Science communication for AI researchers - a short introduction

NeurIPS 2023

Monday 11 December

12:45 – 13:45 Talk

14:00 – 16:00 Drop-in

AIhub

Science communication for AI researchers

Professor Tom Dietterich
Oregon State University

Dr Lucy Smith
AIhub.org

- AIhub is a non-profit (UK charity) dedicated to connecting the AI community to the public by providing free, high-quality information.
- We are supported by many leading AI organisations.

Association for the Advancement of Artificial Intelligence

acm SIGAI

ICML International Conference On Machine Learning

NEURAL INFORMATION PROCESSING SYSTEMS

CLAIRE

EurAi

Ai COMMUNICATIONS

RoboCup

Artificial Intelligence

AIhub

news articles opinions education

https://aihub.org

aihuborg@gmail.com

@aihuborg

aihuborg

aihuborg

@Aihuborg

What we'll cover

- Why science communication matters
- Different ways to do science communication
- Working with the media
- Communicating via social media
- Writing a blog post
- Tips on explaining complex concepts
- How to find and use suitable images
- How to avoid AI hype
- Unconventional ways to do science communication

AIhub

Aims

- By the end of the session, you should be ready to plan and write a blog post.

AIhub

MACHINE LEARNING Science and Technology

PAPER

OPEN ACCESS

On the capacity and superposition of minima in neural network loss function landscapes

Maximilian P Niroomand*, John W R Morgan, Conor T Cafolla and David J Wales*

Department of Chemistry, University of Cambridge, Cambridge, United Kingdom

* Authors to whom any correspondence should be addressed.

Keywords: ensemble learning, interpretability, loss function landscape, theoretical chemistry

Abstract

Minima of the loss function landscape (LFL) of a neural network are locally optimal sets of weights that extract and process information from the input data to make outcome predictions. In underparameterised networks, the capacity of the weights may be insufficient to fit all the relevant information. We demonstrate that different local minima specialise in certain aspects of the learning problem, and process the input information differently. This effect can be exploited using a meta-network in which the predictive power from multiple minima of the LFL is combined to produce a better classifier. With this approach, we can increase the area under the receiver operating characteristic curve by around 20% for a complex learning problem. We propose a theoretical basis for combining minima and show how a meta-network can be trained to select the representative that is used for classification of a specific data item. Finally, we present an analysis of symmetry-equivalent solutions to machine learning problems, which provides a systematic means to improve the efficiency of this approach.

1. Introduction

Deep learning with neural networks is a high-dimensional, non-convex optimisation problem for a loss function landscape (LFL). The coordinates of a *minimum* in the LFL are a *set of weights* for the machine learning model and a locally optimal *solution* to the learning problem, and these terms will therefore be used interchangeably throughout. It follows that the coordinates of the global minimum of the LFL are the weights that produce the lowest possible value of the loss function for the training data. The aim of machine learning is usually for the model to find a set of weights that fit the training data, but also generalise well to unseen testing data. Our approach extends this view. Instead of looking at just one minimum of the LFL, we are interested in the expressive power of multiple minima. To analyse how different minima extract and process information from the input data, we survey numerous low-lying minima of the LFL. Here, we employ tools from the energy landscape approach (Wales 2003) to gain new insight into machine learning LFLs (Ballard *et al* 2017). We note that the role of local minima is somewhat different in ML landscapes compared to molecular systems. While in a molecular energy landscape only minima provide valid configurations for a stable molecule, this restriction does not apply to LFLs for machine learning. In fact, some low-lying non-minima will have a smaller loss value and higher classification accuracy than a high-lying minimum. Here, we are interested in developing a better understanding of the capacity of diverse minima of the LFL, and we show that by combining the expressive power of different minima, we can build a better classifier. The compact form of this predictor provides a balance between accuracy and efficiency, which will be useful in applications where evaluation is a computational bottleneck.

1.1. Background

Machine learning models are structurally limited in the amount of data they can fit: their capacity is finite. The most commonly known measure of capacity is perhaps the Vapnik-Chervonenkis (VC) dimension (Vapnik and Chervonenkis 1971, Vapnik *et al* 1994). The higher the VC dimension, the more complex are the

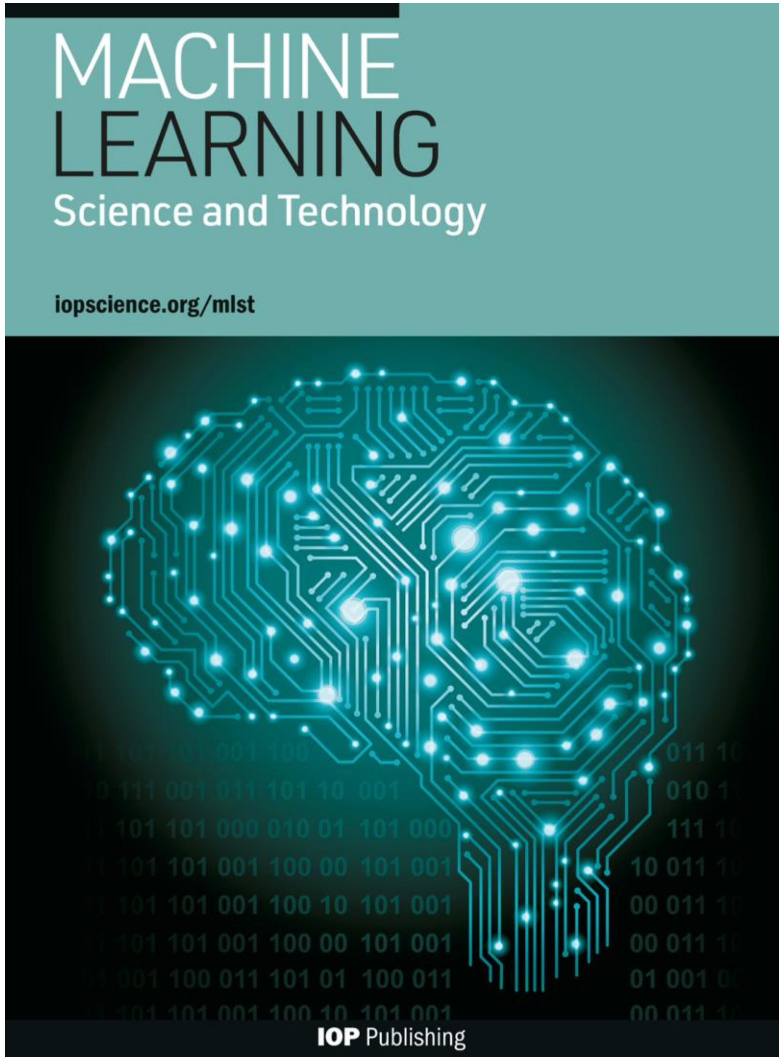

MACHINE
LEARNING
Science and Technology

OPEN ACCESS

PAPER

On the capacity and superposition of minima in neural network loss function landscapes

Maximilian P Niroomand*, John W R Morgan, Conor T Cafolla and David J Wales*
Department of Chemistry, University of Cambridge, Cambridge, United Kingdom
* Authors to whom any correspondence should be addressed.

Keywords: ensemble learning, interpretability, loss function landscape, theoretical chemistry

Abstract
Minima of the loss function landscape (LFL) of a neural network are locally optimal sets of weights that extract and process information from the input data to make outcome predictions. In underparameterised networks, the capacity of the weights may be insufficient to fit all the relevant information. We demonstrate that different local minima specialise in certain aspects of the learning problem, and process the input information differently. T
a meta-network in which the predictive power from multiple mi
produce a better classifier. With this approach, we can increase t
operating characteristic curve by around 20% for a complex lear
theoretical basis for combining minima and show how a meta-n
representative that is used for classification of a specific data item
symmetry-equivalent solutions to machine learning problems, w
to improve the efficiency of this approach.

1. Introduction

Deep learning with neural networks is a high-dimensional, non-conve
function landscape (LFL). The coordinates of a minimum in the LFL ar
learning model and a locally optimal solution to the learning problem, an
interchangeably throughout. It follows that the coordinates of the global minimum of the LFL are the weights that produce the lowest possible value of the loss function for the training data. The aim of machine learning is usually for the model to find a set of weights that fit the training data, but also generalise well to unseen testing data. Our approach extends this view. Instead of looking at just one minimum of the LFL, we are interested in the expressive power of multiple minima. To analyse how different minima extract and process information from the input data, we survey numerous low-lying minima of the LFL. Here, we employ tools from the energy landscape approach (Wales 2003) to gain new insight into machine learning LFLs (Ballard *et al* 2017). We note that the role of local minima is somewhat different in ML landscapes compared to molecular systems. While in a molecular energy landscape only minima provide valid configurations for a stable molecule, this restriction does not apply to LFLs for machine learning. In fact, some low-lying non-minima will have a smaller loss value and higher classification accuracy than a high-lying minimum. Here, we are interested in developing a better understanding of the capacity of diverse minima of the LFL, and we show that by combining the expressive power of different minima, we can build a better classifier. The compact form of this predictor provides a balance between accuracy and efficiency, which will be useful in applications where evaluation is a computational bottleneck.

1.1. Background
Machine learning models are structurally limited in the amount of data they can fit: their capacity is finite. The most commonly known measure of capacity is perhaps the Vapnik-Chervonenkis (VC) dimension (Vapnik and Chervonenkis 1971, Vapnik *et al* 1994). The higher the VC dimension, the more complex are the

© 2022 The Author(s). Published by IOP Publishing Ltd

1000 views

Causal Confounds in Sequential Decision Making

AUTHORS
Gokul Swamy

AFFILIATIONS
RI, CMU

PUBLISHED
November 28, 2022

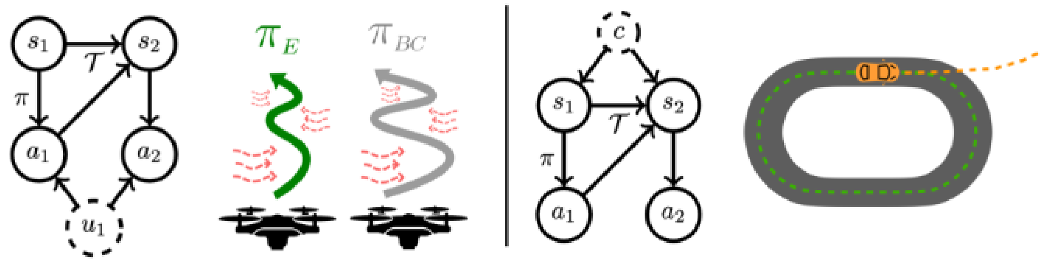

(a) (b) (c) (d)

A standard assumption in sequential decision making is that we observe everything required to make good decisions. In practice however, this isn't always the case. We discuss two specific examples (temporally correlated noise (a) and unobserved contexts (c)) that have stymied the use of IL/RL algorithms (in autonomous helicopters (b) and self-driving (d)). We derive provably correct algorithms for both of these problems that scale to continuous control problems.

Reinforcement Learning (RL) and Imitation Learning (IL) methods have achieved impressive results in recent years like beating the world champion at Go or controlling stratospheric balloons. Usually, these results are on problems where we either a) observe the full state or b) are able to faithfully execute our intended actions on the system. However, we frequently have to contend with situations where this isn't the case: our self-driving car might miss a person's hand gestures or persistent wind might make it difficult to fly our quadcopter perfectly straight. These sorts of situations can cause standard IL approaches to perform poorly ([1], [2]). In causal inference, we call a random variable that we don't observe that influences a relationship we'd like

Causal Confounds in Sequential Decision Making

AUTHORS
Gokul Swamy

AFFILIATIONS
RI, CMU

PUBLISHED
November 28, 2022

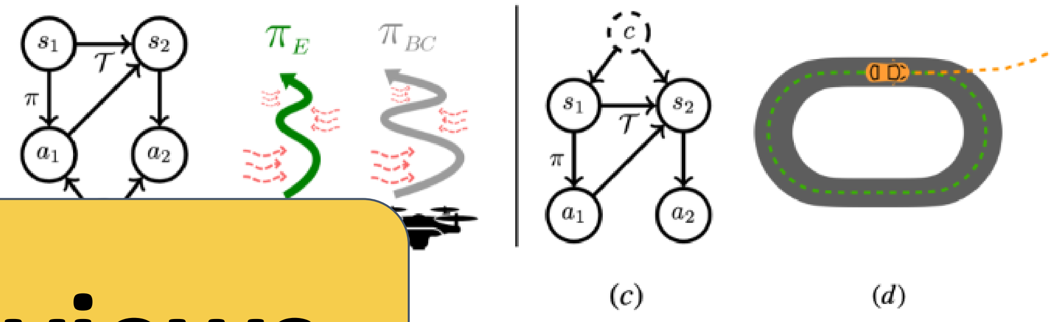

8000 views

making is that we observe everything required to make good decisions. In . We discuss two specific examples (temporally correlated noise (a) and unobserved contexts (c)) that have stymied the use of IL/RL algorithms (in autonomous helicopters (b) and self-driving (d)). We derive provably correct algorithms for both of these problems that scale to continuous control problems.

Reinforcement Learning (RL) and Imitation Learning (IL) methods have achieved impressive results in recent years like beating the world champion at Go or controlling stratospheric balloons. Usually, these results are on problems where we either a) observe the full state or b) are able to faithfully execute our intended actions on the system. However, we frequently have to contend with situations where this isn't the case: our self-driving car might miss a person's hand gestures or persistent wind might make it difficult to fly our quadcopter perfectly straight. These sorts of situations can cause standard IL approaches to perform poorly ([1], [2]). In causal inference, we call a random variable that we don't observe that influences a relationship we'd like

AIhub

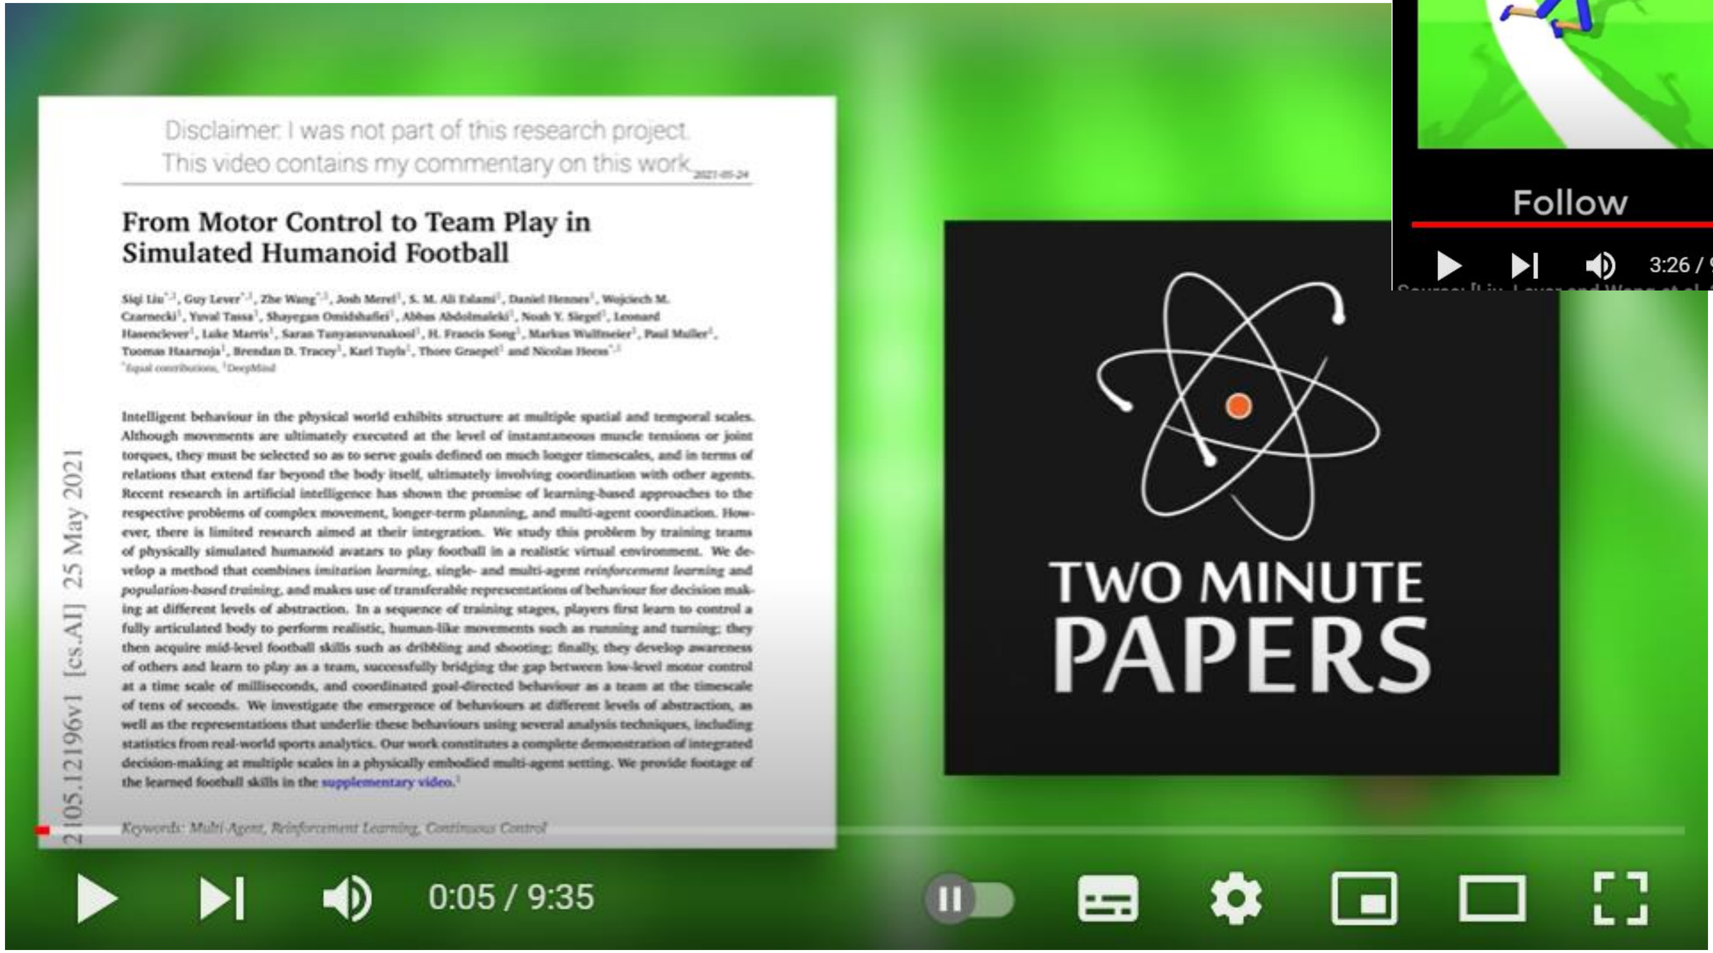

https://youtu.be/HTON7odbW0o

AIhub

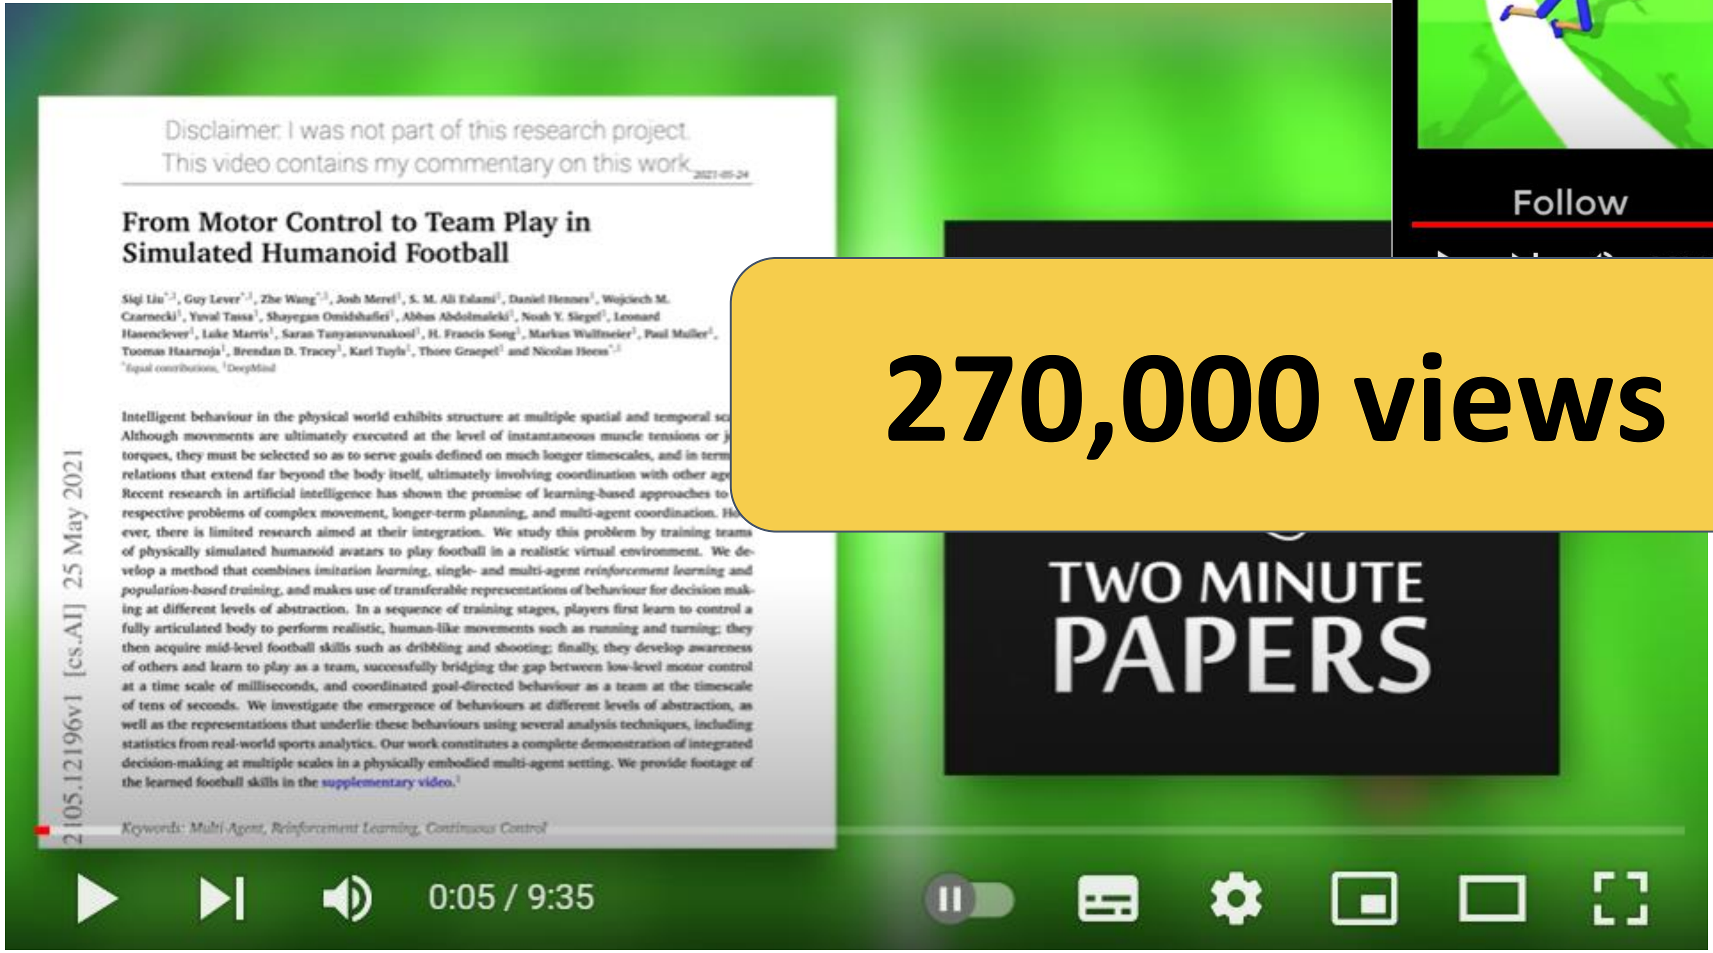

270,000 views

https://youtu.be/HTON7odbW0o

AIhub

The Guardian
News website of the year

Artificial intelligence (AI)

ChatGPT maker OpenAI releases 'not fully reliable' tool to detect AI generated content

OpenAI is calling on educators to give their feedback on how the tool is used, amid rising concerns around AI-assisted cheating at universities

Josh Taylor
@joshgnosis
Wed 1 Feb 2023 03.58 GMT

ChatGPT creator, OpenAI, has released a tool to detect AI generated content Photograph: Lionel Bonaventure/AFP/Getty Images

OpenAI, the research laboratory behind AI program ChatGPT, has released a tool designed to detect whether text has been written by artificial intelligence, but warns it's not completely reliable – yet.

In a blog post on Tuesday, OpenAI linked to a new classifier tool that has been trained to distinguish between text written by a human and that written by a variety of AI, not just ChatGPT.

Open AI researchers said that while it was "impossible to reliably detect all AI-written text", good classifiers could pick up signs that text was written by AI. The tool could be useful in cases where AI was used for "academic

AIhub

The Guardian
News website of the year

Artificial intelligence (AI)

ChatGPT maker OpenAI releases 'not fully reliable' tool to detect AI generated content

OpenAI is calling on educators to give their feedback on how the tool is used, amid rising concerns around AI-assisted cheating at universities

Josh Taylor
@joshgnosis
Wed 1 Feb 2023 03.58 GMT

t Photograph: Lionel

OpenAI, the research laboratory behind AI program ChatGPT, has released a tool designed to detect whether text has been written by artificial intelligence, but warns it's not completely reliable - yet.

In a blog post on Tuesday, OpenAI linked to a new classifier tool that has been trained to distinguish between text written by a human and that written by a variety of AI, not just ChatGPT.

Open AI researchers said that while it was "impossible to reliably detect all AI-written text", good classifiers could pick up signs that text was written by AI. The tool could be useful in cases where AI was used for "academic

> 1 million views?
Circulation of 9 million

AIhub

Why science communication matters

- Inspiring the next generation
- Science for society
- Transparency
- De-hyping science
- Adding value to the research

Why science communication matters
(https://youtu.be/hHFltr_j4FI)

Have you done any science communication before?

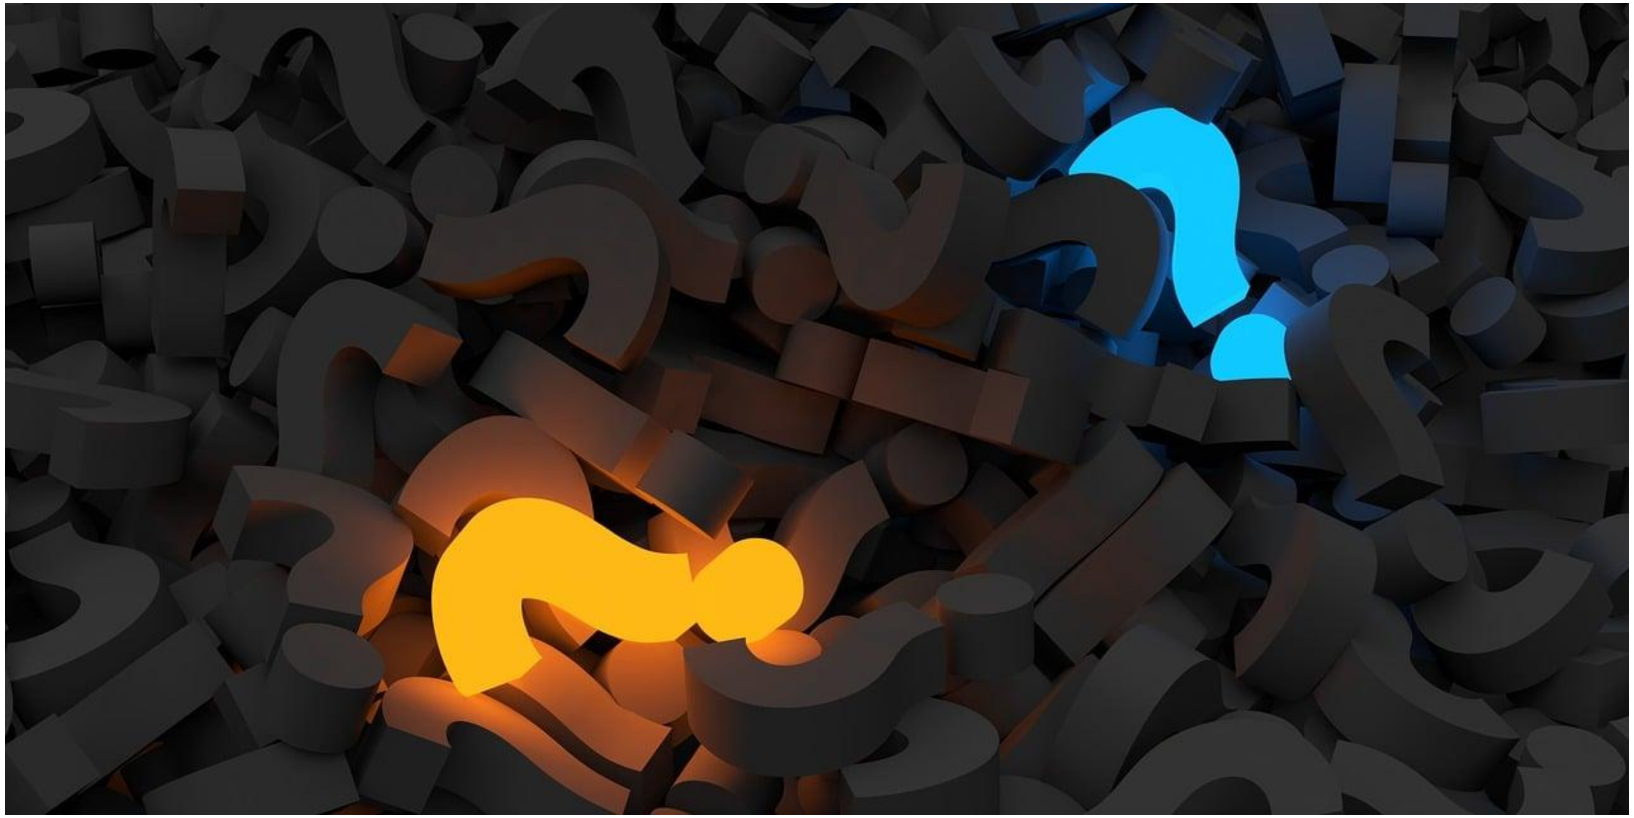

Different ways to do science communication

- TV, public talks, radio
- Blog posts
- Collaboration with artists
- Workshops
- Competitions
- Social media
- Podcasts
- Exhibitions

Different ways to do science communication (https://youtu.be/Jb8eRfItOLE)

AIhub

Working with the media

- Your press office
- Science journalists

How to approach the media: An interview with Evan Ackerman
(*https://youtu.be/5kslhRzoDRw*)

Working with the media - some tips

- Pitching
  - Tell a story - broader implications of your work (don't just state results)
  - A video / image can be helpful in "selling" the story
- Types of questions science journalists may ask
  - Background to the research - where did the idea come from?
  - Context - state of the field, previous work, challenges, future plans
- Tips for answering
  - Make answers accessible
  - Avoid technical specifics
  - Connect your research to real-world issues or applications
- Use your University Press Office

AIhub

The importance of owning your sci-comm

- When someone else reports on your work you lose control over the content.

Robot Programmed To Fall In Love With a Girl Goes Too Far
TECHNOLOGY NOVEMBER 27, 2013

News / Technology
Toshiba unveils the creepy robot that could one day steal your job
By Mary Jordan | 3:20pm Jan 8, 2015

DOOM AND GLOOM Creepy AI predicts what the Apocalypse will look like after scientists reset Doomsday Clock for 2023

Miracle robot will revolutionise brain surgery for epilepsy sufferers

'KILLER' BOTS Rogue superhuman AI 'could kill everyone' and wipe out human race... the tech should be controlled like nukes

INNOVATION
This "Psychic Robot" Can Read Your Mind

A starting point to communicating directly: social media

A starting point to communicating directly: social media

- Ways to use social media for your research:
  - Passive
  - Active

How using social media can benefit your research - passive

- Follow other researchers in the field.
  - Who do they follow?
  - Follow their followers.
  - Build your network.
- Find out about events / workshops / other interesting content.
- Find out about grants / positions / opportunities.
- Follow journalists.

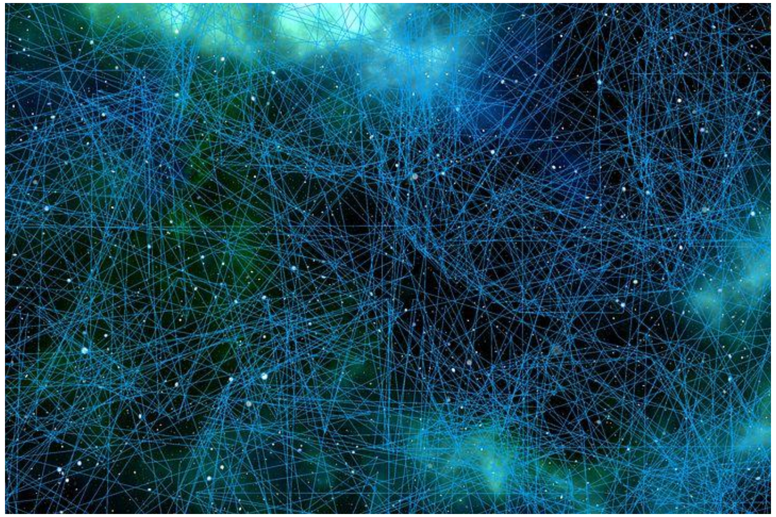

How using social media can benefit your research - active

- Use to promote your research.
    - Can be a great tool for refining your message.
    - How would you compress your research into a tweet, or thread?
- Engage in constructive discussions.
- Build connections with other researchers, journalists, organisations.
- Feel part of a community.
- Amplify the voices of others.

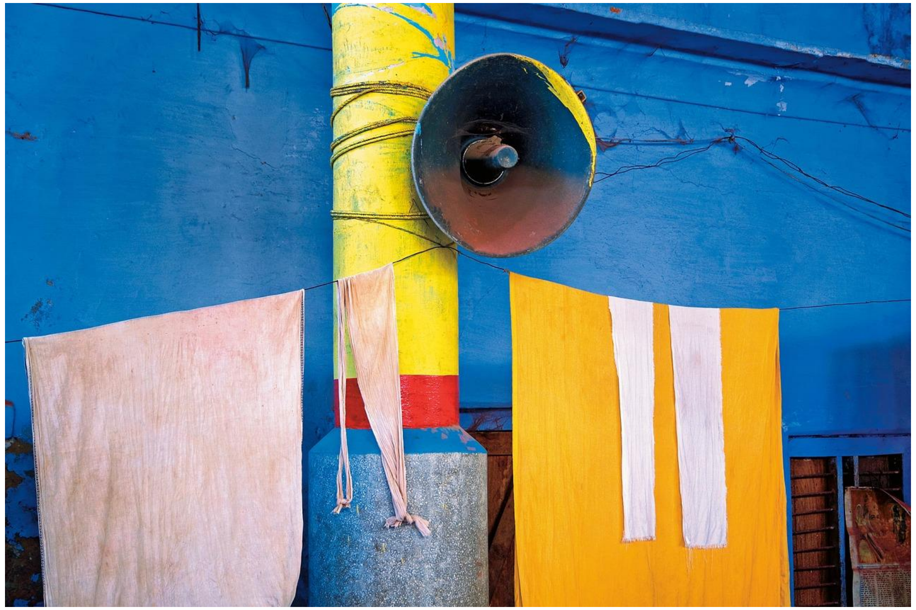

Caveats

- Can be easy to get sucked into controversies and arguments.
- Short-form of tweets (for example) often not conducive to in-depth discussions.

AIhub

Finding your story

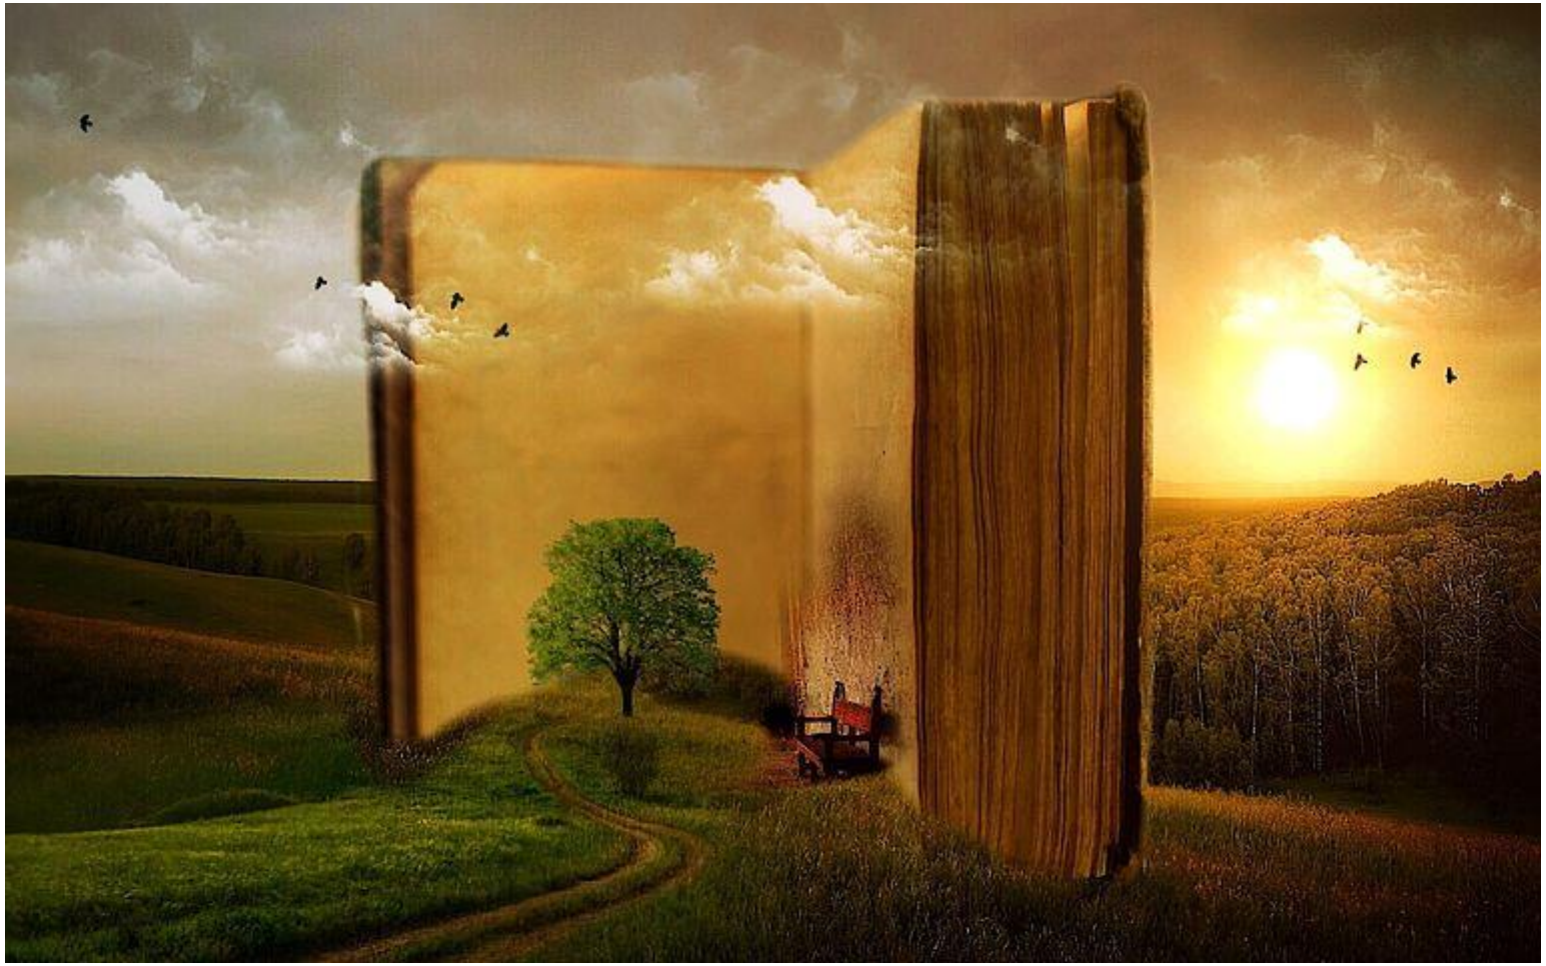

AIhub

Your story as tweets

The questions:

- What problem are you trying to solve? *
- Why is it important?
- How does this relate to people's lives?
- What is the current state of the field?
- What's the contribution of your research? *
- What are the implications of your findings?
- What challenges did you face?
- What are the limitations of your contribution?
- What are you planning next?

(* minimum starting point for communication on a social media platform)

AIhub

Example from a ML research paper

IOP Publishing   *Mach. Learn.: Sci. Technol.* **3** (2022) 045034   https://doi.org/10.1088/2632-2153/aca23d

MACHINE LEARNING Science and Technology

CrossMark

OPEN ACCESS

RECEIVED
1 July 2022

REVISED
17 October 2022

ACCEPTED FOR PUBLICATION
11 November 2022

PUBLISHED
29 December 2022

PAPER

Self-supervised learning of materials concepts from crystal structures via deep neural networks

Yuta Suzuki[1,2,6], Tatsunori Taniai[3], Kotaro Saito[2,4], Yoshitaka Ushiku[3] and Kanta Ono[1,2,5,*]

1 The Graduate University for Advanced Studies (SOKENDAI), Ibaraki, Japan
2 Institute of Materials Structure Science (IMSS), High Energy Accelerator Research Organization (KEK), Ibaraki, Japan
3 OMRON SINIC X Corporation, Tokyo, Japan
4 Randeft, Inc., Tokyo, Japan
5 Department of Applied Physics, Osaka University, Osaka, Japan
6 Current affiliation: Advanced R&D and Engineering Company, TOYOTA MOTOR CORPORATION, Shizuoka, Japan.
* Author to whom any correspondence should be addressed.

E-mail: ono@ap.eng.osaka-u.ac.jp

Keywords: materials informatics, deep metric learning, crystal structure, self-supervised learning

Supplementary material for this article is available online

a

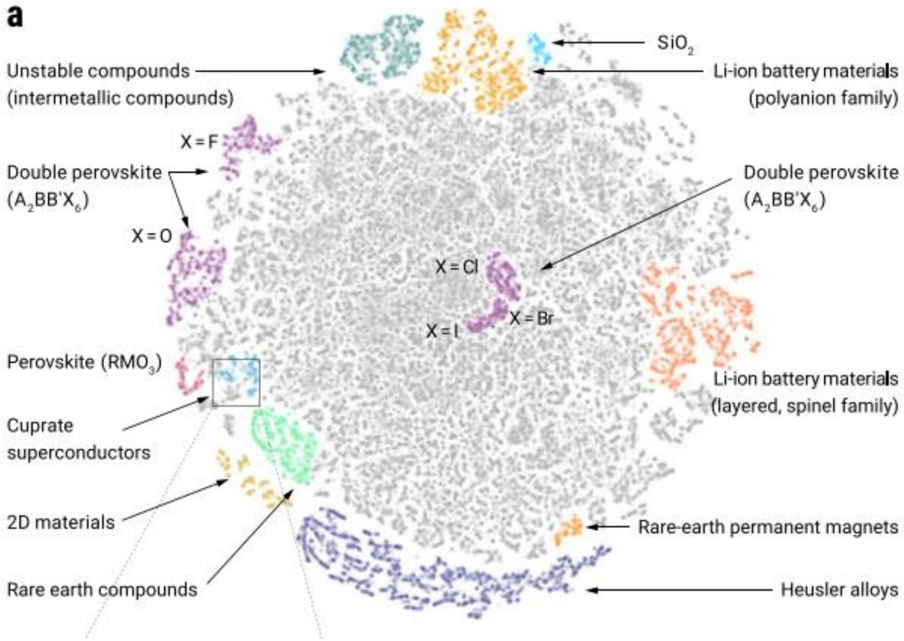

AIhub

Your story as tweets

The questions:

- What problem are you trying to solve?
- Why is it important?
- How does this relate to people's lives?
- What is the current state of the field?
- What's the contribution of your research?
- What are the implications of your findings?
- What challenges did you face?
- What are the limitations of your contribution?
- What are you planning next?

Materials discovery is a slow process that involves searching through a vast space of potential structures. Key to accelerating this process is understanding how the structure of a material affects its function. Suzuki *et al* have used ML to better understand, and map, this relationship.

AIhub

Your story as tweets

The questions:

- What problem are you trying to solve?
- Why is it important?
- How does this relate to people's lives?
- What is the current state of the field?
- What's the contribution of your research?
- What are the implications of your findings?
- What challenges did you face?
- What are the limitations of your contribution?
- What are you planning next?

The discovery of new materials is essential to making progress in many of the technological challenges we face, such as the development of more efficient solar cells or batteries, and clean water production.

AIhub

Turning your tweets into a blog post

What makes a good story?

- Pitched at the right level for the audience.

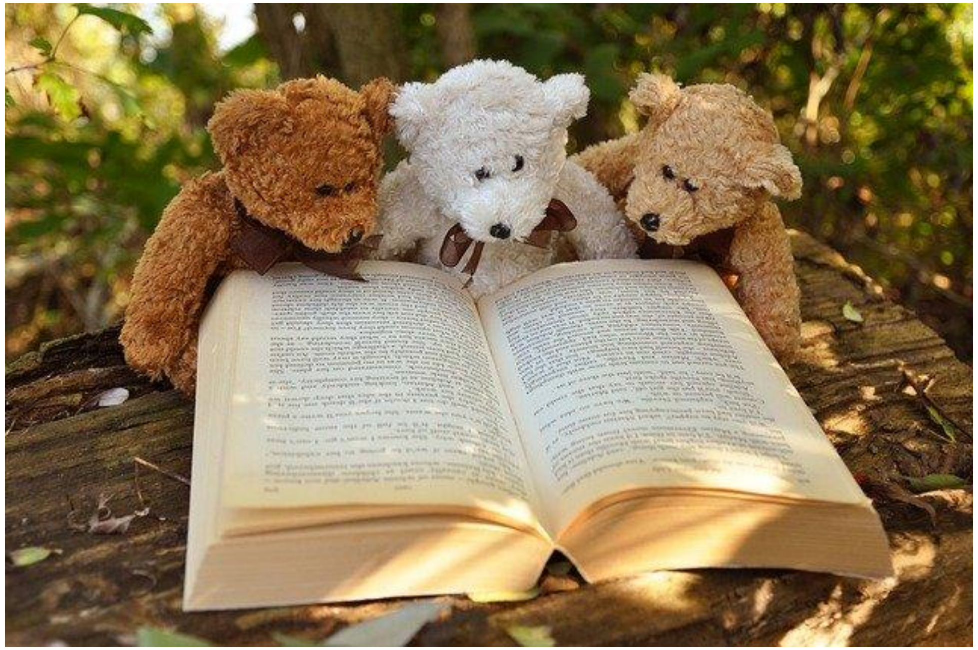

What makes a good story?

- Pitched at the right level for the audience.
- Connects with the audience:
    - Contains a link to application(s) from the real world.
    - Touches on a lived experience / passion / problem.

What makes a good story?

- Pitched at the right level for the audience.
- Connects with the audience:
  - Contains a link to application(s) from the real world.
  - Touches on a lived experience / passion / problem.
- Contains a true representation of the work and its implications, avoiding hype and inflated expectations

What makes a good story?

- Pitched at the right level for the audience.
- Connects with the audience:
  - Contains a link to application(s) from the real world.
  - Touches on a lived experience / passion / problem.
- Contains a true representation of the work and its implications, avoiding hype and inflated expectations
- Has a structure and natural flow.

What makes a good story?

- Pitched at the right level for the audience.
- Connects with the audience:
  - Contains a link to application(s) from the real world.
  - Touches on a lived experience / passion / problem.
- Contains a true representation of the work and its implications, avoiding hype and inflated expectations
- Has a structure and natural flow.
- Poses a question/hypothesis at the beginning, which the author goes about answering throughout the post.

Tips on writing a blog post

- As mentioned, first establish who your audience is - this determines the level to pitch the post.

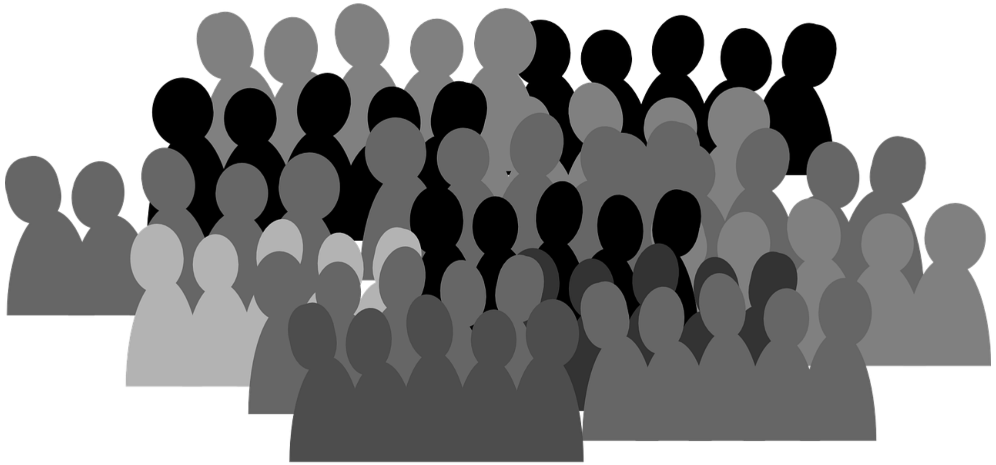

Tips on writing a blog post

- First establish who your audience is - this determines the level to pitch the post.
- Think about your key message - what do you want to convey to the audience?

Tips on writing a blog post

- First establish who your audience is - this determines the level to pitch the post.
- Think about your key message - what do you want to convey to the audience?
- Have your bullet point summaries to hand. Do these have a logical flow? Do you need to add others, or change the order?

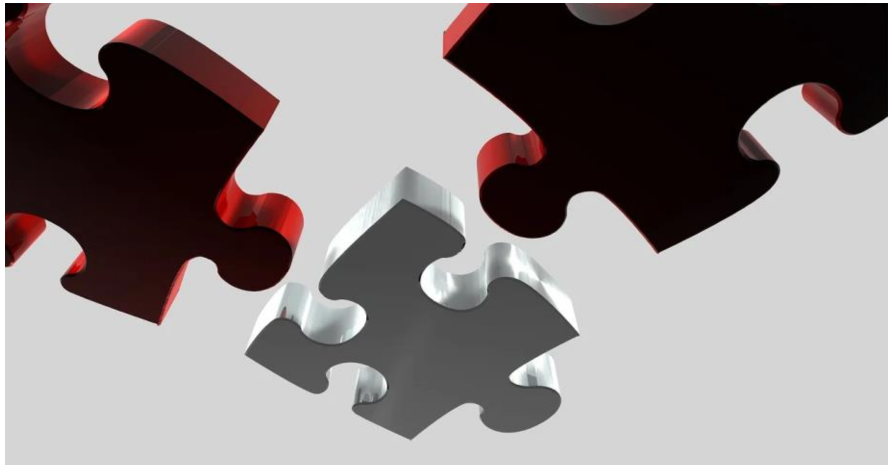

Tips on writing a blog post

- First establish who your audience is - this determines the level to pitch the post.
- Think about your key message - what do you want to convey to the audience?
- Have your bullet point summaries to hand. Do these have a logical flow? Do you need to add others, or change the order?
- Expand your tweets into paragraphs. Clarify, explain and give examples. More on this to follow…

AIhub

Tips on writing a blog post

- First establish who your audience is - this determines the level to pitch the post.
- Think about your key message - what do you want to convey to the audience?
- Have your bullet point summaries to hand. Do these have a logical flow? Do you need to add others, or change the order?
- Expand your tweets into paragraphs. Clarify, explain and give examples. More on this to follow…
- Add images/diagrams/videos to help explain key concepts.

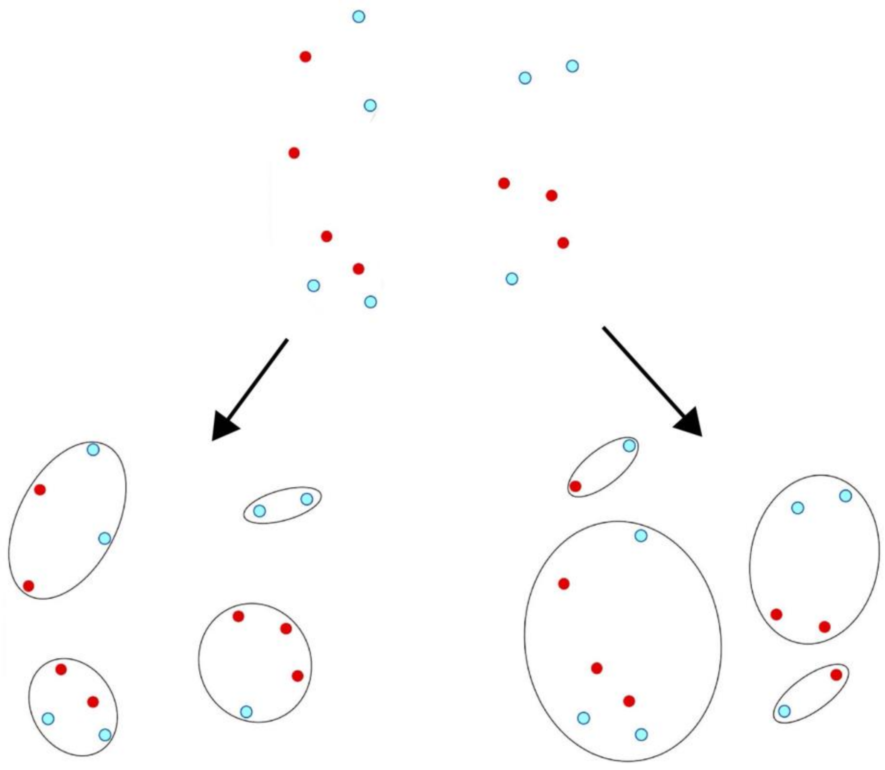

Credit: Max Springer

Tips on writing a blog post

- First establish who your audience is - this determines the level to pitch the post.
- Think about your key message - what do you want to convey to the audience?
- Have your bullet point summaries to hand. Do these have a logical flow? Do you need to add others, or change the order?
- Expand your tweets into paragraphs. Clarify, explain and give examples. More on this to follow…
- Add images/diagrams/videos to help explain key concepts.
- Read, re-read and seek feedback.

The next step: expand your tweets into paragraphs

AIhub

Our example research paper

IOP Publishing Mach. Learn.: Sci. Technol. 3 (2022) 045034 https://doi.org/10.1088/2632-2153/aca23d

MACHINE LEARNING
Science and Technology

CrossMark

OPEN ACCESS

RECEIVED
1 July 2022

REVISED
17 October 2022

ACCEPTED FOR PUBLICATION
11 November 2022

PUBLISHED
29 December 2022

PAPER

Self-supervised learning of materials concepts from crystal structures via deep neural networks

Yuta Suzuki[1,2,6], **Tatsunori Taniai**[3], **Kotaro Saito**[2,4], **Yoshitaka Ushiku**[3] **and Kanta Ono**[1,2,5,*]

1 The Graduate University for Advanced Studies (SOKENDAI), Ibaraki, Japan
2 Institute of Materials Structure Science (IMSS), High Energy Accelerator Research Organization (KEK), Ibaraki, Japan
3 OMRON SINIC X Corporation, Tokyo, Japan
4 Randeft, Inc., Tokyo, Japan
5 Department of Applied Physics, Osaka University, Osaka, Japan
6 Current affiliation: Advanced R&D and Engineering Company, TOYOTA MOTOR CORPORATION, Shizuoka, Japan.
* Author to whom any correspondence should be addressed.

E-mail: ono@ap.eng.osaka-u.ac.jp

Keywords: materials informatics, deep metric learning, crystal structure, self-supervised learning

Supplementary material for this article is available online

Turning your tweets into a blog post

Materials discovery is a slow process that involves searching through a vast space of potential structures. Key to accelerating this process is understanding how the structure of a material affects its function. Suzuki *et al* have used ML to better understand, and map, this relationship.

Turning your tweets into a blog post

Materials discovery is a slow process that involves searching through a vast space of potential structures. Key to accelerating this process is understanding how the structure of a material affects its function. Suzuki *et al* have used ML to better understand, and map, this relationship.

Imagine you are working on developing a new material for an efficient battery. Where do you start? How do you go about finding that material? What structure would give you the properties you are looking for?

AIhub

Turning your tweets into a blog post

Materials discovery is a slow process that involves searching through a vast space of potential structures. Key to accelerating this process is understanding how the structure of a material affects its function. Suzuki *et al* have used ML to better understand, and map, this relationship.

Imagine you are working on developing a new material for an efficient battery. Where do you start? How do you go about finding that material? What structure would give you the properties you are looking for? In the past, this would have involved a time-consuming experimental fabrication process, most likely informed by theoretical models. Given the sparsity of materials in a vast search space, the process of discovering and fabricating a new material could take many years.

Turning your tweets into a blog post

Materials discovery is a slow process that involves searching through a vast space of potential structures. Key to accelerating this process is understanding how the structure of a material affects its function. Suzuki *et al* have used ML to better understand, and map, this relationship.

Imagine you are working on developing a new material for an efficient battery. Where do you start? How do you go about finding that material? What structure would give you the properties you are looking for? In the past, this would have involved a time-consuming experimental fabrication process, most likely informed by theoretical models. Given the sparsity of materials in a vast search space, the process of discovering and fabricating a new material could take many years.

Now, imagine you could accelerate part of this process and narrow your search. The key to doing this is through an understanding the relationships between the structures of materials and their functional properties, as the diverse properties of materials are determined by their structures.

AIhub

Turning your tweets into a blog post

Materials discovery is a slow process that involves searching through a vast space of potential structures. Key to accelerating this process is understanding how the structure of a material affects its function. Suzuki *et al* have used ML to better understand, and map, this relationship.

Imagine you are working on developing a new material for an efficient battery. Where do you start? How do you go about finding that material? What structure would give you the properties you are looking for? In the past, this would have involved a time-consuming experimental fabrication process, most likely informed by theoretical models. Given the sparsity of materials in a vast search space, the process of discovering and fabricating a new material could take many years.

Now, imagine you could accelerate part of this process and narrow your search. The key to doing this is through an understanding the relationships between the structures of materials and their functional properties, as the diverse properties of materials are determined by their structures. In their research, Suzuki *et al* used machine learning (ML) techniques to create a map of the materials space and measure the similarity between materials.

AIhub

Simplifying complex concepts

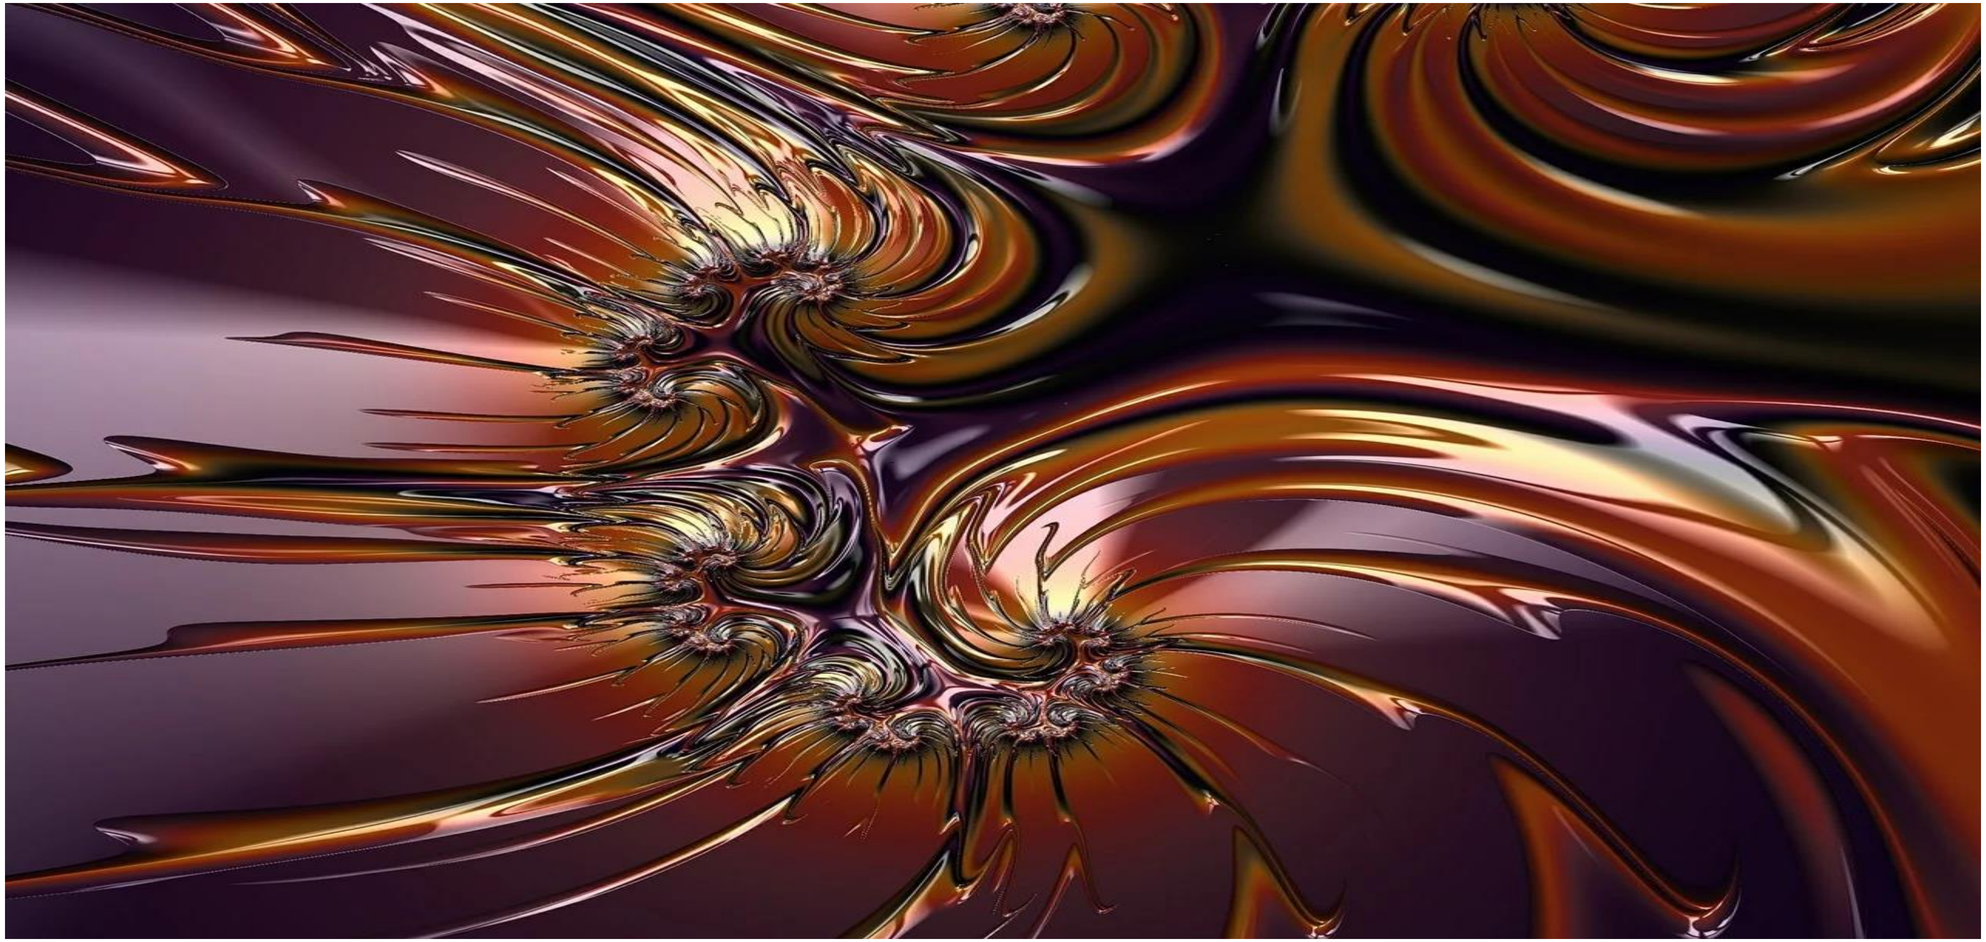

AIhub

Levels of complexity

CrossMark

OPEN ACCESS

RECEIVED
1 July 2022

PAPER

Self-supervised learning of materials concepts from crystal structures via deep neural networks

Yuta Suzuki[1,2,6], Tatsunori Taniai[3], Kotaro Saito[2,4], Yoshitaka Ushiku[3] and Kanta Ono[1,2,5,*]

Using a couple of sentences about their method and contribution as an example.

Level 1: suitable for a ML/physics audience.

- Suzuki *et al* have used a self-supervised deep learning approach to learn material embeddings from crystal structures of over 120 000 materials. This enabled them to capture relationships between the structure of a material and its properties.

AIhub

Levels of complexity

CrossMark

OPEN ACCESS

RECEIVED
1 July 2022

PAPER

Self-supervised learning of materials concepts from crystal structures via deep neural networks

Yuta Suzuki[1,2,6], Tatsunori Taniai[3], Kotaro Saito[2,4], Yoshitaka Ushiku[3] and Kanta Ono[1,2,5,*]

Using a sentence about their method and contribution as an example.

Level 1: suitable for a ML/physics audience.

- Suzuki *et al* have used a self-supervised deep learning approach to learn material embeddings from crystal structures of over 120 000 materials. This enabled them to capture relationships between the structure of a material and its properties.

AIhub

Levels of complexity

CrossMark

OPEN ACCESS

RECEIVED
1 July 2022

PAPER

Self-supervised learning of materials concepts from crystal structures via deep neural networks

Yuta Suzuki[1,2,6], Tatsunori Taniai[3], Kotaro Saito[2,4], Yoshitaka Ushiku[3] and Kanta Ono[1,2,5,*]

Level 2: suitable for a tech/science-savvy audience (e.g readers of Wired of MIT Tech Review)

- Suzuki and colleagues have used a deep neural network (a type of machine learning algorithm) to better understand relationships between the structure of a material and its properties. Such properties could include superconductivity, or magnetism, for example. The researchers trained their model on 120 000 known materials and the algorithm learned the key features of each material, then mapped that material to a point in a multi-dimensional space. The closer two materials are to one another in this space, the greater the similarity between their properties.

Levels of complexity

PAPER

CrossMark

OPEN ACCESS

RECEIVED
1 July 2022

Self-supervised learning of materials concepts from crystal structures via deep neural networks

Yuta Suzuki[1,2,6], Tatsunori Taniai[3], Kotaro Saito[2,4], Yoshitaka Ushiku[3] and Kanta Ono[1,2,5,*]

Suzuki and colleagues have used a deep neural network (a type of machine learning algorithm) to better understand relationships between the structure of a material and its properties. Such properties could include superconductivity, or magnetism, for example. The researchers trained their model on 120 000 known materials and the algorithm learned a representation of each material, mapping each one to a point in a multi-dimensional space. The closer two materials are to one another in this space, the greater the similarity between their properties.

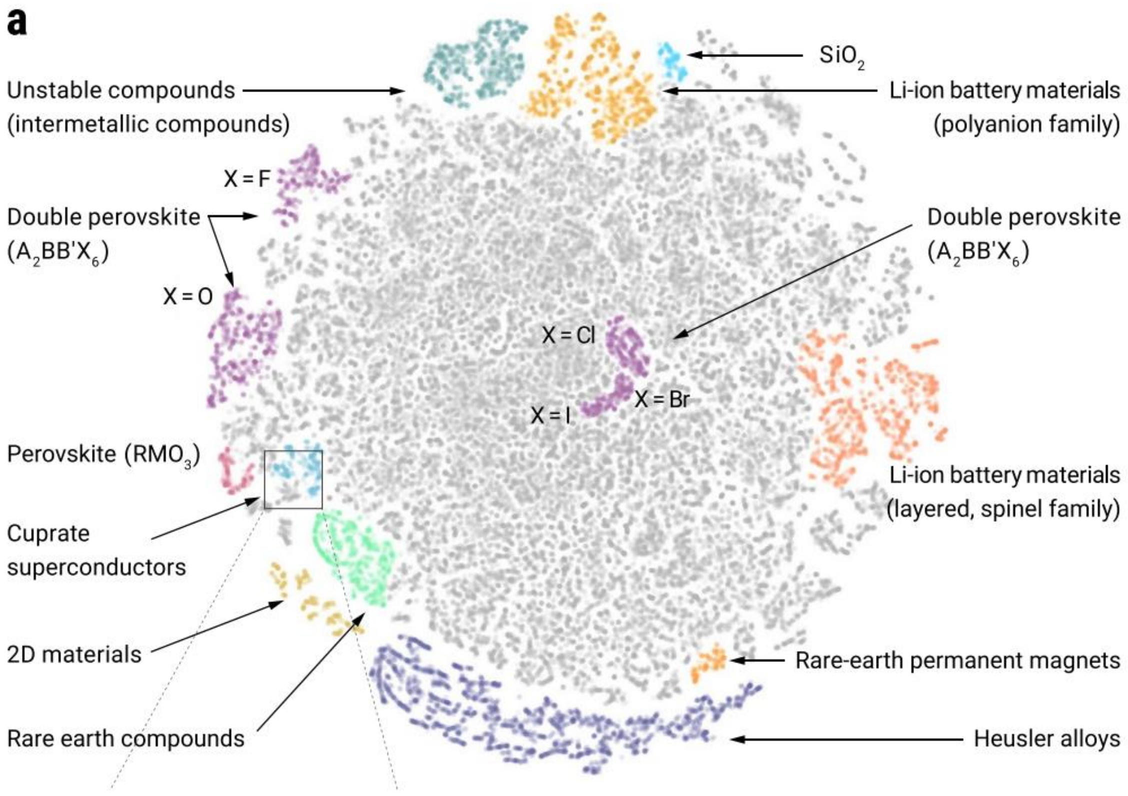

Use of an image to illustrate the mapping.

AIhub

Levels of complexity

CrossMark

OPEN ACCESS

RECEIVED
1 July 2022

PAPER

Self-supervised learning of materials concepts from crystal structures via deep neural networks

Yuta Suzuki[1,2,6], Tatsunori Taniai[3], Kotaro Saito[2,4], Yoshitaka Ushiku[3] and Kanta Ono[1,2,5,*]

Level 3: suitable for a more general audience.

- Researchers have used a machine learning algorithm to better understand materials and their properties. Such properties could include superconductivity, or magnetism, for example. The algorithm was fed data about over 120 000 different materials and used this information to group the materials according to the similarity of their properties. The method for clustering similar materials is like that used for recommender systems ("you've seen this film, so here's another you may like"). However, instead of the algorithm suggesting films similar to those you've seen before, it can indicate materials with similar properties.

AIhub

Creating a portfolio of media

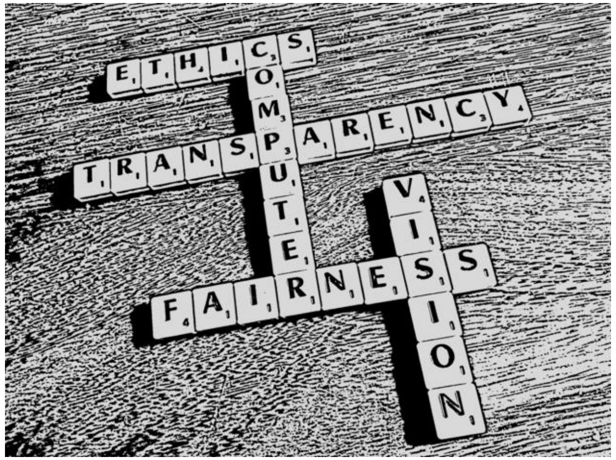

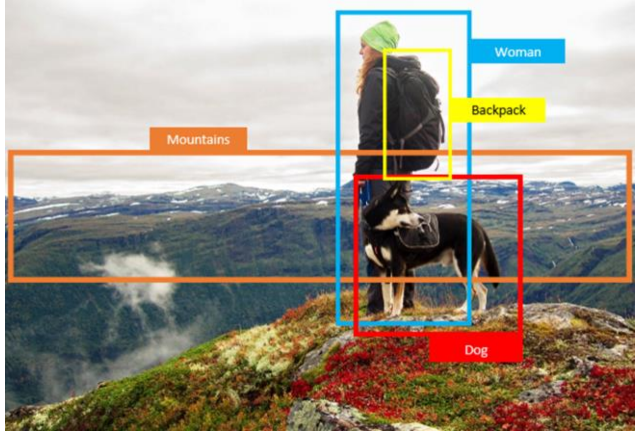

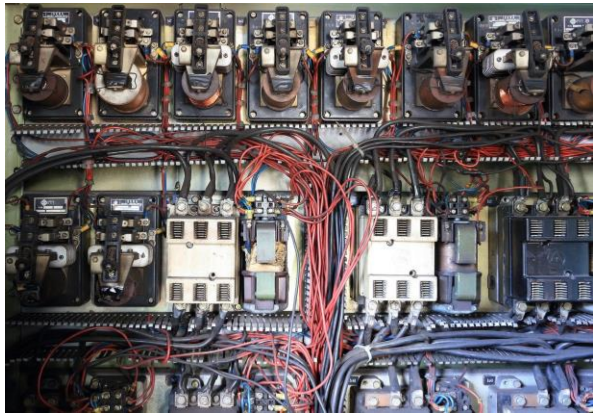

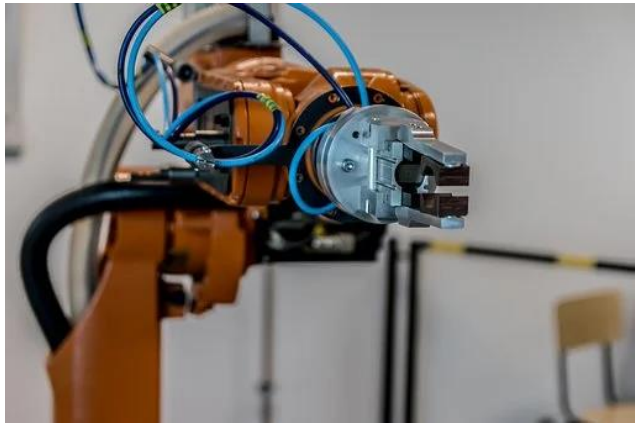

Creating a portfolio of media

Enhancing your blog posts with images and videos is important for two reasons.

- It can help increase the visual impact of your work.

Creating a portfolio of media

Enhancing your blog posts with images and videos is important for two reasons.

- It can help increase the visual impact of your work.
- It aids the understanding of concepts you are describing.

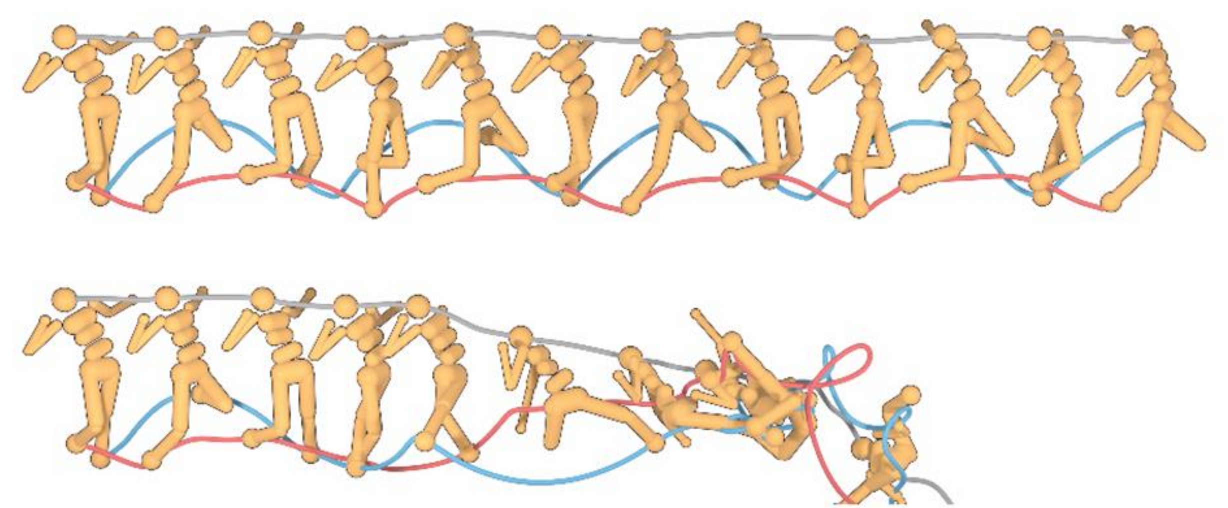

Credit: Michael Janner. From BAIR blog.

Creating a portfolio of media

- Option 1: use photos, graphs, images from your own research.

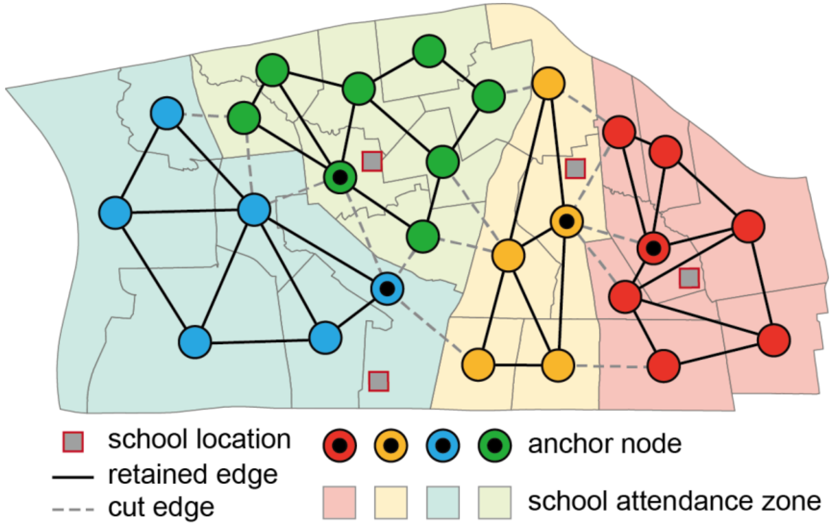

Credit: Fanglan Chen

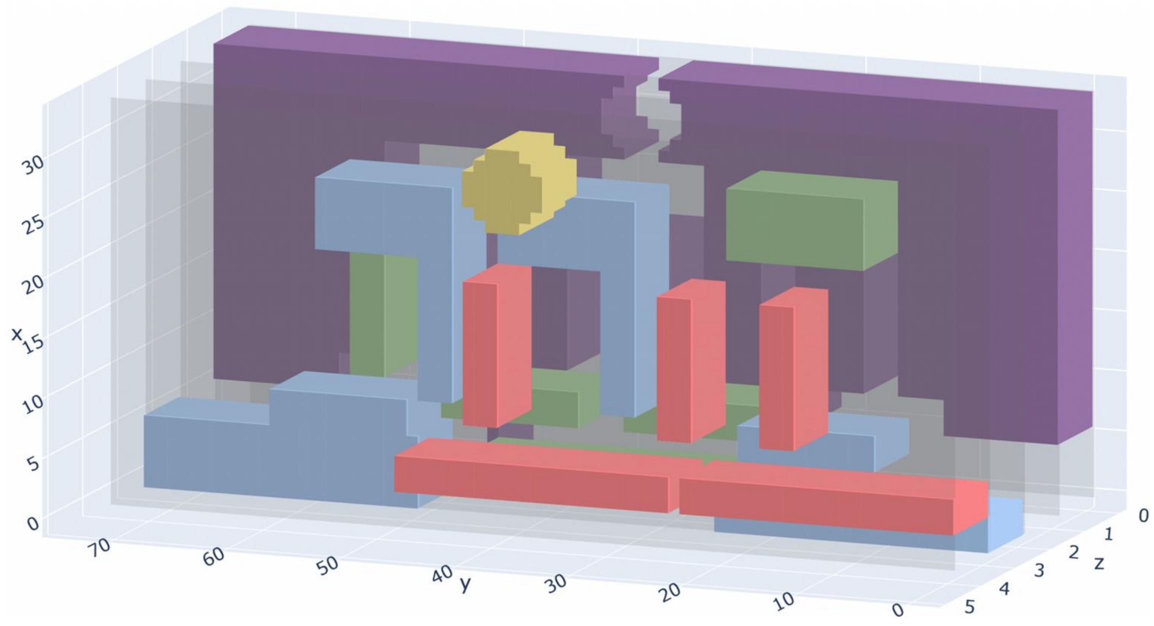

Credit: Matthew Stephenson and Frederic Abraham

Credit: Guillem Alenya

Creating a portfolio of media

- Option 1: use photos, graphs, images from your own research.
- Option 2: create your own images.

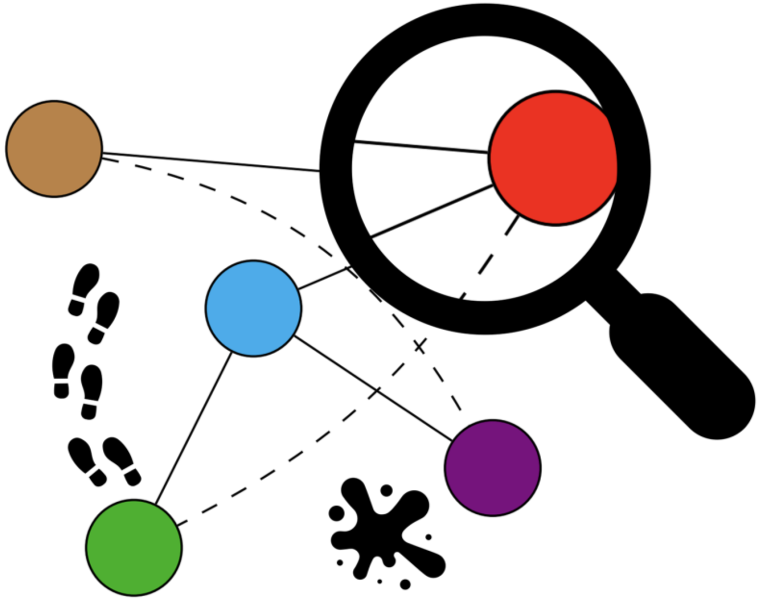

Credit: Ramon Fernández Mir and Lauren Nicole DeLong

AIhub

Creating a portfolio of media

- Option 1: use photos, graphs, images from your own research.
- Option 2: create your own images.
- Option 3: buy stock images.

Creating a portfolio of media

- Option 1: use photos, graphs, images from your own research.
- Option 2: create your own images.
- Option 3: buy stock images.
- Option 4: use images freely available online.

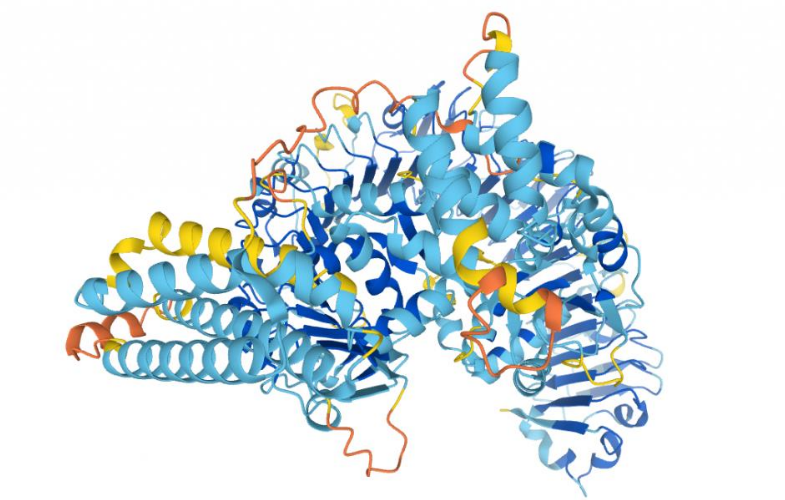

AIhub

Creating a portfolio of media

- Option 1: use photos, graphs, images from your own research.
- Option 2: create your own images.
- Option 3: buy stock images.
- Option 4: use images freely available online.
    - Be sure to check the license conditions for reproducing the image.

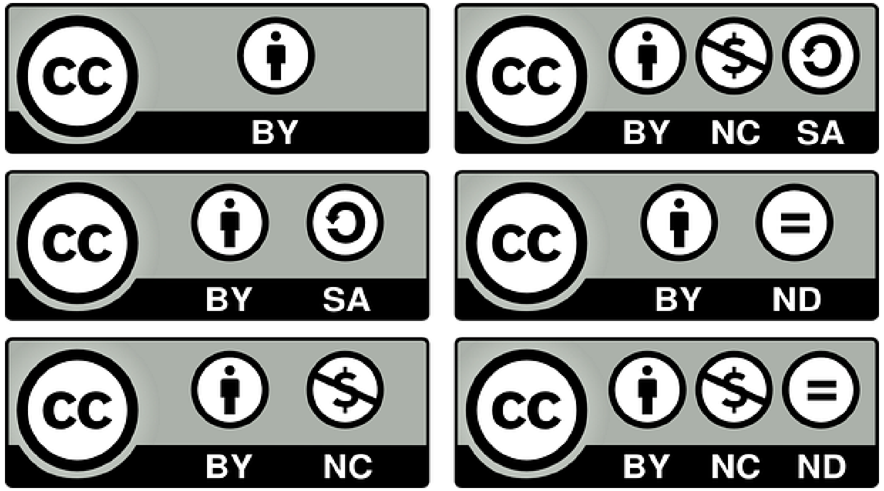

Creative commons licenses

- Creative Commons is a nonprofit organization that helps overcome legal obstacles to the sharing of knowledge and creativity.
- They provide Creative Commons licenses and public domain tools that give every person and organization in the world a free, simple, and standardized way to grant copyright permissions for creative and academic works; ensure proper attribution; and allow others to copy, distribute, and make use of those works

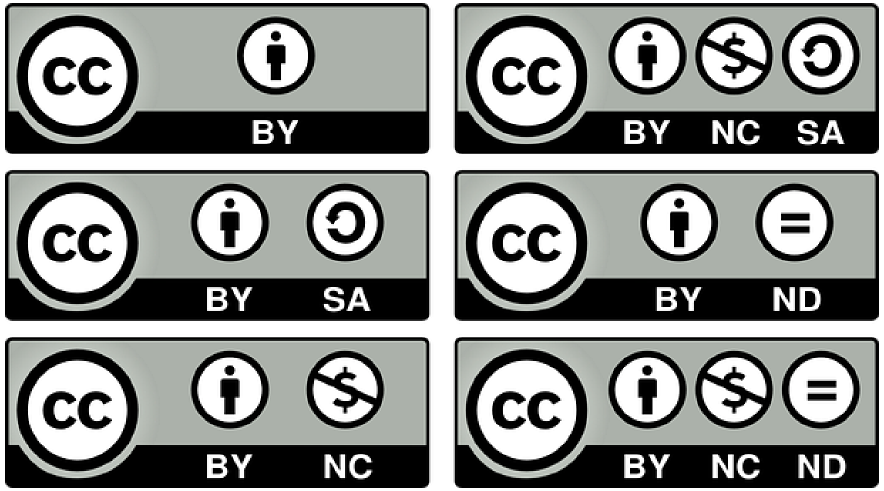

Creating a portfolio of media

- https://pixabay.com/
- https://unsplash.com/
- https://snappygoat.com/
- https://www.pexels.com/
- https://burst.shopify.com/
- https://www.flickr.com/

Better Images of AI

Better Images of AI

https://betterimagesofai.org/

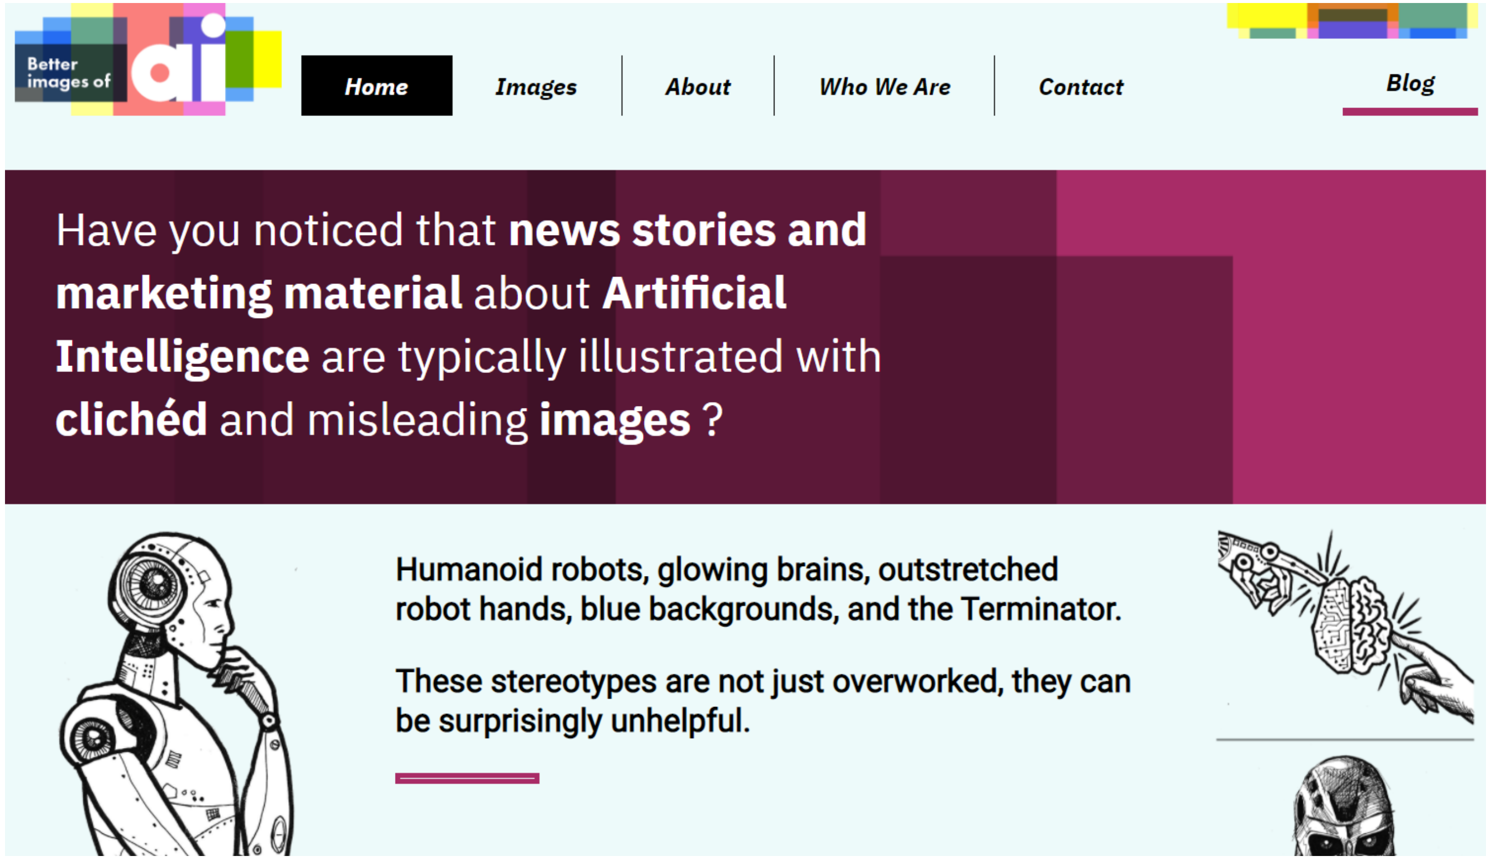

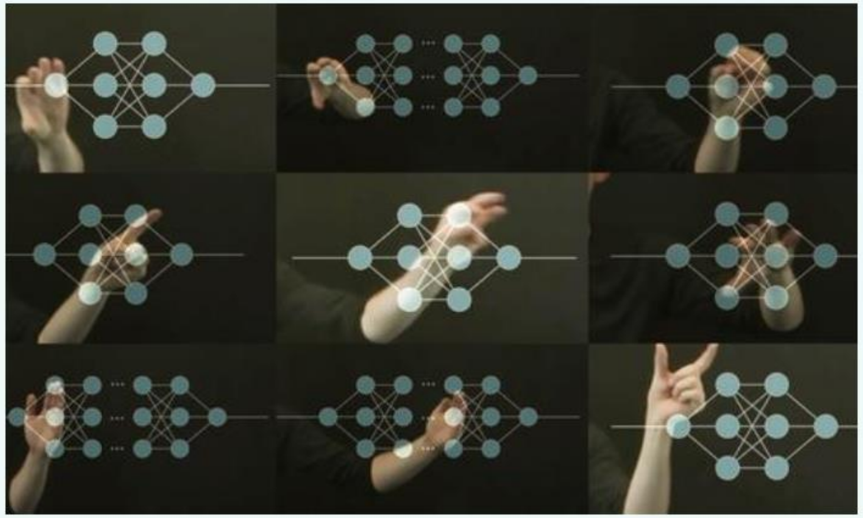

Explainable AI - Alexa Steinbrück

Fish reversed - Rens Dimmendaal & David Clode

Autonomous Driving - Anton Grabolle

Quantified Human - Alan Warburton

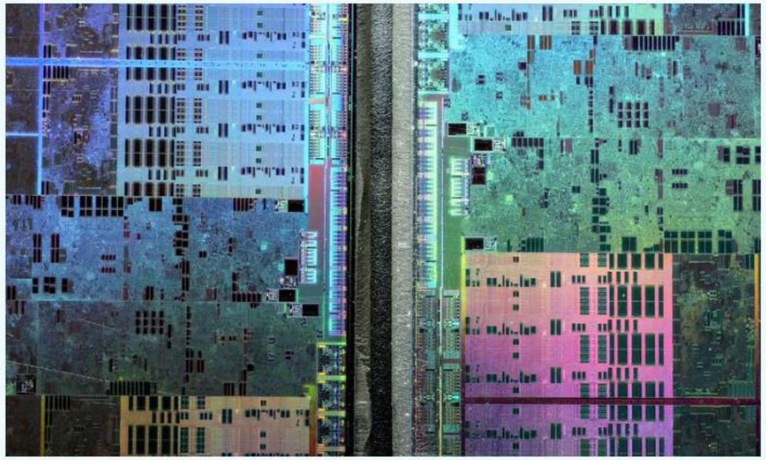

GPU shot etched 5 - Fritzchens Fritz

AIhub

Hype

Whilst it can be good to create a buzz around your research, too much hype tends to:

- Set inflated expectations about the technology,
- Drive unnecessary fears in the general public,
- Detracts from meaningful discussions about the actual aspects of the technology that we need to be concerned about.

AIhub

Tips for avoiding hype in your sci-comm

AIhub

Tips for avoiding hype in your sci-comm

- Don’t exaggerate the impact of your work:
  - Be specific about your contribution.
  - Make any limitations clear.

Tips for avoiding hype in your sci-comm

- Don’t exaggerate the impact of your work:
    - Be specific about your contribution.
    - Make any limitations clear.
- Avoid anthropomorphism.

Tips for avoiding hype in your sci-comm

- Don’t exaggerate the impact of your work:
  - Be specific about your contribution.
  - Make any limitations clear.
- Avoid anthropomorphism.
- Prioritise scientific accuracy in your headline.

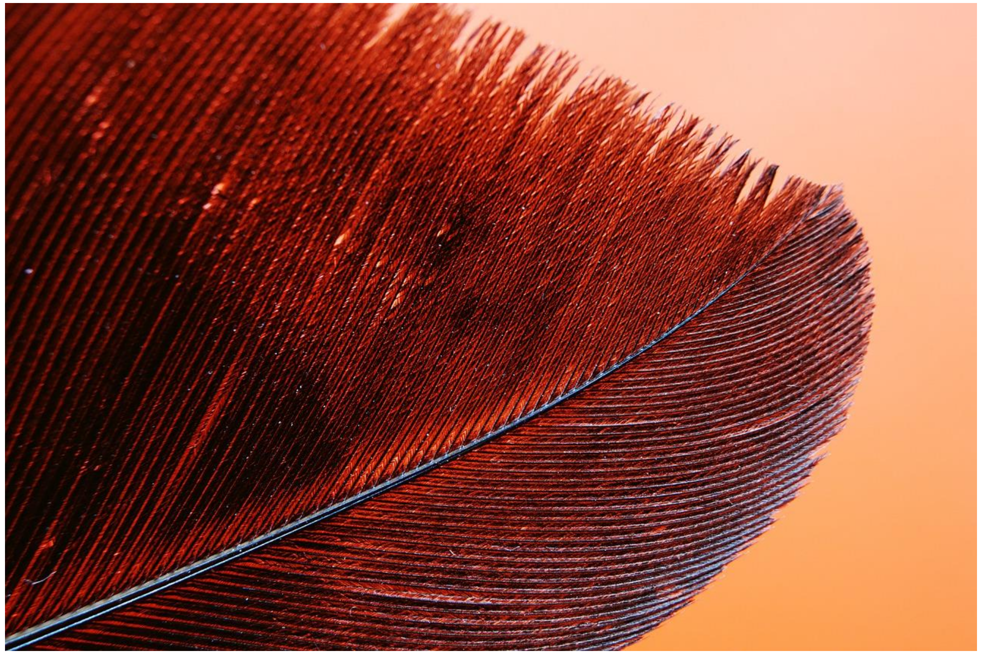

Tips for avoiding hype in your sci-comm

- Don’t exaggerate the impact of your work:
  - Be specific about your contribution.
  - Make any limitations clear.
- Avoid anthropomorphism.
- Prioritise scientific accuracy in your headline.
- Choose relevant images: avoid stereotypical images of robots from science fiction!

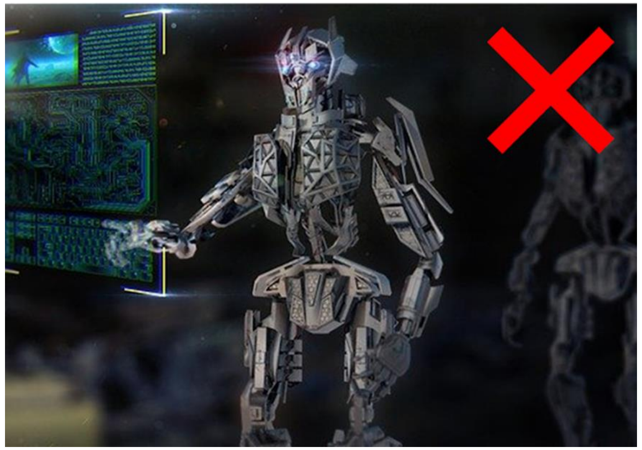

Check your blog post / tweets for hype

Unconventional ways of doing sci-comm

AIhub

Unconventional ways of doing sci-comm: *swarm escape!*

AIhub

Unconventional ways of doing sci-comm: *swarm escape!*

AIhub

Unconventional ways of doing sci-comm

Some examples to think about:

- Photograph essay
- Comic
- Stand-up monologue
- Short film
- Sci-fi book
- Food dish
- Escape room
- Sitcom
- Dance
- Theatre play
- Painting
- Sculpture
- Music festival performance
- Children's book
- Video game

Who is your audience, and could any of these formats help you communicate better?

Are there any aspects of your research that work with any of these formats?

AIhub

Next steps

- Try out some of the exercises from this talk.
- From 2-4pm: an informal session to discuss any ideas you have regarding sci-comm.
- Interested in covering NeurIPS for AIhub?
- Reach out to us - we can work with you to help you shape your story.

aihuborg@gmail.com

https://aihub.org

AIhub

Acknowledgements

Dr Daniel Carrillo-Zapata
Robohub, scicomm.io

Professor Sabine Hauert,
University of Bristol

AIhub

news articles opinions education

https://aihub.org

aihuborg@gmail.com

@aihuborg

aihuborg

aihuborg

@Aihuborg

Questions?

Feedback form

https://forms.gle/hXtoTZp3X85bvS1U8