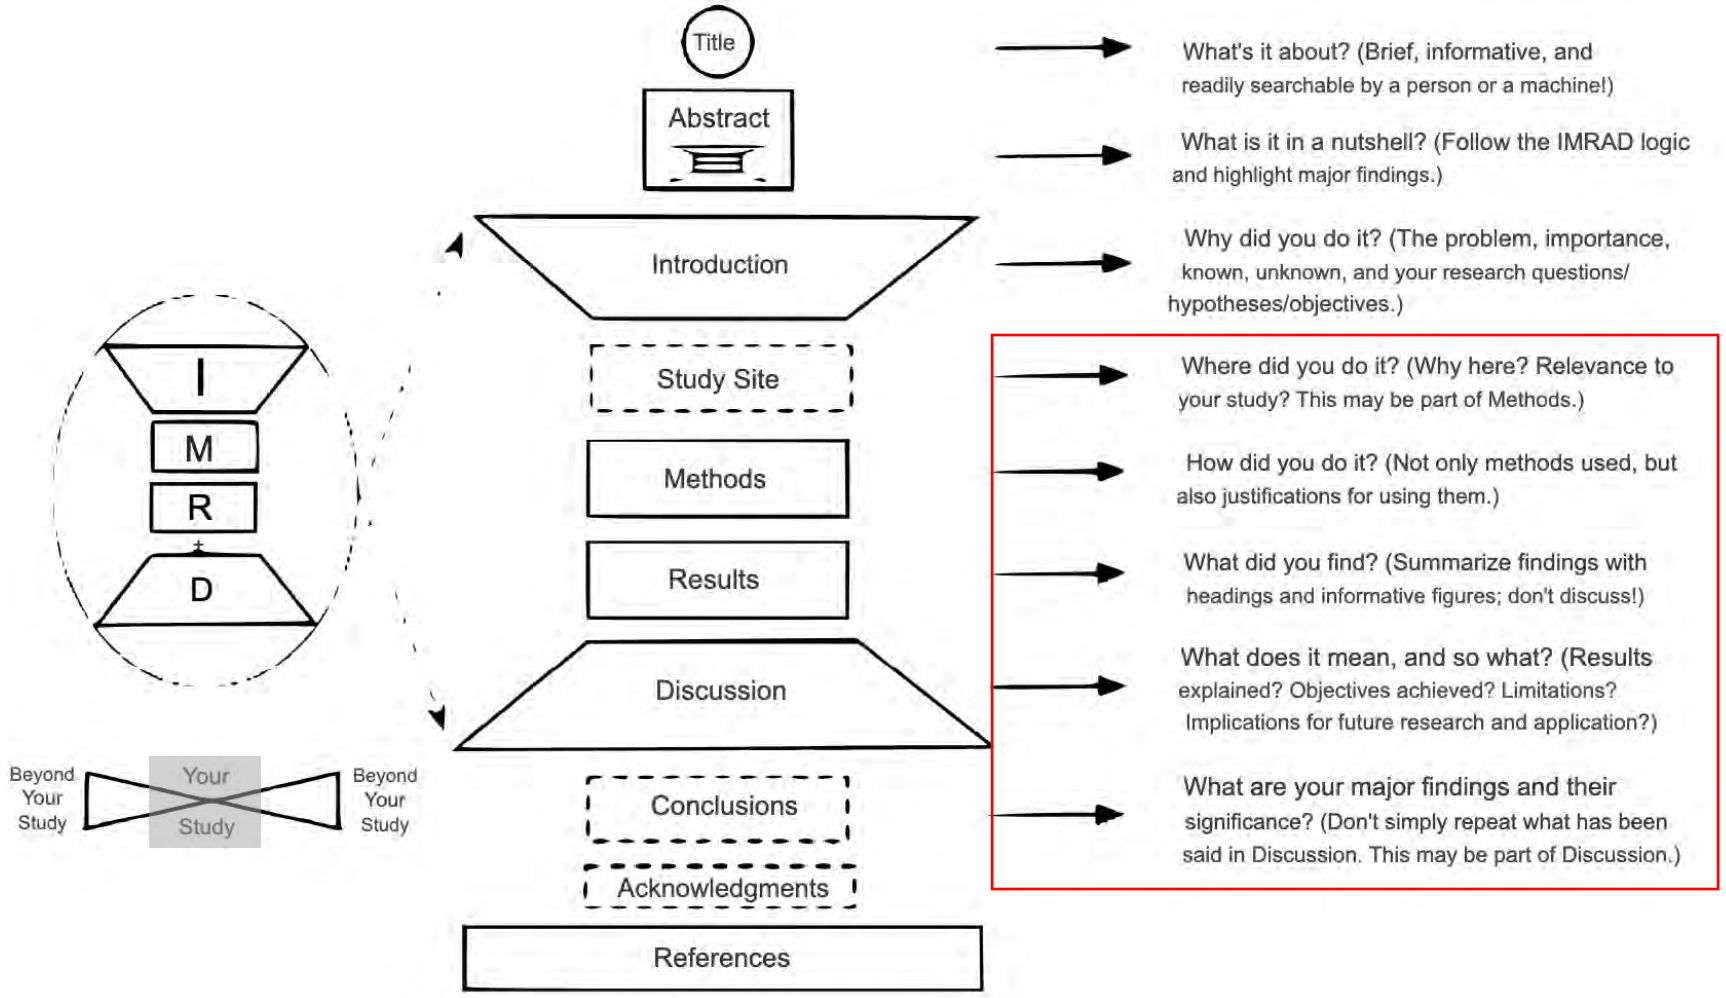

Reinschrift
SCIENCE COMMUNICATION

Wu, J. Improving the writing of research papers: IMRAD and beyond. *Landscape Ecology* **26**, 1345 (2011).

Typical structure of research papers

- Title
- Abstract
- Introduction
- Results
- Discussion
- Conclusions
- **Detailed methods**
- **Supplementary Information/Material**

Reinschrift
SCIENCE COMMUNICATION

Methods and Supplementary Material

Methods

Put a ‘competent person’ in a position to be able to repeat the study.

→ ‘Recipe’

Supplementary Material

Present additional detailed material to support and further substantiate the findings and conclusions made in the main paper.

→ ‘Shopping list’

→ No additional findings or conclusions

Reinschrift
SCIENCE COMMUNICATION

Constructing your text

- Manuscript level
- Section level
- Paragraph level
- Sentence level
- Word level

Structure

Using AI intelligently

- In the attached paper, does the abstract follow an hourglass structure?

- Taking information from elsewhere in the paper, can you suggest a new, more hourglass-like structure for the abstract (fewer than 200 words)?

- Can you please draft a Nature-style abstract following their specific structure?
  1) One or two sentences providing a basic introduction to the field, comprehensible to a scientist in any discipline.
  2) Two to three sentences of more detailed background, comprehensible to scientists in related disciplines.
  3) One sentence clearly stating the general problem being addressed by this particular study.
  4) One sentence summarizing the main result (with the words “here we show” or their equivalent).
  5) Two or three sentences explaining what the main result reveals in direct comparison to what was thought to be the case previously, or how the main result adds to previous knowledge.
  6) One or two sentences to put the results into a more general context.
  7) Two or three sentences to provide a broader perspective, readily comprehensible to a scientist in any discipline, may be included in the first paragraph if the editor considers that the accessibility of the paper is significantly enhanced by their inclusion. Under these circumstances, the length of the paragraph can be up to 300 words.

Using AI intelligently

- Taking the role of a *Nature* editor, would you consider the novelty and significance as being sufficient for sending this manuscript out to peer review?

- Let's continue on the *Nature* path nonetheless. Can you please sketch out a suitable structure for a revised full manuscript, based on the information in the original paper?

Constructing your text

- Manuscript level
- Section level
- Paragraph level
- Sentence level
- Word level

Structure

Flow

Creating flow

Structure: Backbone of your story

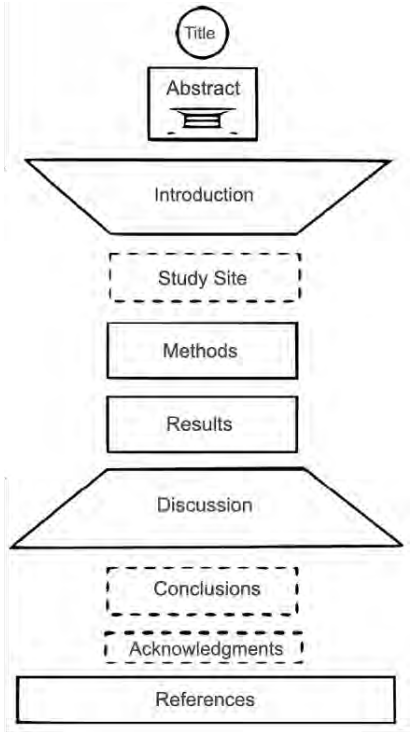

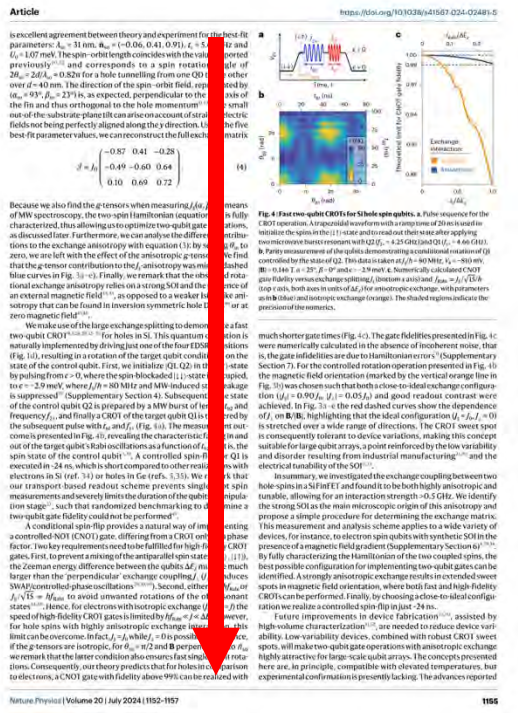

Left – From: Wu, J. *Landscape Ecology* **26**, 1345 (2011).
Right – *Nature Physics* **20**, 1152 (2024).

Creating flow

Flow: Construct a coherent narrative

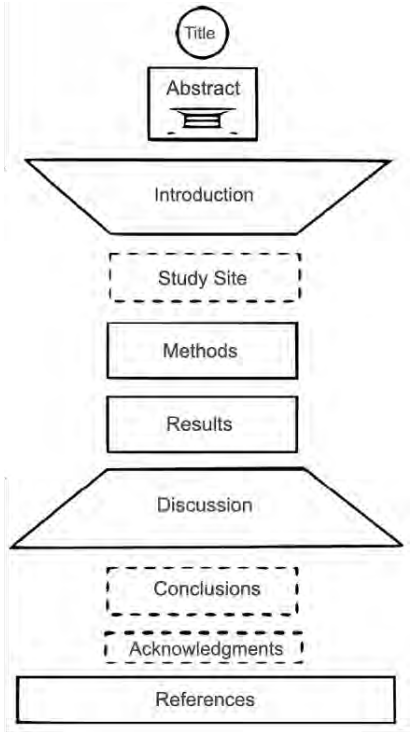

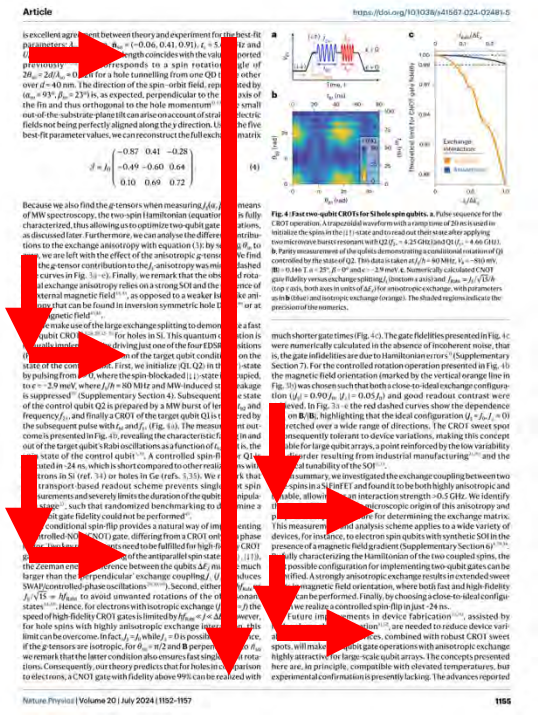

Left – From: Wu, J. *Landscape Ecology* **26**, 1345 (2011).
Right – *Nature Physics* **20**, 1152 (2024).

Constructing your text

- Manuscript level
- Section level
- Paragraph level
- Sentence level
- Word level

Creating flow

Paragraph level: Guiding your reader through the text

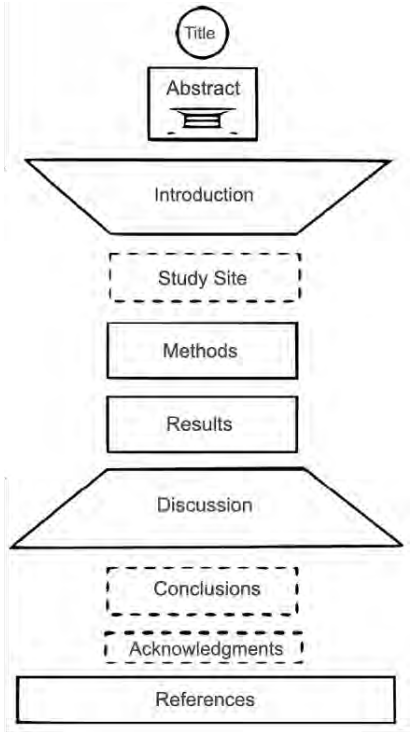

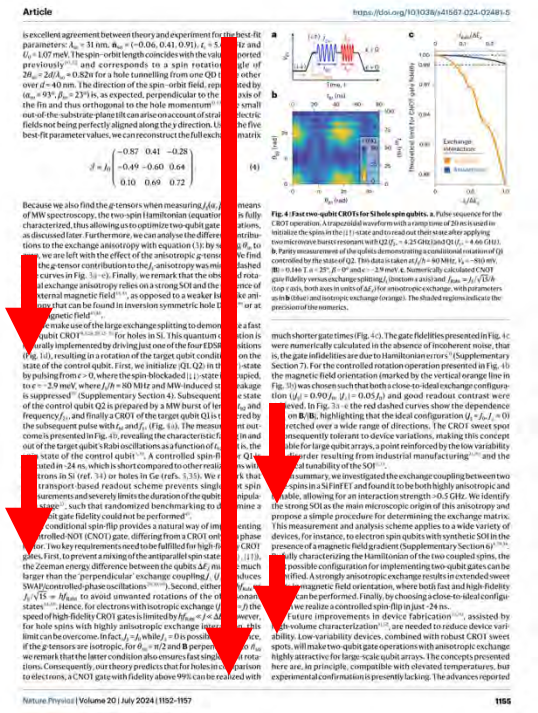

Left – From: Wu, J. *Landscape Ecology* **26**, 1345 (2011).
Right – *Nature Physics* **20**, 1152 (2024).

Guiding your reader through the text

Article

Highly accurate protein structure prediction with AlphaFold

https://doi.org/10.1038/s41586-021-03819-2

Received: 11 May 2021

Accepted: 12 July 2021

Published online: 15 July 2021

Open access

Check for updates

John Jumper[1,4✉], Richard Evans[1,4], Alexander Pritzel[1,4], Tim Green[1,4], Michael Figurnov[1,4], Olaf Ronneberger[1,4], Kathryn Tunyasuvunakool[1,4], Russ Bates[1,4], Augustin Žídek[1,4], Anna Potapenko[1,4], Alex Bridgland[1,4], Clemens Meyer[1,4], Simon A. A. Kohl[1,4], Andrew J. Ballard[1,4], Andrew Cowie[1,4], Bernardino Romera-Paredes[1,4], Stanislav Nikolov[1,4], Rishub Jain[1,4], Jonas Adler[1], Trevor Back[1], Stig Petersen[1], David Reiman[1], Ellen Clancy[1], Michal Zielinski[1], Martin Steinegger[2,3], Michalina Pacholska[1], Tamas Berghammer[1], Sebastian Bodenstein[1], David Silver[1], Oriol Vinyals[1], Andrew W. Senior[1], Koray Kavukcuoglu[1], Pushmeet Kohli[1] & Demis Hassabis[1,4✉]

Proteins are essential to life, and understanding their structure can facilitate a mechanistic understanding of their function. Through an enormous experimental effort[1–4], the structures of around 100,000 unique proteins have been determined[5], but this represents a small fraction of the billions of known protein sequences[6,7]. Structural coverage is bottlenecked by the months to years of painstaking effort required to determine a single protein structure. Accurate computational approaches are needed to address this gap and to enable large-scale structural bioinformatics. Predicting the three-dimensional structure that a protein will adopt based solely on its amino acid sequence—the structure prediction component of the 'protein folding problem'[8]—has been an important open research problem for more than 50 years[9]. Despite recent progress[10–14], existing methods fall far short of atomic accuracy, especially when no homologous structure is available. Here we provide the first computational method that can regularly predict protein structures with atomic accuracy even in cases in which no similar structure is known. We validated an entirely redesigned version of our neural network-based model, AlphaFold, in the challenging 14th Critical Assessment of protein Structure Prediction (CASP14)[15], demonstrating accuracy competitive with experimental structures in a majority of cases and greatly outperforming other methods. Underpinning the latest version of AlphaFold is a novel machine learning approach that incorporates physical and biological knowledge about protein structure, leveraging multi-sequence alignments, into the design of the deep learning algorithm.

The development of computational methods to predict three-dimensional (3D) protein structures from the protein sequence has proceeded along two complementary paths that focus on either the physical interactions or the evolutionary history. The physical interaction programme heavily integrates our understanding of molecular driving forces into either thermodynamic or kinetic simulation of protein physics[16] or statistical approximations thereof[17]. Although theoretically very appealing, this approach has proved highly challenging for even moderate-sized proteins due to the computational intractability of molecular simulation, the context dependence of protein stability and the difficulty of producing sufficiently accurate models of protein physics. The evolutionary programme has provided an alternative in recent years, in which the constraints on protein structure are derived from bioinformatics analysis of the evolutionary history of proteins, homology to solved structures[18,19] and pairwise evolutionary correlations[20–24]. This bioinformatics approach has benefited greatly from the steady growth of experimental protein structures deposited in the Protein Data Bank (PDB)[5], the explosion of genomic sequencing and the rapid development of deep learning techniques to interpret these correlations. Despite these advances, contemporary physical and evolutionary-history-based approaches produce predictions that are far short of experimental accuracy in the majority of cases in which a close homologue has not been solved experimentally and this has limited their utility for many biological applications.

In this study, we develop the first, to our knowledge, computational approach capable of predicting protein structures to near experimental accuracy in a majority of cases. The neural network AlphaFold that we developed was entered into the CASP14 assessment (May–July 2020; entered under the team name 'AlphaFold2' and a completely different model from our CASP13 AlphaFold system[10]). The CASP assessment is carried out biennially using recently solved structures that have not been deposited in the PDB or publicly disclosed so that it is a blind test

[1]DeepMind, London, UK. [2]School of Biological Sciences, Seoul National University, Seoul, South Korea. [3]Artificial Intelligence Institute, Seoul National University, Seoul, South Korea. [4]These authors contributed equally: John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, Alex Bridgland, Clemens Meyer, Simon A. A. Kohl, Andrew J. Ballard, Andrew Cowie, Bernardino Romera-Paredes, Stanislav Nikolov, Rishub Jain, Demis Hassabis
✉e-mail: jumper@deepmind.com; dhcontact@deepmind.com

Nature **596**, 583–589 (2021).

Reinschrift
SCIENCE COMMUNICATION

Guiding your reader through the text

The development of computational methods to predict three-dimensional (3D) protein structures from the protein sequence has proceeded along two complementary paths that focus on either the physical interactions or the evolutionary history. The physical interaction programme heavily integrates our understanding of molecular driving forces into either thermodynamic or kinetic simulation of protein physics[16] or statistical approximations thereof[17]. Although theoretically very appealing, this approach has proved highly challenging for even moderate-sized proteins due to the computational intractability of molecular simulation, the context dependence of protein stability and the difficulty of producing sufficiently accurate models of protein physics. The evolutionary programme has provided an alternative in recent years, in which the constraints on protein structure are derived from bioinformatics analysis of the evolutionary history of proteins, homology to solved structures[18,19] and pairwise evolutionary correlations[20–24]. This bioinformatics approach has benefited greatly from the steady growth of experimental protein structures deposited in the Protein Data Bank (PDB)[5], the explosion of genomic sequencing and the rapid development of deep learning techniques to interpret these correlations. Despite these advances, contemporary physical and evolutionary-history-based approaches produce predictions that are far short of experimental accuracy in the majority of cases in which a close homologue has not been solved experimentally and this has limited their utility for many biological applications.

In this study, we develop the first, to our knowledge, computational approach capable of predicting protein structures to near experimental accuracy in a majority of cases. The neural network AlphaFold that we developed was entered into the CASP14 assessment (May–July 2020; entered under the team name 'AlphaFold2' and a completely different model from our CASP13 AlphaFold system[10]). The CASP assessment is carried out biennially using recently solved structures that have not been deposited in the PDB or publicly disclosed so that it is a blind test

Nature **596**, 583–589 (2021).

Guiding your reader through the text

for the participating methods, and has long served as the gold-standard assessment for the accuracy of structure prediction[25,26].

In CASP14, AlphaFold structures were vastly more accurate than competing methods. AlphaFold structures had a median backbone accuracy of 0.96 Å r.m.s.d.$_{95}$ (Cα root-mean-square deviation at 95% residue coverage) (95% confidence interval = 0.85–1.16 Å) whereas the next best performing method had a median backbone accuracy of 2.8 Å r.m.s.d.$_{95}$ (95% confidence interval = 2.7–4.0 Å) (measured on CASP domains; see Fig. 1a for backbone accuracy and Supplementary Fig. 14 for all-atom accuracy). As a comparison point for this accuracy, the width of a carbon atom is approximately 1.4 Å. In addition to very accurate domain structures (Fig. 1b), AlphaFold is able to produce highly accurate side chains (Fig. 1c) when the backbone is highly accurate and considerably improves over template-based methods even when strong templates are available. The all-atom accuracy of AlphaFold was 1.5 Å r.m.s.d.$_{95}$ (95% confidence interval = 1.2–1.6 Å) compared with the 3.5 Å r.m.s.d.$_{95}$ (95% confidence interval = 3.1–4.2 Å) of the best alternative method. Our methods are scalable to very long proteins with accurate domains and domain-packing (see Fig. 1d for the prediction of a 2,180-residue protein with no structural homologues). Finally, the model is able to provide precise, per-residue estimates of its reliability that should enable the confident use of these predictions.

We demonstrate in Fig. 2a that the high accuracy that AlphaFold demonstrated in CASP14 extends to a large sample of recently released PDB structures; in this dataset, all structures were deposited in the PDB after our training data cut-off and are analysed as full chains (see Methods, Supplementary Fig. 15 and Supplementary Table 6 for more details). Furthermore, we observe high side-chain accuracy when the backbone prediction is accurate (Fig. 2b) and we show that our confidence measure, the predicted local-distance difference test (pLDDT), reliably predicts the Cα local-distance difference test (lDDT-Cα) accuracy of the corresponding prediction (Fig. 2c). We also find that the global superposition metric template modelling score (TM-score)[27] can be accurately estimated (Fig. 2d). Overall, these analyses validate that the high accuracy and reliability of AlphaFold on CASP14 proteins also transfers to an uncurated collection of recent PDB submissions, as would be expected (see Supplementary Methods 1.15 and Supplementary Fig. 11 for confirmation that this high accuracy extends to new folds).

The AlphaFold network

AlphaFold greatly improves the accuracy of structure prediction by incorporating novel neural network architectures and training procedures based on the evolutionary, physical and geometric constraints of protein structures. In particular, we demonstrate a new architecture to jointly embed multiple sequence alignments (MSAs) and pairwise features, a new output representation and associated loss that enable accurate end-to-end structure prediction, a new equivariant attention

Nature **596**, 583–589 (2021).

‘Signposting’

Help the reader navigate the text by setting markers (leading sentences)

Introducing background information
- "Historically, ..."
- "In the context of ..."

Presenting methods or procedures
- "Our approach involves ..."
- "By employing ..."

Highlighting results or findings
- "Our results indicate ..."
- "We observed that ..."

Discussing implications or interpretations
- "These results suggest ..."
- "From this data, we infer ..."

Shifting to new topics or perspectives
- "By contrast, ..."
- "Having established X, we can demonstrate that ... "

Summarizing or concluding
- "In conclusion, ..."
- "Overall, ..."

Highlighting limitations or future directions
- "One limitation of this study is ..."
- "Future research should explore ..."
- "Further investigation is needed to ..."

Reinschrift
SCIENCE COMMUNICATION

Other examples: https://canvas.hull.ac.uk/courses/371/pages/signposting

‘Signposting’

But never use ‘blank’ posts

Vague or uninformative introductions
- "In this section, …"
- "It is important to note that ..."
- "As we move on, …"

Redundant or repetitive phrases
- "As mentioned earlier, ..."
- "It is clear that …"
- "It goes without saying that …"

Ambiguous transitions
- "Moreover, ..."
- "Additionally, ..."
- "Furthermore, ..."

Clichéd or overused phrases
- "The fact of the matter is …"
- "With that being said, ..."
- "All things considered, …"

Reinschrift
SCIENCE COMMUNICATION

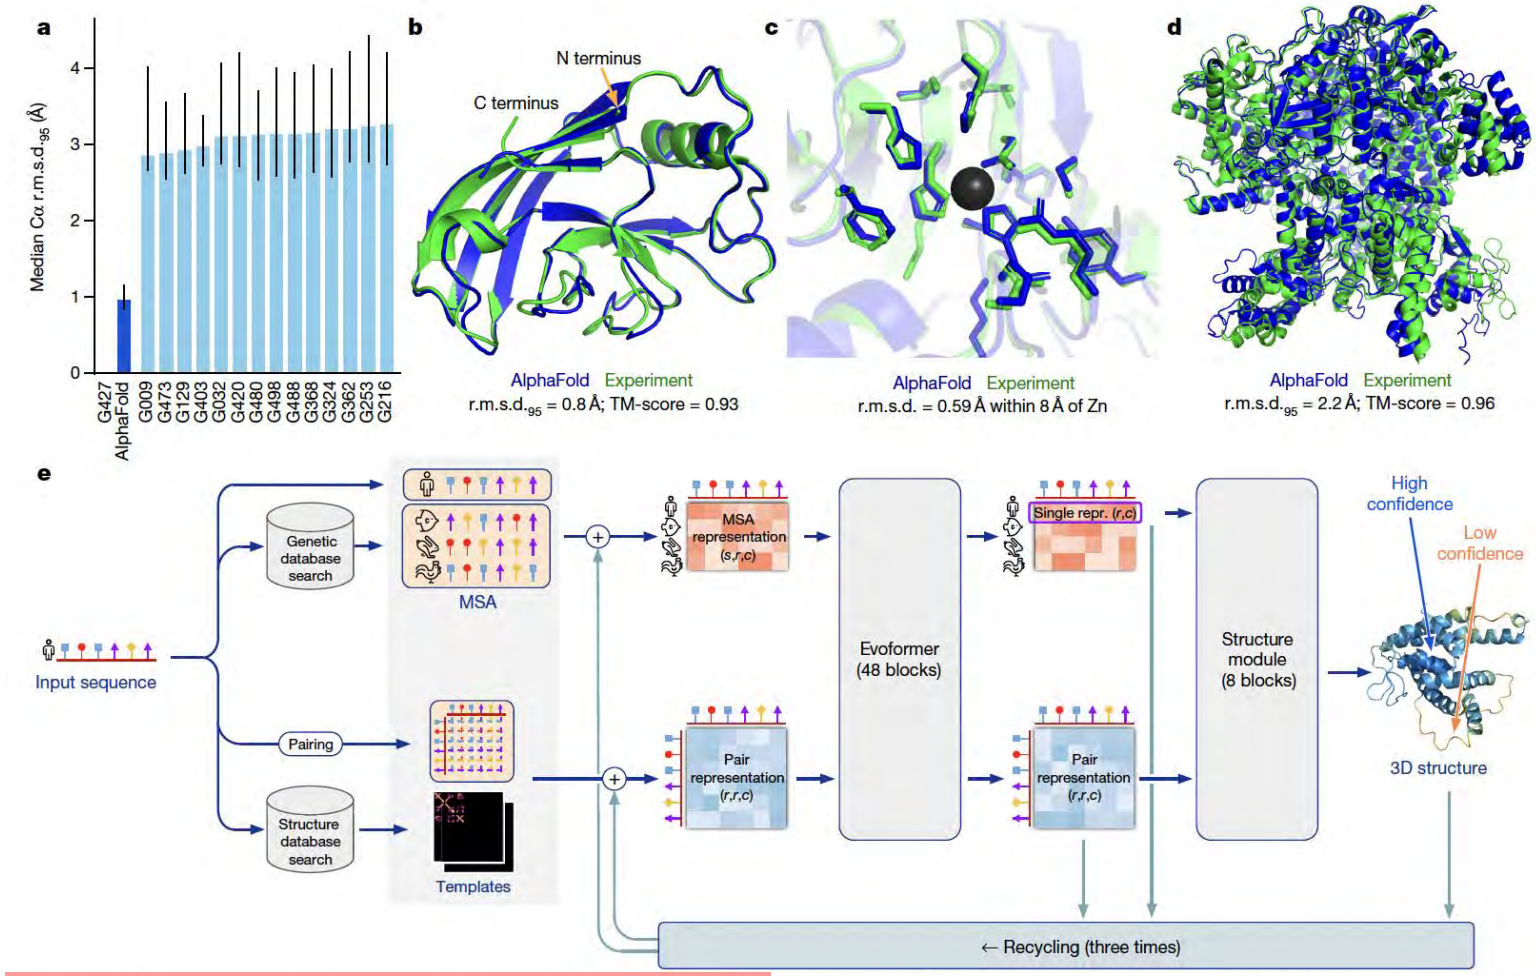

Fig. 1 | AlphaFold produces highly accurate structures. a, The performance of AlphaFold on the CASP14 dataset ($n = 87$ protein domains) relative to the top-15 entries (out of 146 entries), group numbers correspond to the numbers assigned to entrants by CASP. Data are median and the 95% confidence interval of the median, estimated from 10,000 bootstrap samples. **b**, Our prediction of CASP14 target T1049 (PDB 6Y4F, blue) compared with the true (experimental) structure (green). Four residues in the C terminus of the crystal structure are *B*-factor outliers and are not depicted. **c**, CASP14 target T1056 (PDB 6YJ1). An example of a well-predicted zinc-binding site (AlphaFold has accurate side chains even though it does not explicitly predict the zinc ion). **d**, CASP target T1044 (PDB 6VR4)—a 2,180-residue single chain—was predicted with correct domain packing (the prediction was made after CASP using AlphaFold without intervention). **e**, Model architecture. Arrows show the information flow among the various components described in this paper. Array shapes are shown in parentheses with *s*, number of sequences (N_{seq} in the main text); *r*, number of residues (N_{res} in the main text); *c*, number of channels.

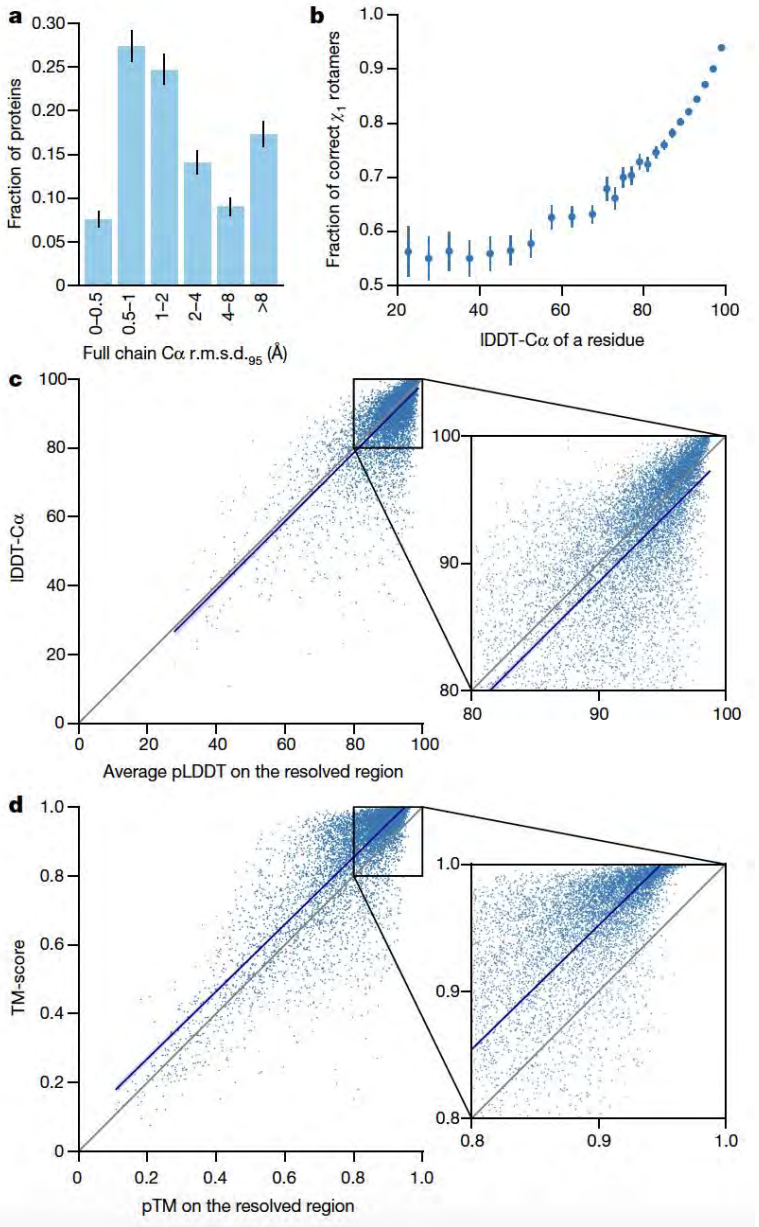

Fig. 2 | Accuracy of AlphaFold on recent PDB structures. The analysed structures are newer than any structure in the training set. Further filtering is applied to reduce redundancy (see Methods). **a**, Histogram of backbone r.m.s.d. for full chains (Cα r.m.s.d. at 95% coverage). Error bars are 95% confidence intervals (Poisson). This dataset excludes proteins with a template (identified by hmmsearch) from the training set with more than 40% sequence identity covering more than 1% of the chain ($n = 3{,}144$ protein chains). The overall median is 1.46 Å (95% confidence interval = 1.40–1.56 Å). Note that this measure will be highly sensitive to domain packing and domain accuracy; a high r.m.s.d. is expected for some chains with uncertain packing or packing errors. **b**, Correlation between backbone accuracy and side-chain accuracy. Filtered to structures with any observed side chains and resolution better than 2.5 Å ($n = 5{,}317$ protein chains); side chains were further filtered to B-factor <30 Å^2. A rotamer is classified as correct if the predicted torsion angle is within 40°. Each point aggregates a range of lDDT-Cα, with a bin size of 2 units above 70 lDDT-Cα and 5 units otherwise. Points correspond to the mean accuracy; error bars are 95% confidence intervals (Student t-test) of the mean on a per-residue basis. **c**, Confidence score compared to the true accuracy on chains. Least-squares linear fit lDDT-Cα = 0.997 × pLDDT − 1.17 (Pearson's $r = 0.76$). $n = 10{,}795$ protein chains. The shaded region of the linear fit represents a 95% confidence interval estimated from 10,000 bootstrap samples. In the companion paper[39], additional quantification of the reliability of pLDDT as a confidence measure is provided. **d**, Correlation between pTM and full chain TM-score. Least-squares linear fit TM-score = 0.98 × pTM + 0.07 (Pearson's $r = 0.85$). $n = 10{,}795$ protein chains. The shaded region of the linear fit represents a 95% confidence interval estimated from 10,000 bootstrap samples.

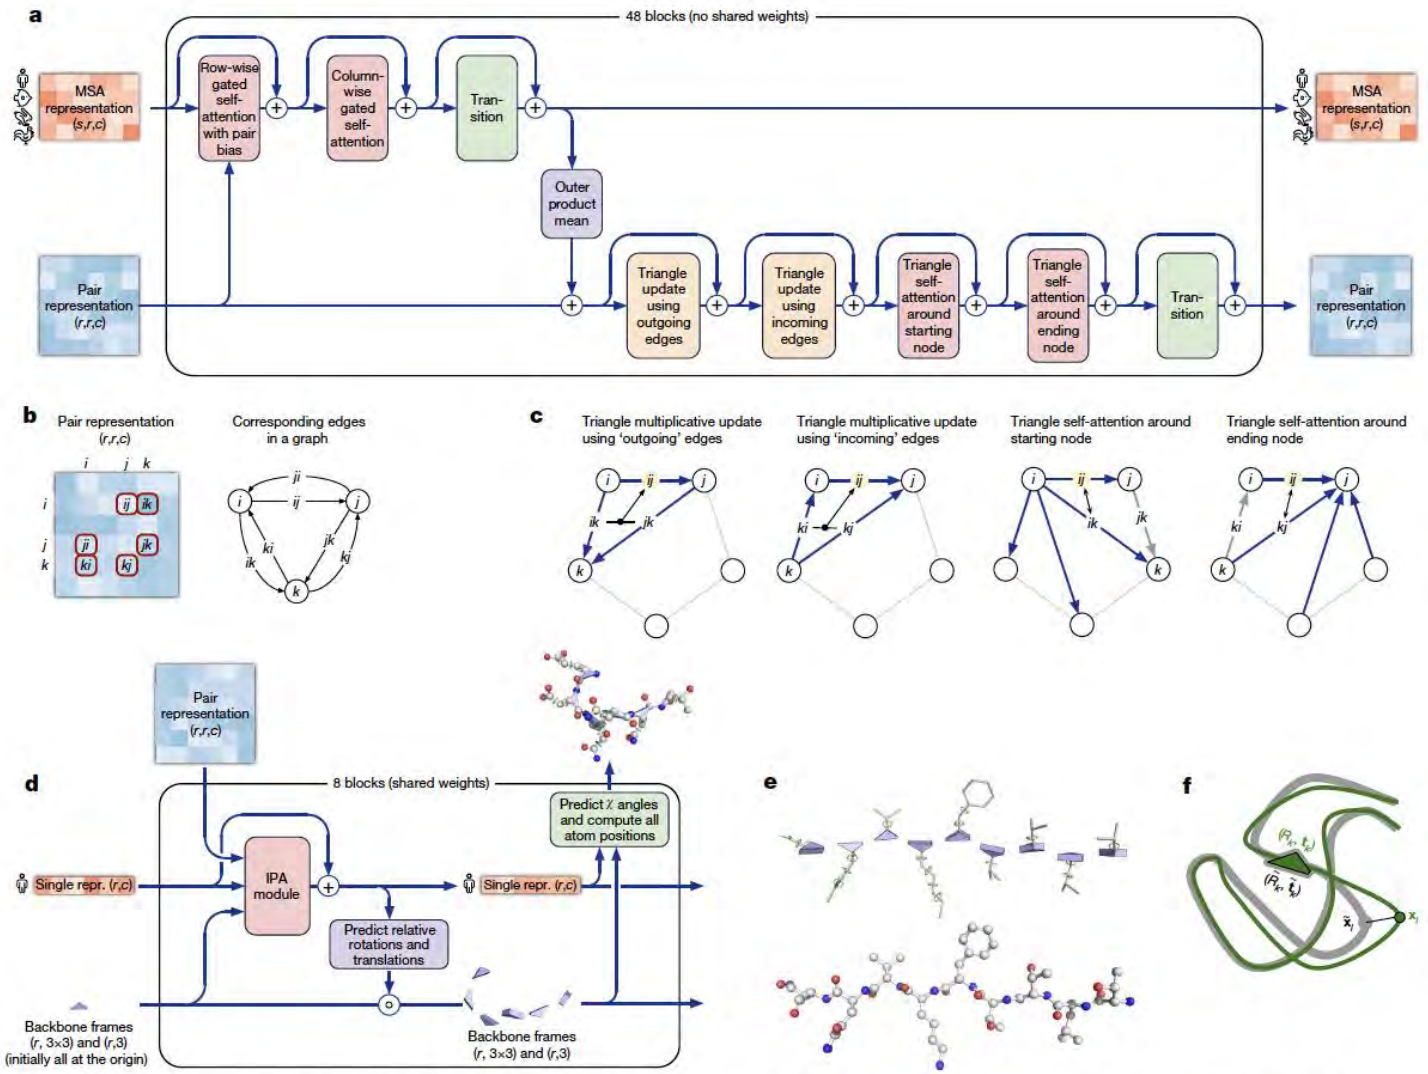

Fig. 3 | Architectural details. a, Evoformer block. Arrows show the information flow. The shape of the arrays is shown in parentheses. **b**, The pair representation interpreted as directed edges in a graph. **c**, Triangle multiplicative update and triangle self-attention. The circles represent residues. Entries in the pair representation are illustrated as directed edges and in each diagram, the edge being updated is *ij*. **d**, Structure module including Invariant point attention (IPA) module. The single representation is a copy of the first row of the MSA representation. **e**, Residue gas: a representation of each residue as one free-floating rigid body for the backbone (blue triangles) and χ angles for the side chains (green circles). The corresponding atomic structure is shown below. **f**, Frame aligned point error (FAPE). Green, predicted structure; grey, true structure; $(R_k, \mathbf{t}_k)$, frames; $\mathbf{x}_i$, atom positions.

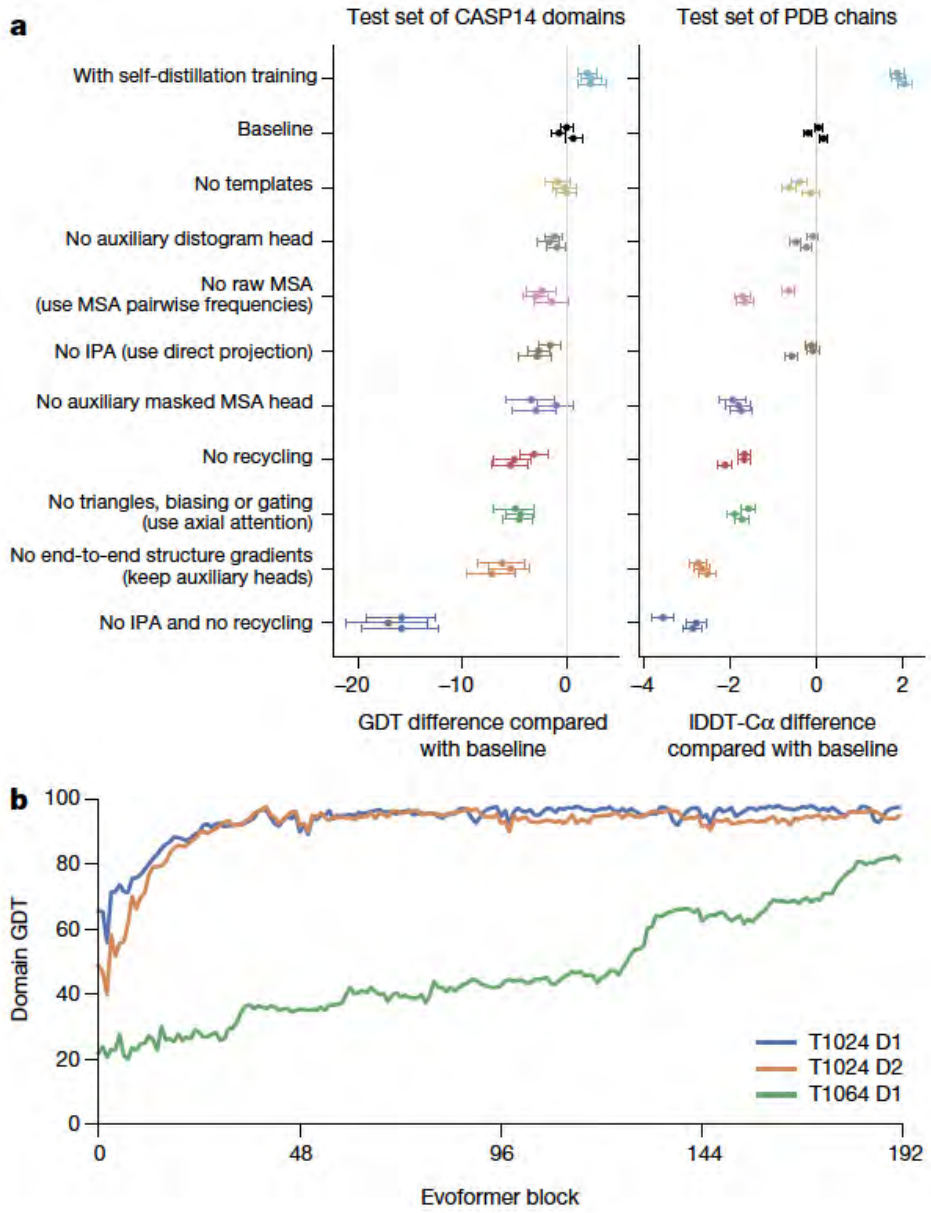

Fig. 4 | Interpreting the neural network. a, Ablation results on two target sets: the CASP14 set of domains ($n = 87$ protein domains) and the PDB test set of chains with template coverage of ≤30% at 30% identity ($n = 2,261$ protein chains). Domains are scored with GDT and chains are scored with lDDT-Cα. The ablations are reported as a difference compared with the average of the three baseline seeds. Means (points) and 95% bootstrap percentile intervals (error bars) are computed using bootstrap estimates of 10,000 samples. **b**, Domain GDT trajectory over 4 recycling iterations and 48 Evoformer blocks on CASP14 targets LmrP (T1024) and Orf8 (T1064) where D1 and D2 refer to the individual domains as defined by the CASP assessment. Both T1024 domains obtain the correct structure early in the network, whereas the structure of T1064 changes multiple times and requires nearly the full depth of the network to reach the final structure. Note, 48 Evoformer blocks comprise one recycling iteration.

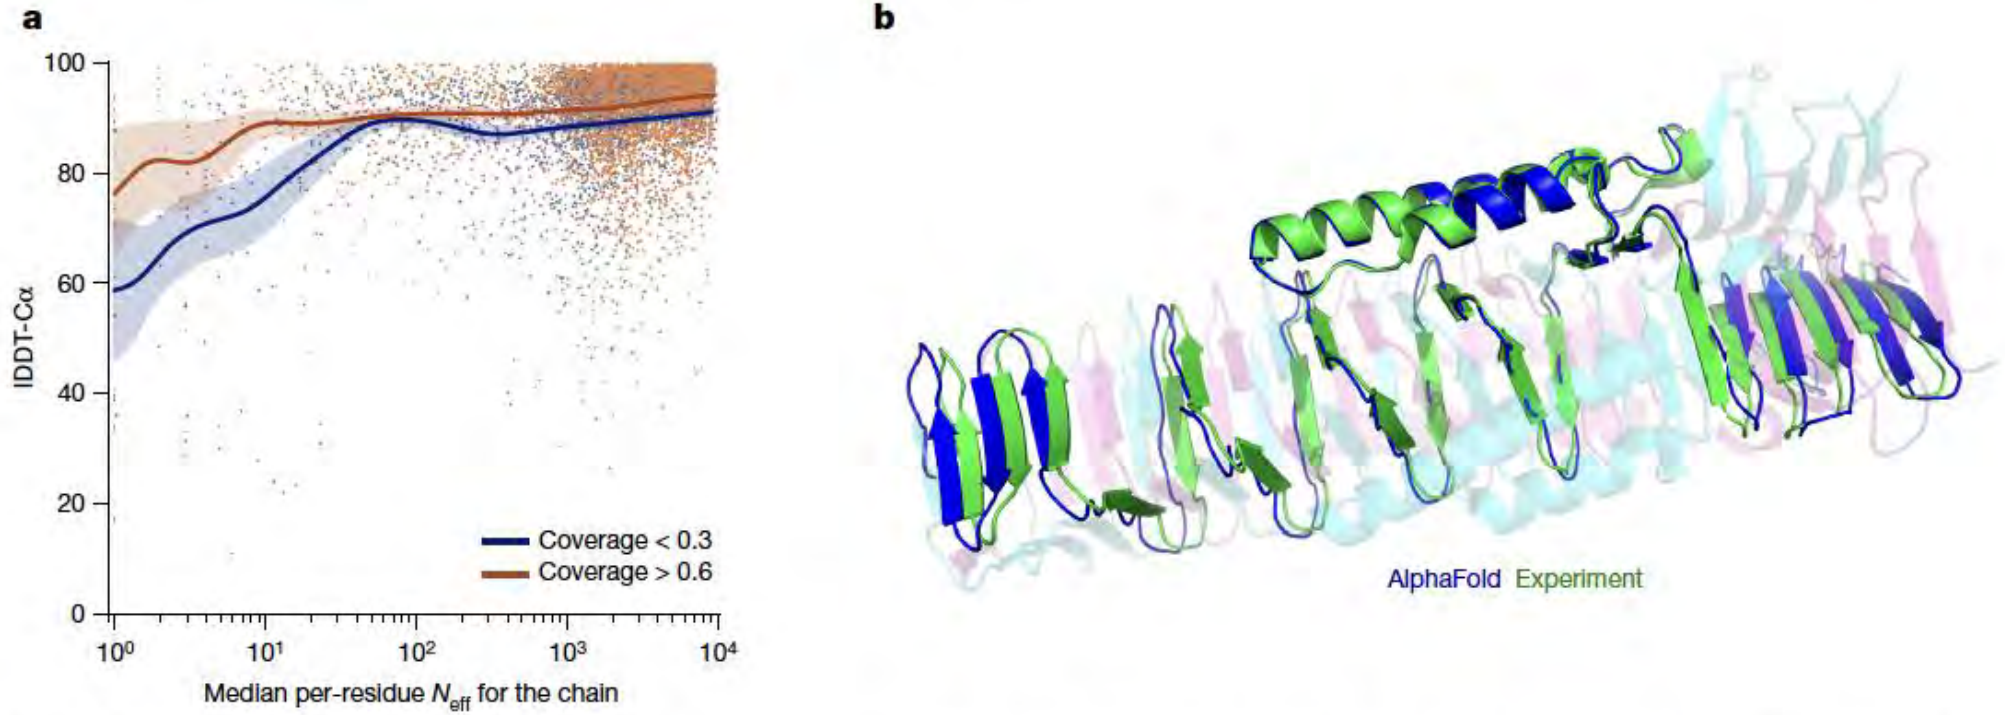

Fig. 5 | Effect of MSA depth and cross-chain contacts. a, Backbone accuracy (lDDT-Cα) for the redundancy-reduced set of the PDB after our training data cut-off, restricting to proteins in which at most 25% of the long-range contacts are between different heteromer chains. We further consider two groups of proteins based on template coverage at 30% sequence identity: covering more than 60% of the chain ($n = 6,743$ protein chains) and covering less than 30% of the chain ($n = 1,596$ protein chains). MSA depth is computed by counting the number of non-gap residues for each position in the MSA (using the N_{eff} weighting scheme; see Methods for details) and taking the median across residues. The curves are obtained through Gaussian kernel average smoothing (window size is 0.2 units in $\log_{10}(N_{eff})$); the shaded area is the 95% confidence interval estimated using bootstrap of 10,000 samples. **b**, An intertwined homotrimer (PDB 6SK0) is correctly predicted without input stoichiometry and only a weak template (blue is predicted and green is experimental).

Constructing your text

- Manuscript level
- Section level
- Paragraph level
- Sentence level
- Word level

Creating flow

How to construct a coherent narrative

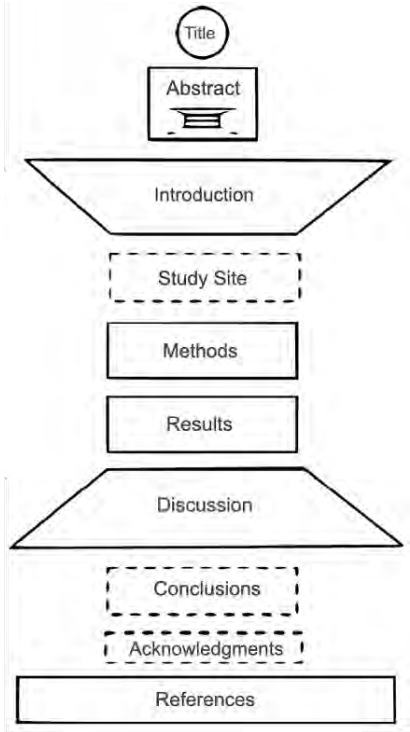

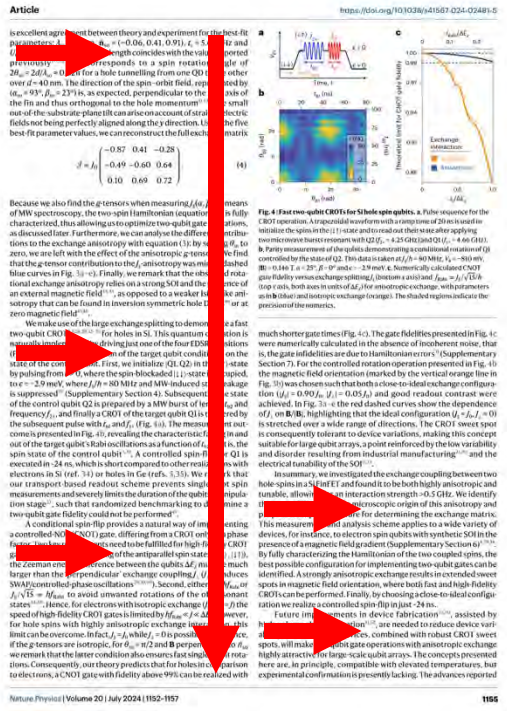

Left – From: Wu, J. *Landscape Ecology* **26**, 1345 (2011).
Right – *Nature Physics* **20**, 1152 (2024).

Reinschrift
SCIENCE COMMUNICATION

One sentence, one message

Three features we have observed—the bidirectionality of the conversion, its phase-preserving nature and the absorption of injected signal power during conversion—provide firm evidence that state transfer is occurring, and that we have accessed the beam-splitter Hamiltonian that is fundamentally capable of noiseless state transfer.

We have observed three key features in the conversion: its bidirectionality, its phase-preserving nature, and the absorption of injected signal power. These features provide firm evidence that state transfer is occurring. They also demonstrate that we have accessed the beam-splitter Hamiltonian, which is fundamentally capable of noiseless state transfer.

One sentence, one message

- **statement** — We have found A. / A was used to do …
- **cause–effect** — Because we have A, there is B. / From A follows B.
  A leads to B. / A enables B. / A is caused by B. / A results from B.
- **contrast** — Whereas A is …, B is … / Although A is …, B is … /
  In contrast to A, B is …
- **comparison** — Whereas A is …, B is …. / Compared to A, B is …

Reinschrift
SCIENCE COMMUNICATION

Connecting sentences

Proteins are essential to life, and understanding their structure can facilitate a mechanistic understanding of their function. Through an enormous experimental effort[1,2,3,4], the structures of around 100,000 unique proteins have been determined[5], but this represents a small fraction of the billions of known protein sequences[6,7]. Structural coverage is bottlenecked by the months to years of painstaking effort required to determine a single protein structure. Accurate computational approaches are needed to address this gap and to enable large-scale structural bioinformatics. Predicting the three-dimensional structure that a protein will adopt based solely on its amino acid sequence—the structure prediction component of the 'protein folding problem'[8]—has been an important open research problem for more than 50 years[9]. Despite recent progress[10,11,12,13,14], existing methods fall far short of atomic accuracy, especially when no homologous structure is available. Here we provide the first computational method that can regularly predict protein structures with atomic accuracy even in cases in which no similar structure is known. We validated an entirely redesigned version of our neural network-based model, AlphaFold, in the challenging 14th Critical Assessment of protein Structure Prediction (CASP14)[15], demonstrating accuracy competitive with experimental structures in a majority of cases and greatly outperforming other methods. Underpinning the latest version of AlphaFold is a novel machine learning approach that incorporates physical and biological knowledge about protein structure, leveraging multi-sequence alignments, into the design of the deep learning algorithm.

Nature **596**, 583–589 (2021).

Thematic progression

$$T_1 \rightarrow R_1$$
$$\downarrow$$
$$T_2 \rightarrow R_2$$

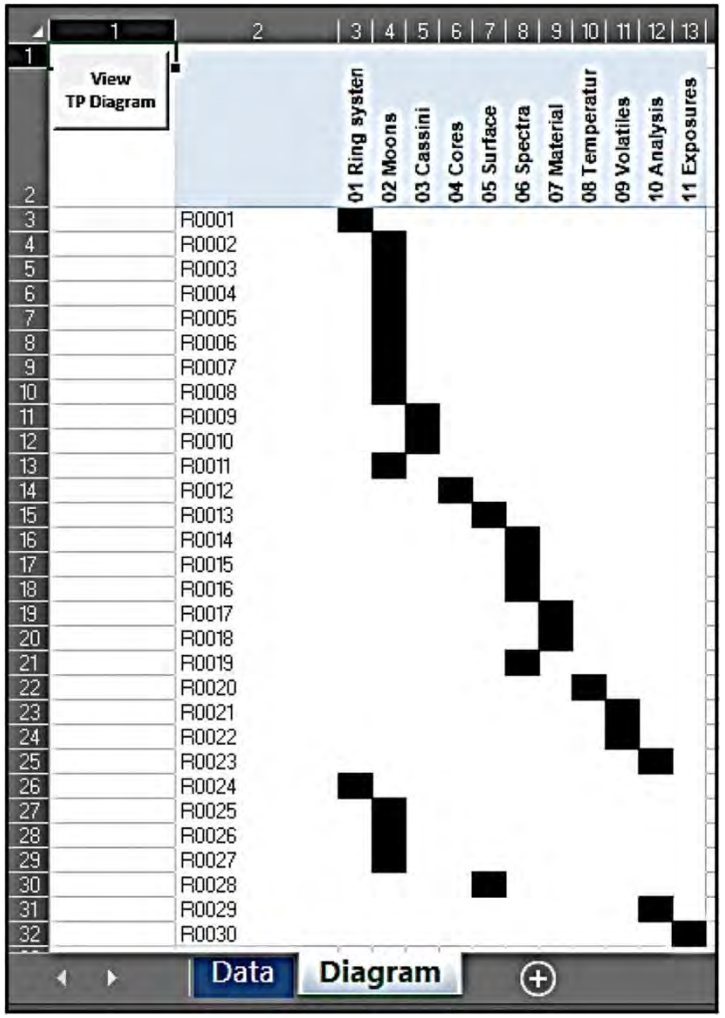

Reinschrift
SCIENCE COMMUNICATION

Leong, P.A. Visualizing texts: a tool for generating thematic-progression diagrams. *Functional Linguist.* **6**, 4 (2019).

Using AI intelligently

- In the attached paper, do the paragraphs have a good signposting structure? Please present a paragraph-by-paragraph analysis, citing the leading sentence for each paragraph, provide a short analysis and suggest an alternative sentence where needed.

- Now, let's focus on the abstract. Is there a clear thematic progression? Also here, please cite each sentence, provide a short analysis and suggest an alternative sentence where needed.

- How is the connection between sentences? Does sentence *n* pick up an element from sentence *n*-1 and pass another on to sentence n+1?

- Next, can you please familiarize yourself with the paper "Visualizing texts: a tool for generating thematic-progression diagrams" (Functional Linguistics 6, 4; 2019) and then provide a graphical representation of the theme–rheme structure of our abstract?

Connecting sentences

In practice

→ Reading aloud

Key points

- Compact sentences: One sentence, one message
- Deliver your message clearly.
- Connect the sentences in a coherent manner.

Reinschrift
SCIENCE COMMUNICATION

Constructing your text

- Manuscript level
- Section level
- Paragraph level
- Sentence level
- Word level

Flow

Creating flow

How to construct a coherent narrative

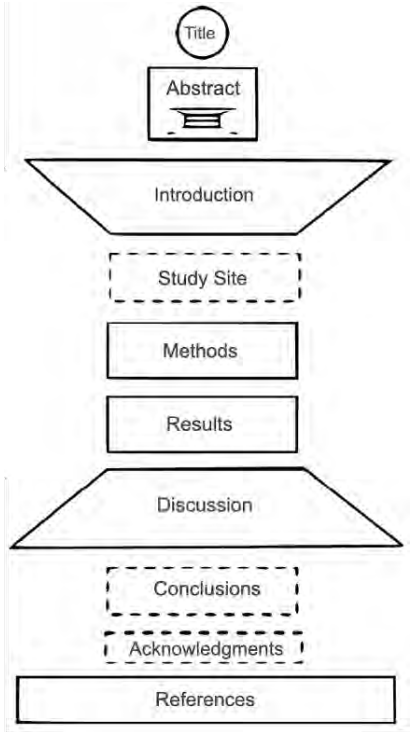

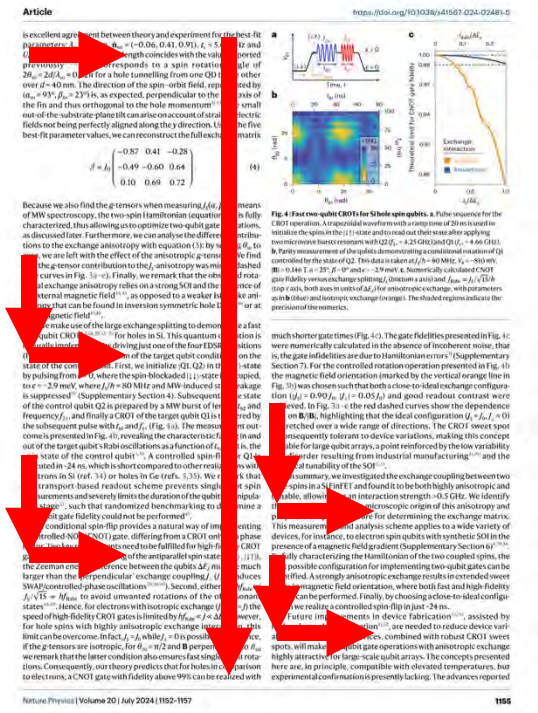

Left – From: Wu, J. *Landscape Ecology* **26**, 1345 (2011).
Right – *Nature Physics* **20**, 1152 (2024).

Reinschrift
SCIENCE COMMUNICATION

Constructing your text

- Manuscript level
- Section level
- Paragraph level
- Sentence level
- Word level

Structure

Flow

Style

Using AI intelligently

- Can we now look again at the full paper and analyse language and style? Specifically, please focus on
  1) finding typos and grammatical errors
  2) identifying awkward expressions and non-native-speaker giveaways
  3) flow and coherence of style and expressions

  Please structure your output along these criteria, with first the original sentence, then your suggestions (new text highlighted with bold face) and finally a short explanation.

- Can you please double-check?

Using AI as a ‘writing companion’: Key points

- The only thing that ultimately matters is content and clarity, not form
  - → Scientific writing is not about aesthetic virtuosity, but about clarity
  - → If AI can assist, then good, but the intellectual framework has to come from you
  - → You have to take responsibility for what is written
  - → Keep in mind that the text is ultimately ‘consumed’ by humans and not by AI

- Learn to interact meaningfully with AI technologies to truly benefit
  - → You need the necessary knowledge of your own (subject matter and language)

Reinschrift
SCIENCE COMMUNICATION

LLMs: A look under the bonnet

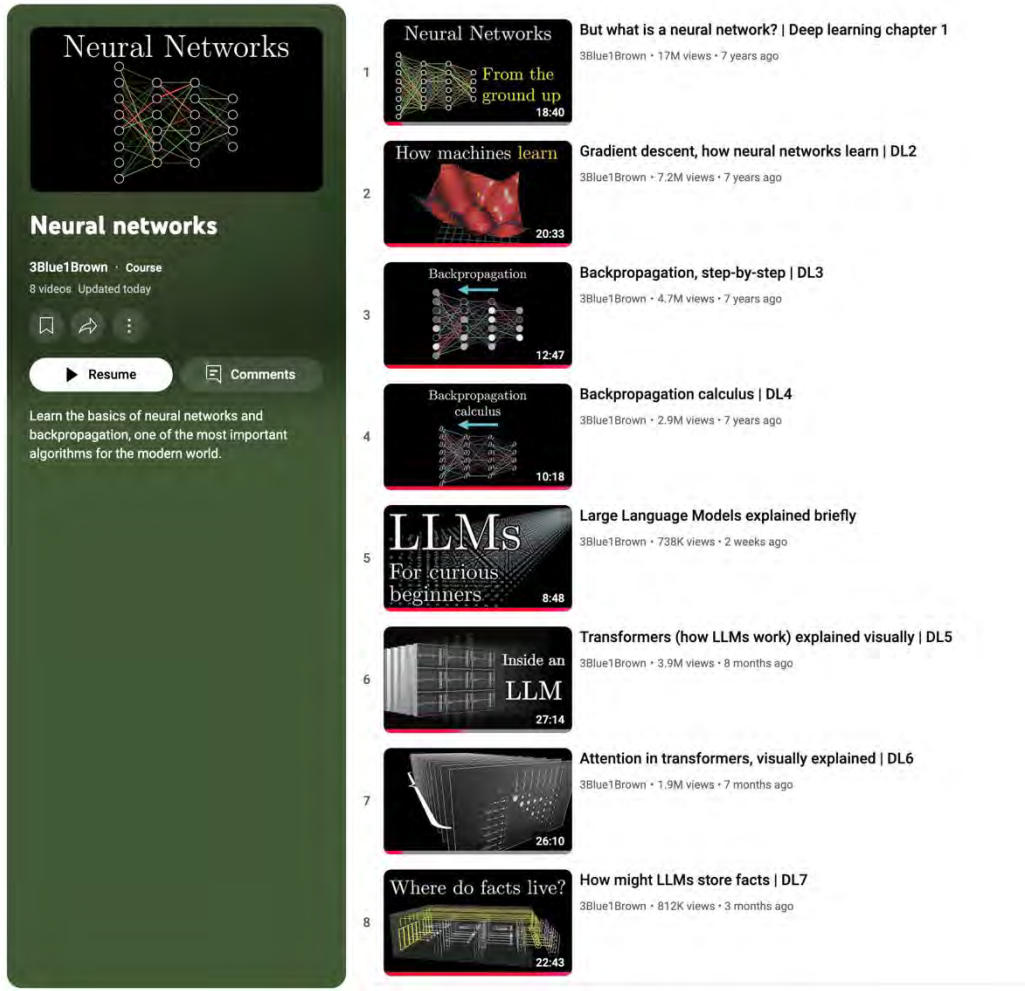

https://youtube.com/playlist?list=PLZHQObOWTQDNU6R1_67000Dx_ZCJB-3pi&si=xZYrAgs5OCm0I1PL

Reinschrift
SCIENCE COMMUNICATION

Use your own intelligence

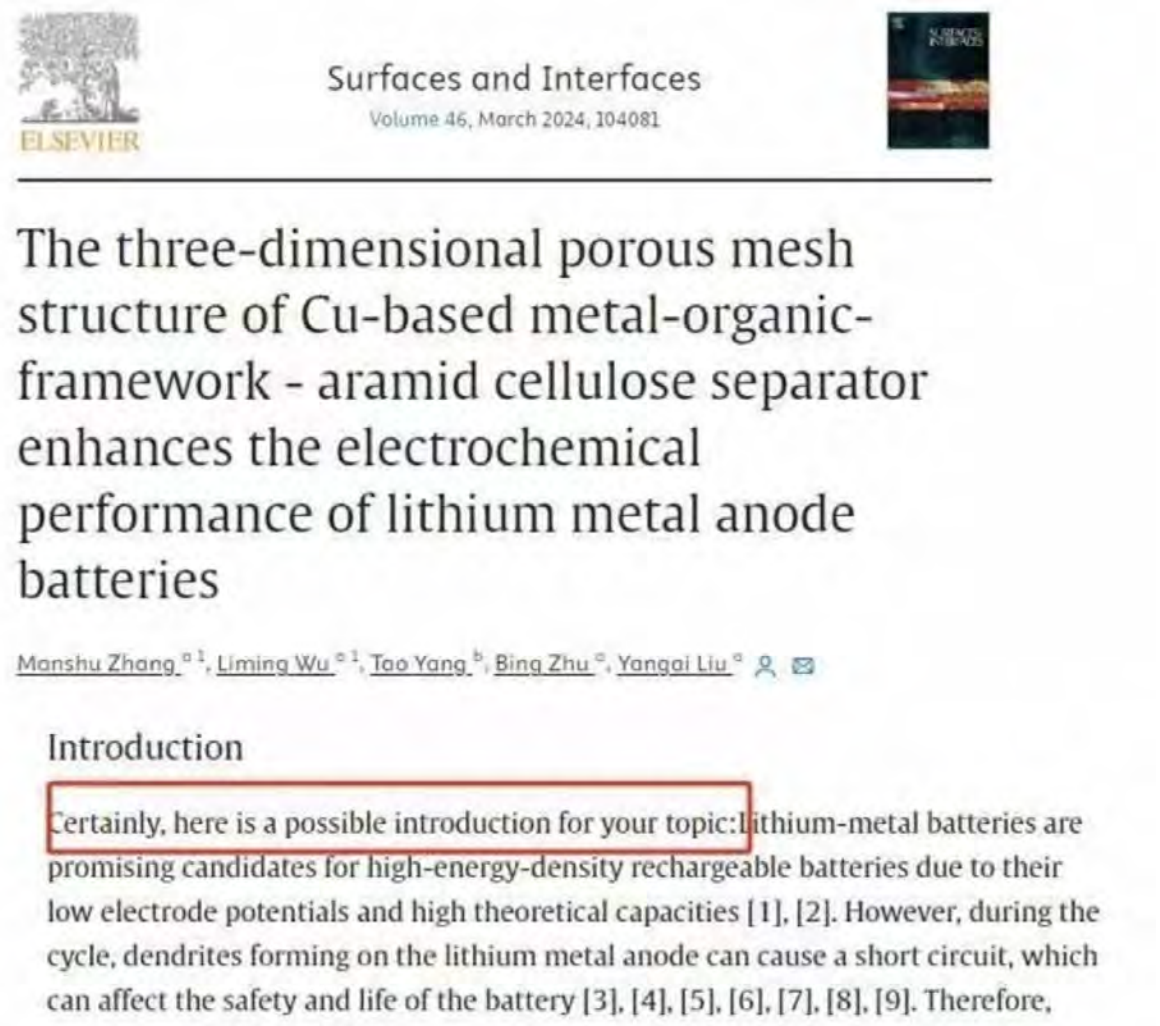

Surfaces and Interfaces
Volume 46, March 2024, 104081

The three-dimensional porous mesh structure of Cu-based metal-organic-framework - aramid cellulose separator enhances the electrochemical performance of lithium metal anode batteries

Manshu Zhang [a 1], Liming Wu [a 1], Tao Yang [b], Bing Zhu [a], Yangai Liu [a]

Introduction

Certainly, here is a possible introduction for your topic:Lithium-metal batteries are promising candidates for high-energy-density rechargeable batteries due to their low electrode potentials and high theoretical capacities [1], [2]. However, during the cycle, dendrites forming on the lithium metal anode can cause a short circuit, which can affect the safety and life of the battery [3], [4], [5], [6], [7], [8], [9]. Therefore,

Use your own intelligence

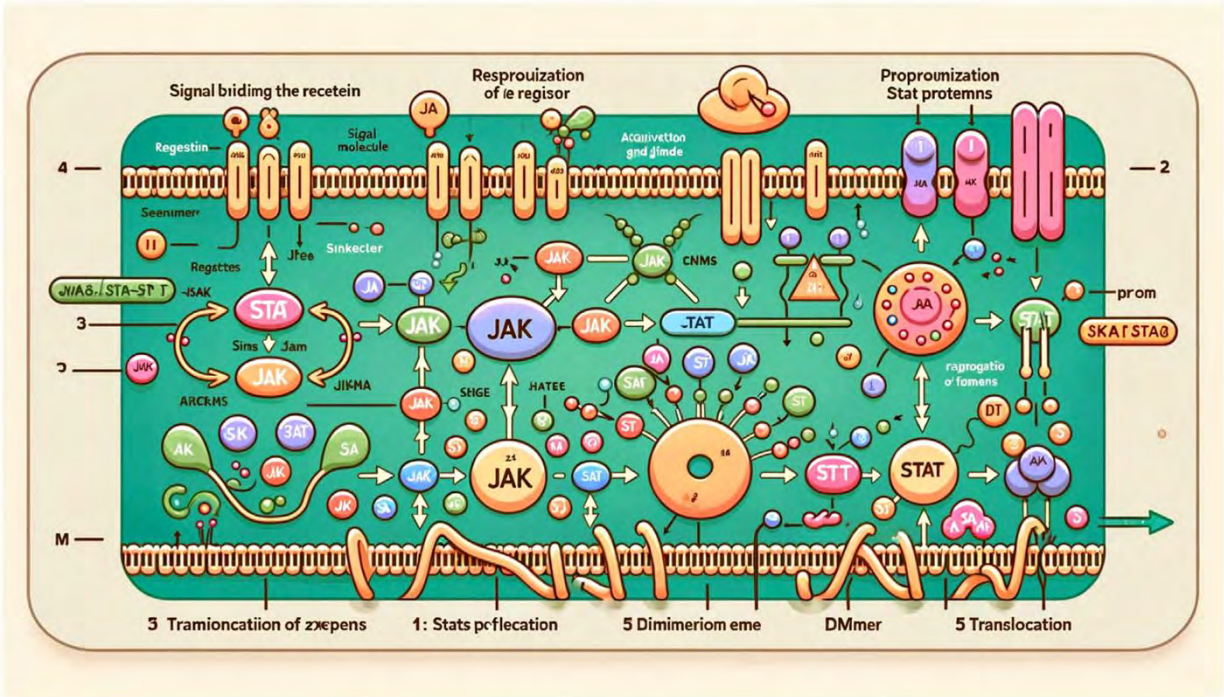

FIGURE 2
Diagram of the JAK-STAT signaling pathway: 1) Signal molecule binding to the receptor, 2) Activation of JAK kinase and phosphorylation of the receptor, 3) Recruitment and phosphorylation of STAT proteins by JAK, 4) Dimerization of STAT proteins, 5) Translocation of STAT dimers into the nucleus and initiation of gene transcription.

Guo, X., Dong, L. and Hao, D., *Front. Cell Dev. Biol.* 11:1339390 (2024).

Gell-Mann amnesia effect

"Briefly stated, the Gell-Mann Amnesia effect is as follows. You open the newspaper to an article on some subject you know well. In Murray's case, physics. In mine, show business. You read the article and see the journalist has absolutely no understanding of either the facts or the issues. Often, the article is so wrong it actually presents the story backward—reversing cause and effect. I call these the "wet streets cause rain" stories. Paper's full of them.

In any case, you read with exasperation or amusement the multiple errors in a story, and then turn the page to national or international affairs, and read as if the rest of the newspaper was somehow more accurate about Palestine than the baloney you just read. You turn the page, and forget what you know."

– *Michael Crichton (1942-2008)*

Reinschrift
SCIENCE COMMUNICATION

https://www.epsilontheory.com/gell-mann-amnesia/
https://en.wikipedia.org/wiki/Michael_Crichton#Gell-Mann_amnesia_effect

Putting AI to good use

- Anatomy of a scientific paper
- Storytelling, the abstract
- Language and style
- Text construction
- The editorial process
- References

Academic writing in the dawning age of AI

My currently favoured tools

- DeepL (Write)
- Claude
- (Microsoft Copilot)
- Undermind