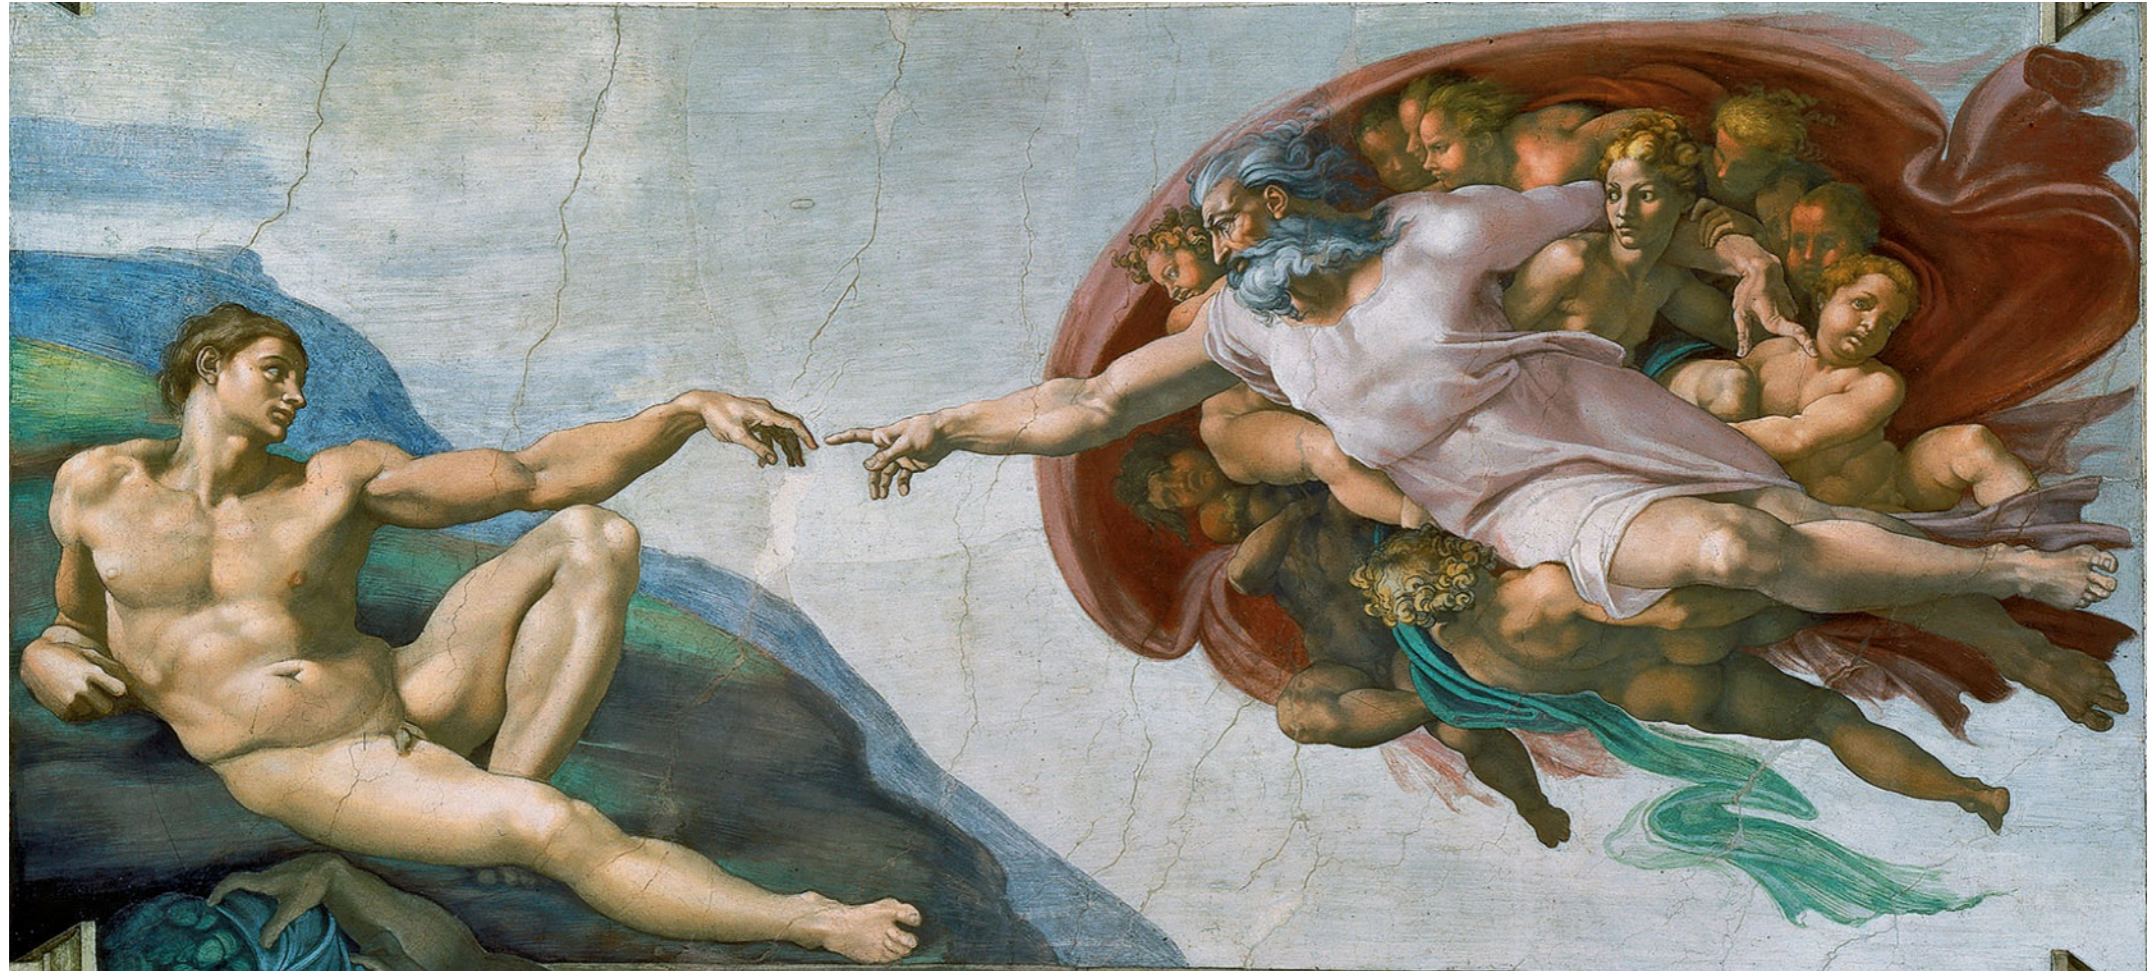


HOW TO WRITE PAPERS SO PEOPLE CAN READ THEM




Derek Dreyer

Max Planck Institute for Software Systems

PLMW@ICFP 2021



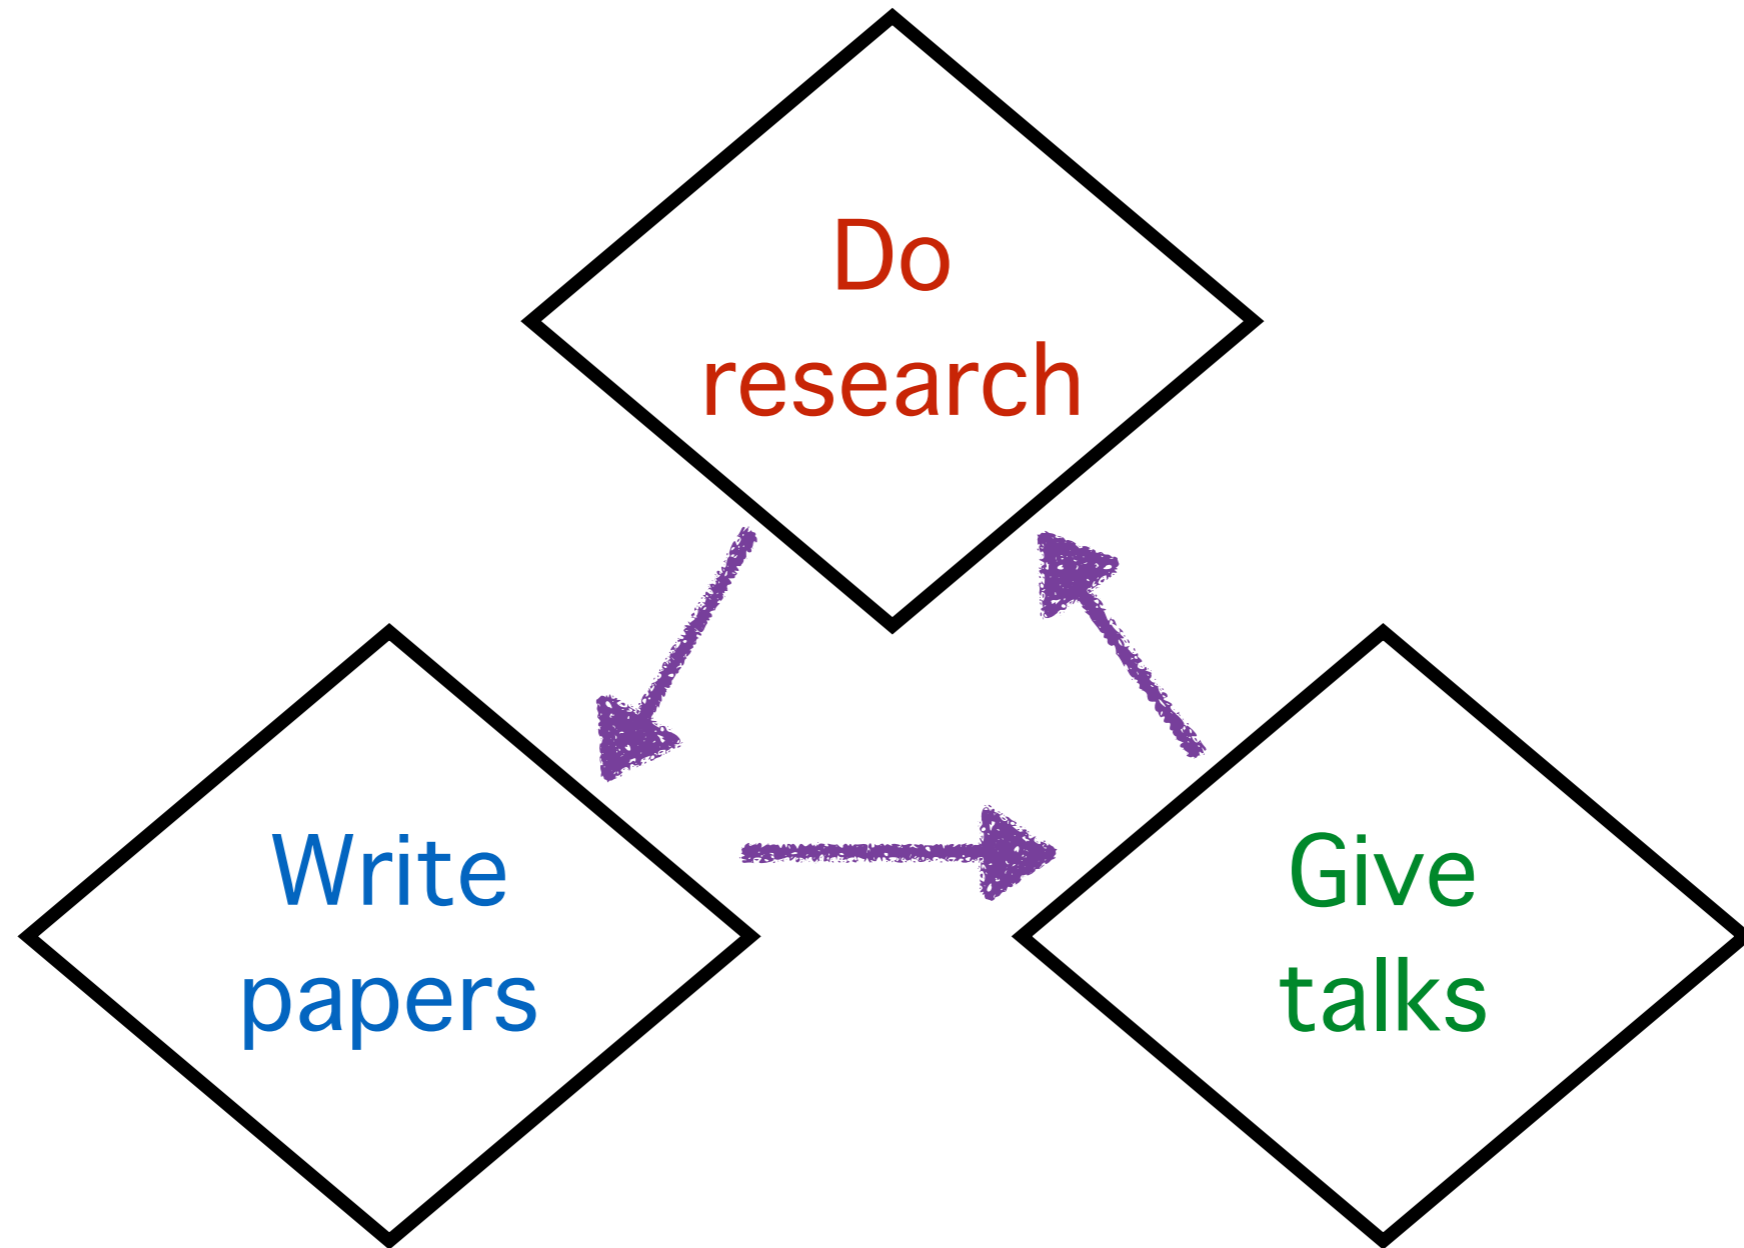
A photograph of a living room. In the center, a small, light-colored dog is lying on a dark brown armchair. The chair has an orange and white striped cushion. To the left of the chair is a large potted plant with long, thin green leaves. To the right is a black side table with a potted plant and some items on it. The floor is made of light-colored wood. In the foreground, several small bottles are scattered on the floor. A white text box is overlaid on the image, containing the text "Check out my blog at herrdreyaer.wordpress.com".

Check out my blog at
herrdreyaer.wordpress.com

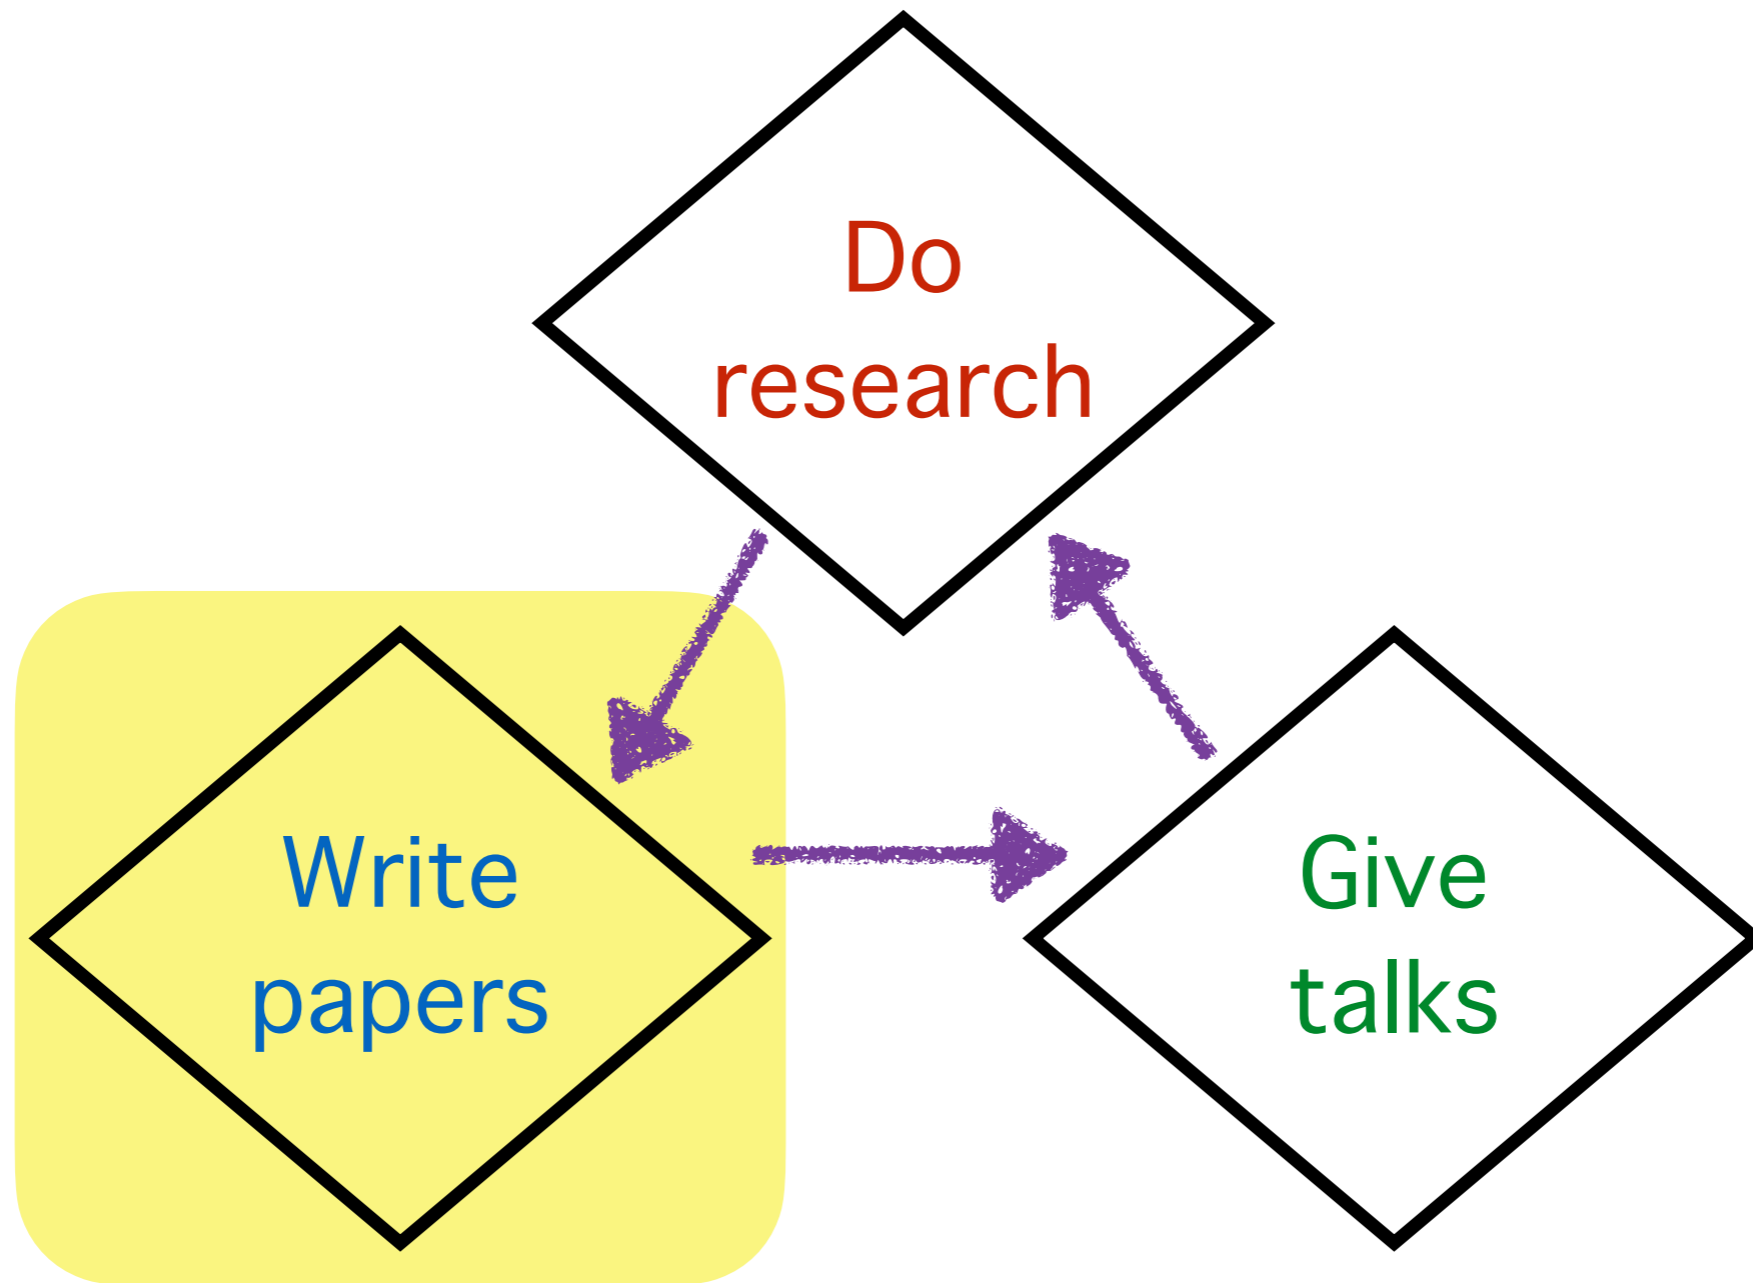
My job as a researcher



My job as a researcher



My job as a researcher



Have you read any
research papers lately?

Have you read any research papers lately?



Have you read any research papers lately?

- You may think you just lack the technical sophistication to understand them.



Have you read any research papers lately?

- You may think you just lack the technical sophistication to understand them.
- But in fact, many papers are **poorly written**.



So if you can write clear, accessible papers...

- People will **enjoy** reading them!
- People will **learn** something from them!
- They will get **accepted** to top conferences!

Fame



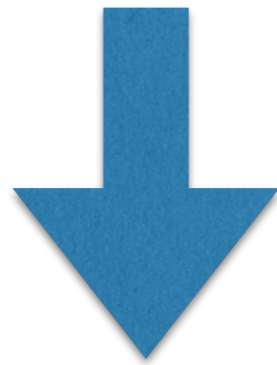
So if you can write clear, accessible papers...

- People will **enjoy** reading them!
- People will **learn** something from them!
- They will get **accepted** to top conferences!

Fame



A piece of research



Writer

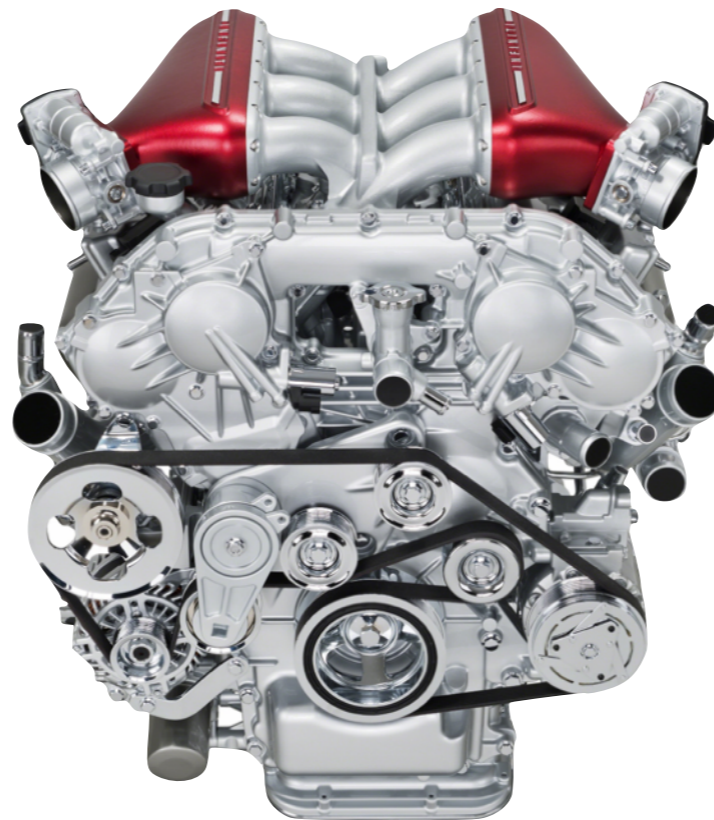


Reader

By downcasting the pre-axial gaskets,
we achieved 47% reduction in XPS latency
on the re-uptake bivalve!



Writer



Reader

By downcasting the pre-axial gaskets,
we achieved 47% reduction in XPS latency
on the re-uptake bivalve!

OK, but what does it **do**,
and why do I **care**?



Writer



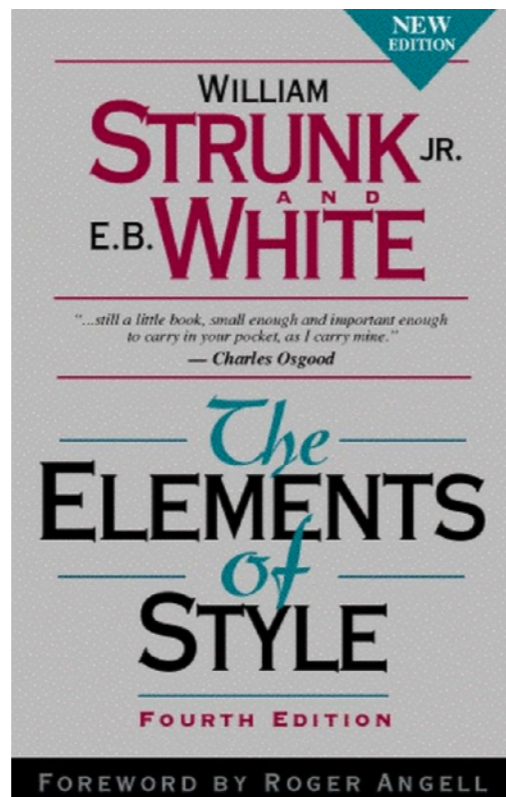
Reader

The good news

- There are **principles** you can follow that will help you write clearer, more readable prose
 - Based on how readers process information

The good news

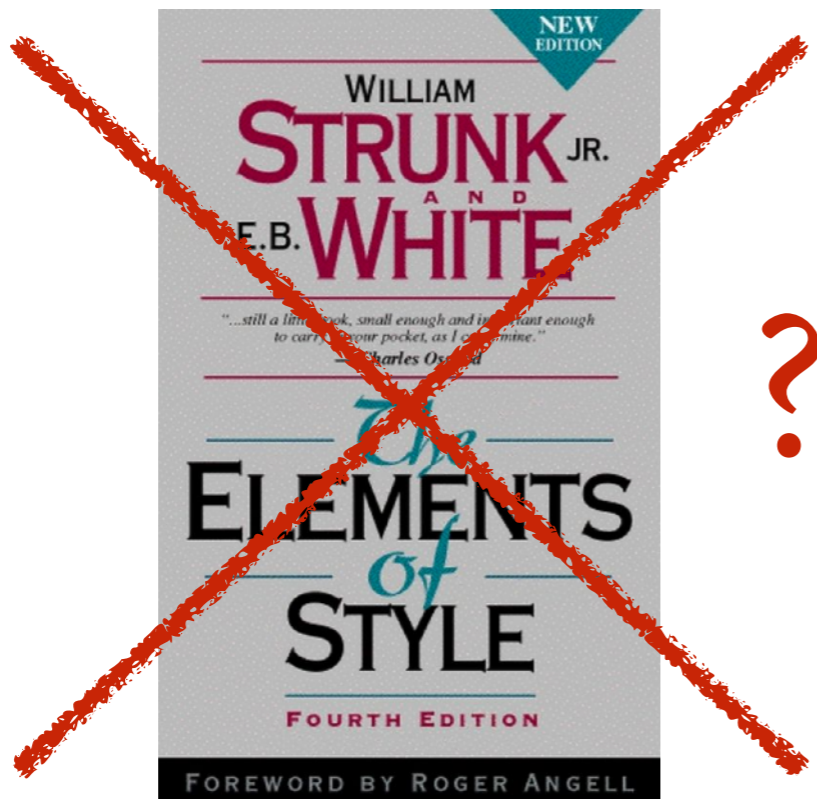
- There are **principles** you can follow that will help you write clearer, more readable prose
 - Based on how readers process information



?

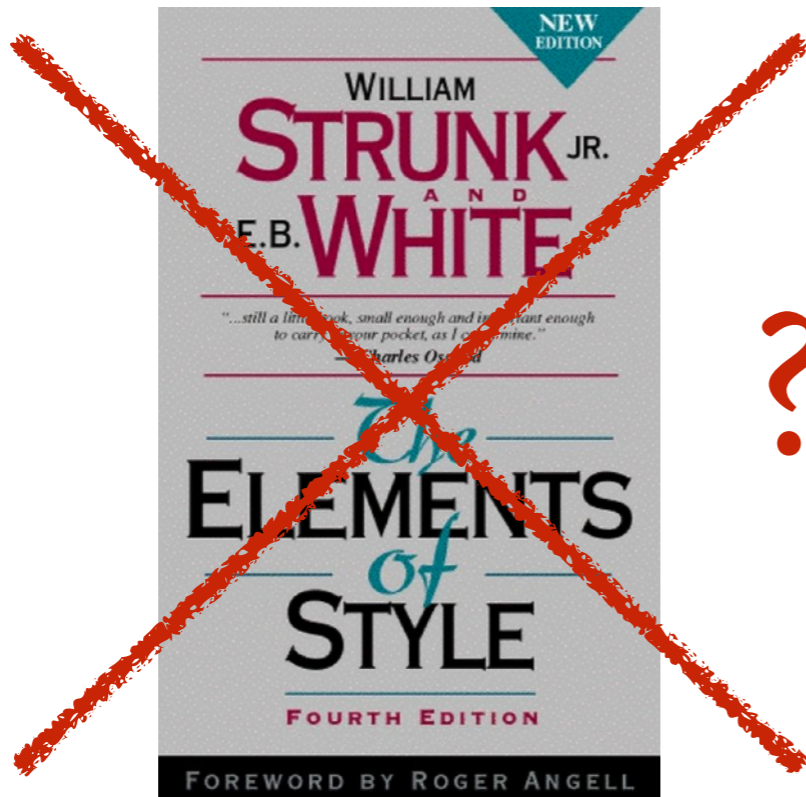
The good news

- There are **principles** you can follow that will help you write clearer, more readable prose
 - Based on how readers process information



The good news

- There are **principles** you can follow that will help you write clearer, more readable prose
 - Based on how readers process information



?

“Be clear”

“Omit needless words”

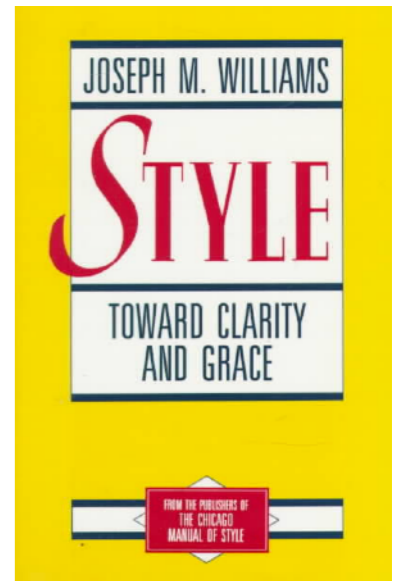
...

The good news

- There are **principles** you can follow that will help you write clearer, more readable prose
 - Based on how readers process information
- These principles are **constructive**:
 - Easy to check if your text satisfies these principles
 - If not, principles suggest improvements

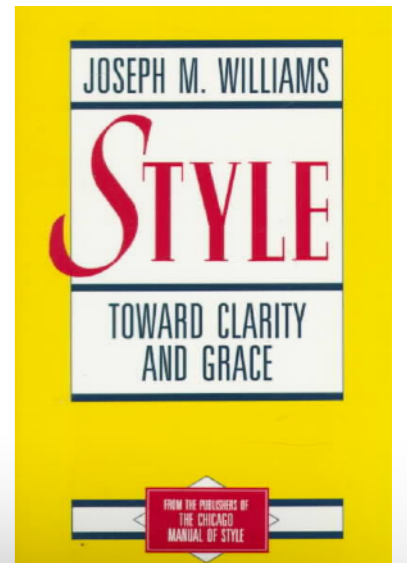
Inspirations for this talk

- **Joseph M. Williams.** *Style: Toward clarity and grace.* 1990. (book)
- **Norman Ramsey.** *Learn technical writing in two hours per week.* (course notes)
 - <http://www.cs.tufts.edu/~nr/pubs/two.pdf>
- **Simon Peyton Jones.** *How to write a great research paper.* (talk)
 - <https://www.microsoft.com/en-us/research/video/how-to-write-a-great-research-paper-3/>



Inspirations for this talk

- **Joseph M. Williams.** *Style: Toward clarity and grace.* 1990. (book)



Talk developed jointly with
Rose Hoberman
@ MPI-SWS



- **Simon Peyton Jones.** *How to write a great research paper.* (talk)



- <https://www.microsoft.com/en-us/research/video/how-to-write-a-great-research-paper-3/>

Sentences & paragraphs

Flow



It should be clear how each sentence and paragraph relates to **the adjacent ones**

Does this text flow?

Does this text flow?

Security proofs of cryptographic protocols are crucial for the security of everyday electronic communication. However, these proofs tend to be complex and difficult to get right. The game-playing technique, originally proposed by Jones et al., follows a code-based approach where the security properties are formulated in terms of probabilistic programs, called games. This is a general design principle for cryptographic proofs to ease their management.

Does this text flow?

Security proofs of cryptographic protocols are crucial for the security of everyday electronic communication.

However, these proofs tend to be complex and difficult to get right. The game-playing technique, originally proposed by Jones et al., follows a code-based approach where the security properties are formulated in terms of probabilistic programs, called games. This is a general design principle for cryptographic proofs to ease their management.

Does this text flow?

Security proofs of cryptographic protocols are crucial for the security of everyday electronic communication.

However, these proofs tend to be complex and difficult to get right.

The game-playing technique, originally proposed by Jones et al., follows a code-based approach where the security properties are formulated in terms of probabilistic programs, called games. This is a general design principle for cryptographic proofs to ease their management.

Does this text flow?

Security proofs of cryptographic protocols are crucial for the security of everyday electronic communication. However, these proofs tend to be complex and difficult to get right. The game-playing technique, originally proposed by Jones et al., follows a code-based approach where the security properties are formulated in terms of probabilistic programs, called games. This is a general design principle for cryptographic proofs to ease their management.

Does this text flow?

Security proofs of cryptographic protocols are crucial for the security of everyday electronic communication. However, these proofs tend to be complex and difficult to get right. The game-playing technique, originally proposed by Jones et al., follows a code-based approach where the security properties are formulated in terms of probabilistic programs, called games. This is a general design principle for cryptographic proofs to ease their management.

Does this text flow?

Security proofs of cryptographic protocols are crucial for the security of everyday electronic communication. However, these proofs tend to be complex and difficult to get right. The game-playing technique, originally proposed by Jones et al., follows a code-based approach where the security properties are formulated in terms of probabilistic programs, called games. This is a general design principle for cryptographic proofs to ease their management.

Does this text flow?

Security proofs of cryptographic protocols are crucial for the security of everyday electronic communication. However, these proofs tend to be complex and difficult to get right. The game-playing technique, originally proposed by Jones et al., follows a code-based approach where the security properties are formulated in terms of probabilistic programs, called games. This is a general design principle for cryptographic proofs to ease their management.

Does this text flow?

Security proofs of cryptographic protocols are crucial for the security of everyday electronic communication.

However, these proofs tend to be complex and difficult to get right.

The game-playing technique, originally proposed by Jones et al., follows a code-based approach



What does this game-playing technique have to do with what came before?

Old to new

- Begin sentences with old info
 - Creates link to earlier text
- End sentences with new info
 - Creates link to the text that follows
 - Also places new info in position of **emphasis**



Applying old-to-new

New information

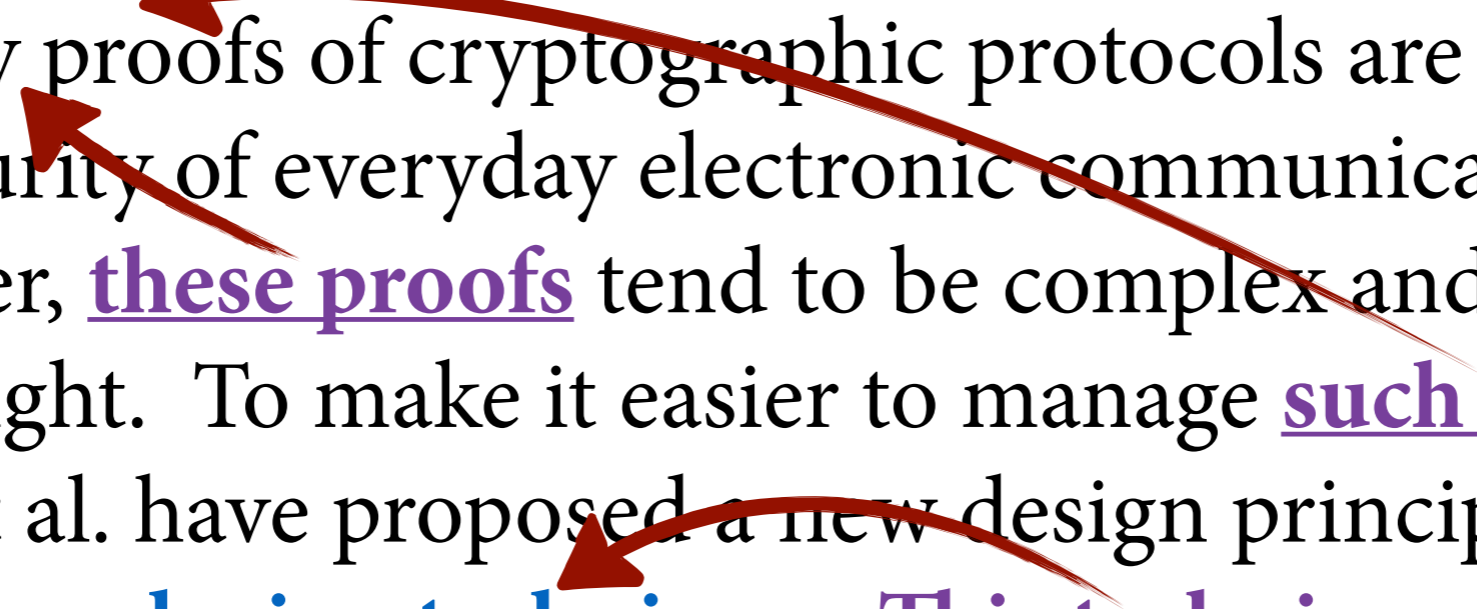
Security proofs of cryptographic protocols are crucial for the security of everyday electronic communication. However, these proofs tend to be complex and difficult to get right. **The game-playing technique, originally proposed by Jones et al., follows a code-based approach where the security properties are formulated in terms of probabilistic programs, called games.** This is a general design principle for cryptographic proofs to ease their management.

Applying old-to-new

Security proofs of cryptographic protocols are crucial for the security of everyday electronic communication. However, these proofs tend to be complex and difficult to get right. To make it easier to manage such proofs, Jones et al. have proposed a new design principle, called the game-playing technique. This technique follows a code-based approach where the security properties are formulated in terms of probabilistic programs, called games.

Old-to-new satisfied

Security proofs of cryptographic protocols are crucial for the security of everyday electronic communication. However, these proofs tend to be complex and difficult to get right. To make it easier to manage such proofs, Jones et al. have proposed a new design principle, called the **game-playing technique**. This technique follows a code-based approach where the security properties are formulated in terms of probabilistic programs, called games.



The diagram consists of two red curved arrows. The first arrow starts from the underlined phrase 'these proofs' and points to the underlined phrase 'such proofs'. The second arrow starts from the underlined phrase 'This technique' and points back to the underlined phrase 'these proofs'. This visualizes a cycle of satisfaction or a transition from an old state to a new one.

But flow is not enough!

What about this text?

Lions and tigers are some of the most dramatic and awe-inspiring species of cats. Most of these large cats, however, are currently facing extinction. A smaller cat that has been more evolutionarily successful is the house cat. Although house cats are currently the most popular pet in the world, they are in many ways anti-social. It would therefore be interesting to study whether house cats can be trained to be more sociable.

What about this text?

Lions and tigers are some of the most dramatic and awe-inspiring species of cats. Most of these large cats, however, are currently facing extinction. A smaller cat that has been more evolutionarily successful is the house cat. Although house cats are currently the most popular pet in the world, they are in many ways anti-social. It would therefore be interesting to study whether house cats can be trained to be more sociable.

What about this text?

Lions and tigers are some of the most dramatic and awe-inspiring species of cats. **Most of these large cats, however, are currently facing extinction.** A smaller cat that has been more evolutionarily successful is the house cat. Although house cats are currently the most popular pet in the world, they are in many ways anti-social. It would therefore be interesting to study whether house cats can be trained to be more sociable.

What about this text?

Lions and tigers are some of the most dramatic and awe-inspiring species of cats. Most of these large cats, however, are currently facing extinction. **A smaller cat that has been more evolutionarily successful is the house cat.** Although house cats are currently the most popular pet in the world, they are in many ways anti-social. It would therefore be interesting to study whether house cats can be trained to be more sociable.

What about this text?

Lions and tigers are some of the most dramatic and awe-inspiring species of cats. Most of these large cats, however, are currently facing extinction. A smaller cat that has been more evolutionarily successful is the house cat. **Although house cats are currently the most popular pet in the world, they are in many ways anti-social.** It would therefore be interesting to study whether house cats can be trained to be more sociable.

What about this text?

Lions and tigers are some of the most dramatic and awe-inspiring species of cats. Most of these large cats, however, are currently facing extinction. A smaller cat that has been more evolutionarily successful is the house cat. Although house cats are currently the most popular pet in the world, they are in many ways anti-social. It would therefore be interesting to study whether house cats can be trained to be more sociable.

What about this text?

Lions and tigers are some of the most dramatic and awe-inspiring species of cats. Most of these large cats, however, are currently facing extinction. A smaller cat that has been more evolutionarily successful is the house cat. Although house cats are currently the most popular pet in the world, they are in many ways anti-social. It would therefore be interesting to study whether house cats can be trained to be more sociable.

What about this text?

Lions and tigers are some of the most dramatic and awe-inspiring species of cats. The large cats, however, are not the only ones. A smaller, less dramatic, but equally successful species of cats are currently being studied. They are in fact, very social. It would therefore be interesting to study whether house cats can be trained to be more sociable.

**Has great flow,
but is incoherent!**

Coherence



It should be clear how each
sentence and paragraph relates to
the big picture

One paragraph, one point

- A paragraph should have one main point, expressed in a single **point sentence**
- **Typically** the point sentence should appear **at or near the beginning of the paragraph**



Get to the
point!

No point sentence

Lions and tigers are some of the most dramatic and awe-inspiring species of cats. Most of these large cats, however, are currently facing extinction. A smaller cat that has been more evolutionarily successful is the house cat. Although house cats are currently the most popular pet in the world, they are in many ways anti-social. It would therefore be interesting to study whether house cats can be trained to be more sociable.

Point sentence up front

There appears to be a negative correlation between the charisma of a species and its ability to survive.

Lions and tigers, for instance, are among the most majestic creatures in the animal kingdom, yet they are currently facing extinction. In contrast, the house cat is evolutionarily quite successful, even though it is mostly known for stupid pet tricks.

Flow & coherence



Create **flow** with **old to new**

Create **coherence** with
one paragraph, one point



Two other principles



- **Name your baby:**
 - Give unique names to things and use them consistently



- **Just in time:**
 - Give information precisely when it is needed, not before

Three other principles



Bonus principle from Rose

Short subjects:

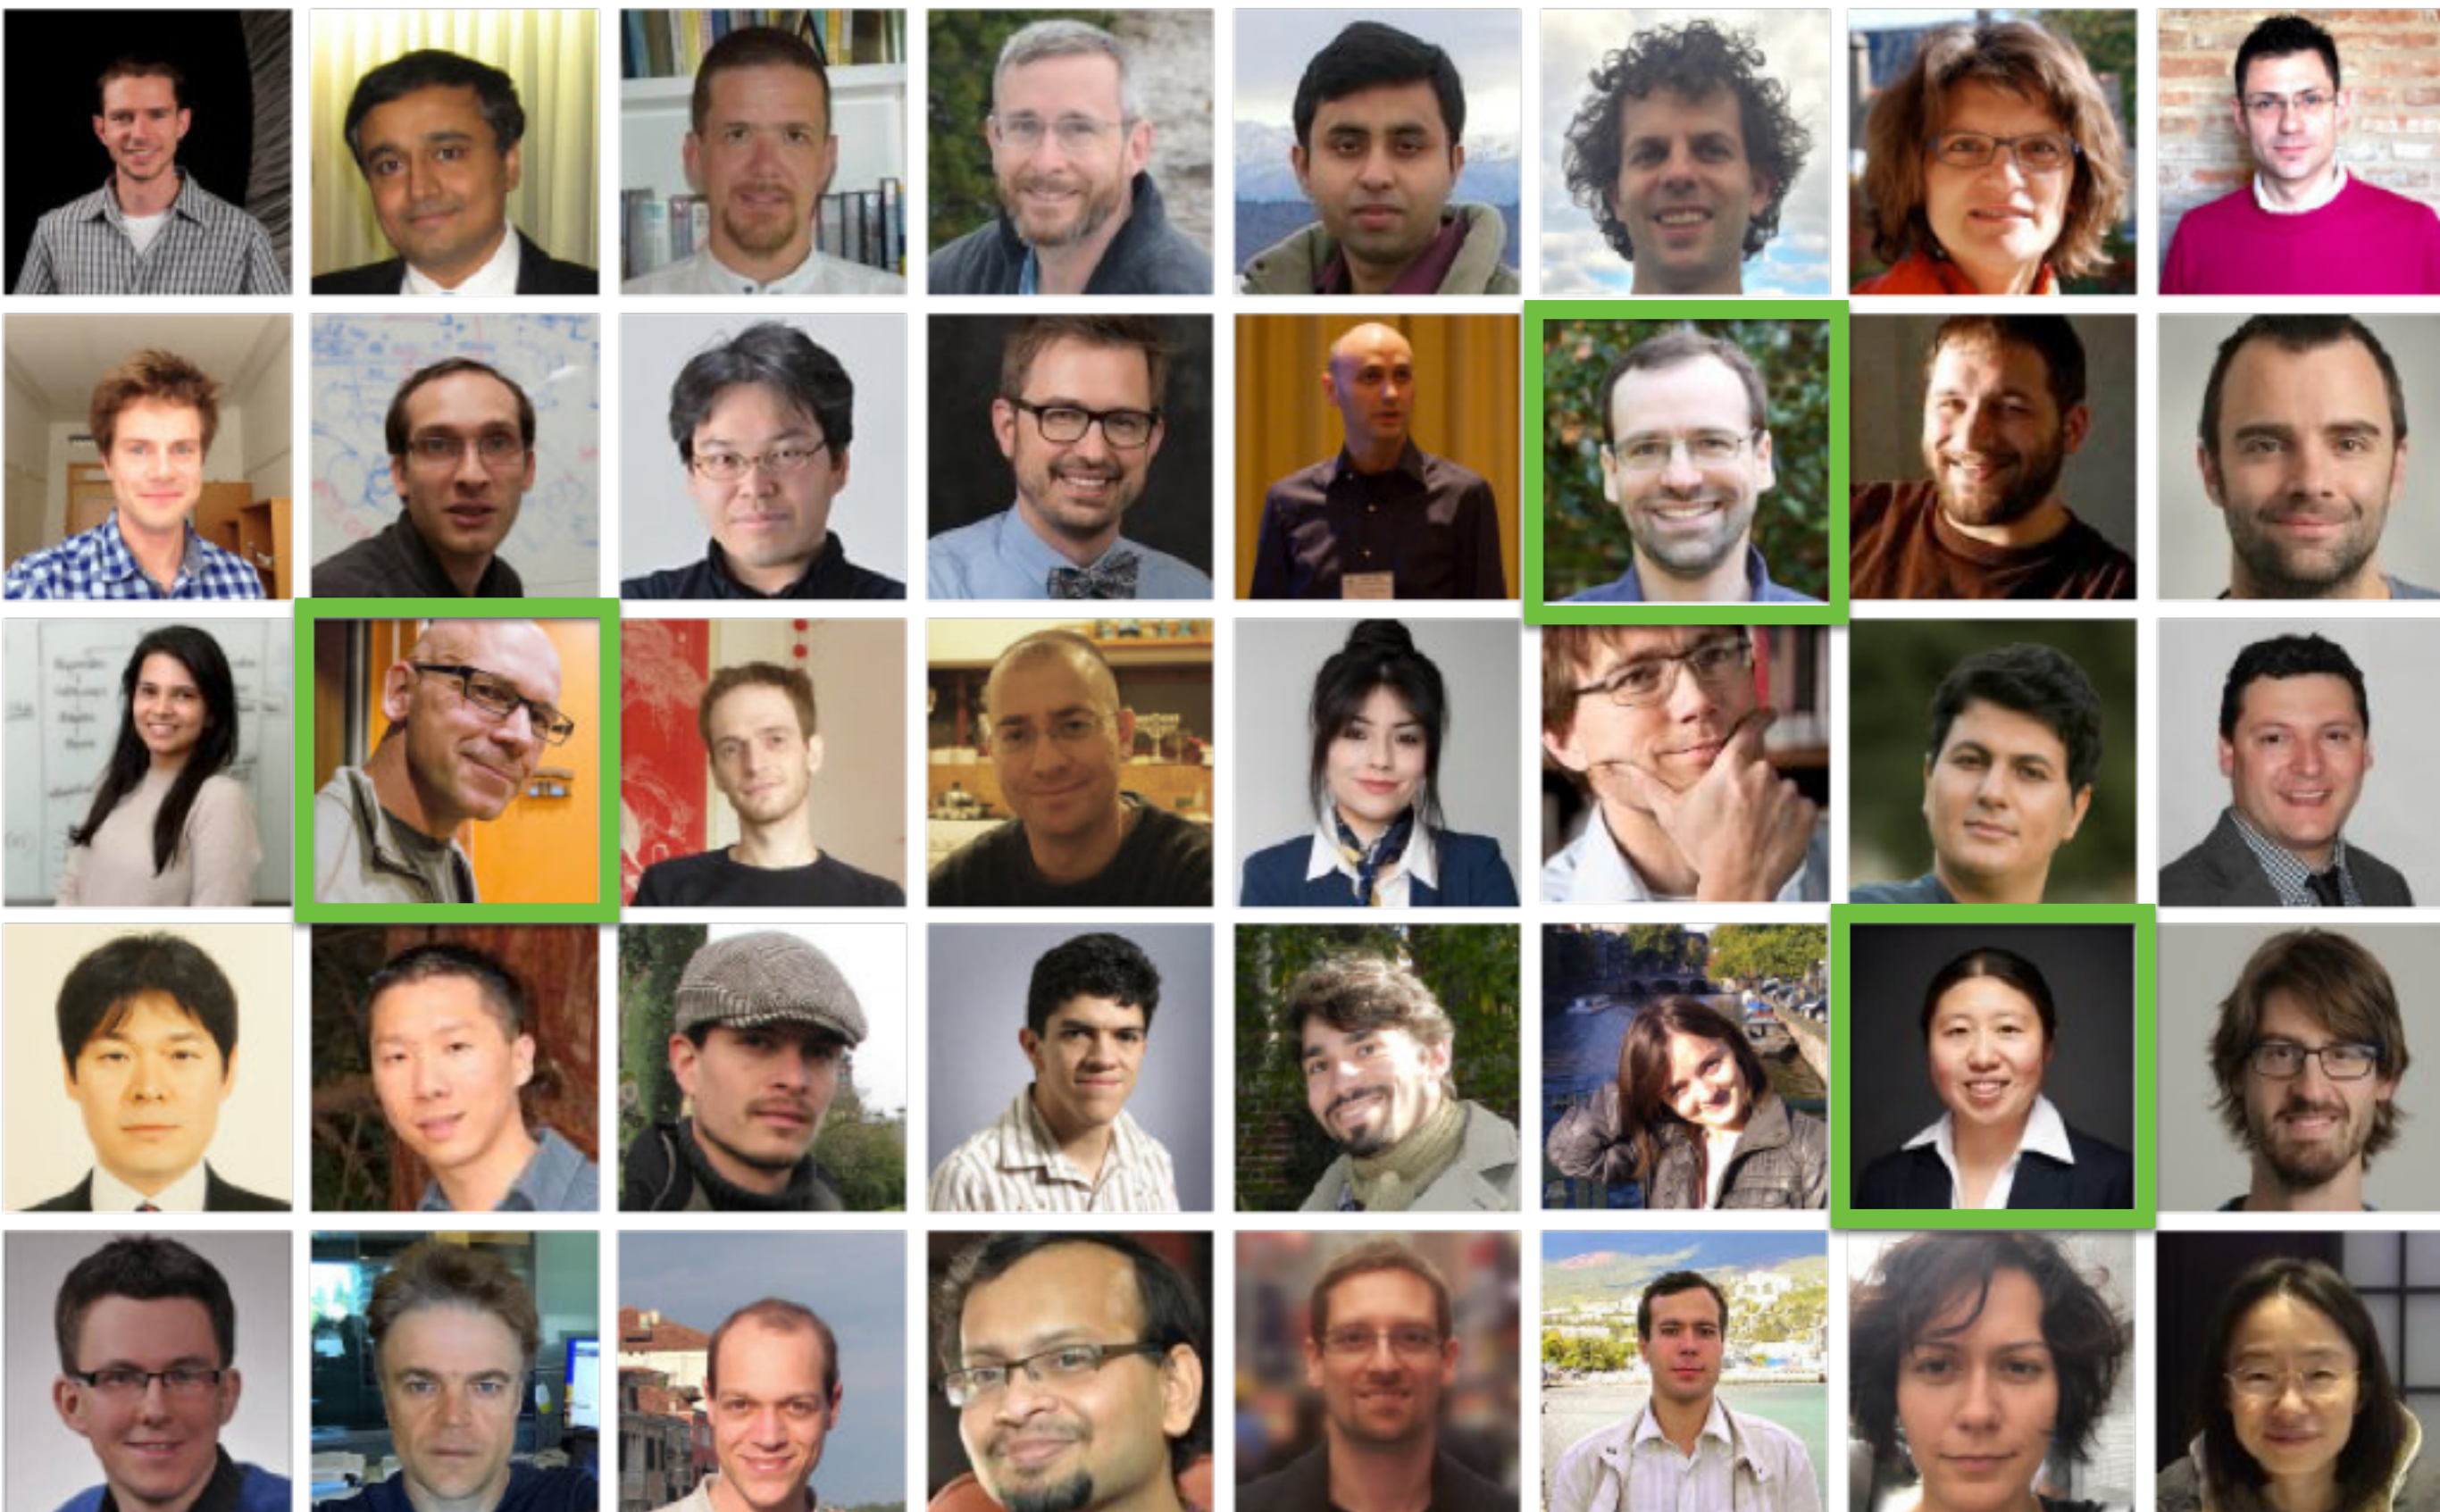
Subject of sentence should be
at most 8 words long



- Give information precisely
when it is needed, not before

Structure of a research paper







Program committee (PC) is diverse:

- Your reviewers may not be “experts”

Review period is short (~2 months):

- ~10-20 papers per PC member
- Often < 1 day to review each paper

Overarching Principle #1

TOP-DOWN

Explain your work at multiple levels of abstraction,
starting at a high level (accessible to non-experts)
and getting progressively more detailed

Overarching Principle #2

Tell them what they
want to know



Overarching Principle #2

Tell them what they
want to know



- ☒ How is your work **important**?
- ☒ How is your work **novel**?
- ☒ How is your work **interesting**?

Overarching Principle #2

Tell them what they
want to know



- ☒ How is your work **important**?
- ☒ How is your work **novel**?
- ☒ How is your work **interesting**?
- ☐ ~~How was your work **challenging**?~~

A structure that works

- **Abstract** (1-2 paragraphs, 1000 readers)
- **Intro** (2-4 pages, 100 readers)
- **Key ideas** (4-6 pages, 50 readers)
- **Technical meat** (8-12 pages, 5 readers)
- **Related work** (1-3 pages, 100 readers)


A structure that works

- **Abstract** (1-2 paragraphs, 1000 readers)
- **Intro** (2-4 pages, 100 readers)
- Key ideas (4-6 pages, 50 readers)
- Technical meat (8-12 pages, 5 readers)
- Related work (1-3 pages, 100 readers)

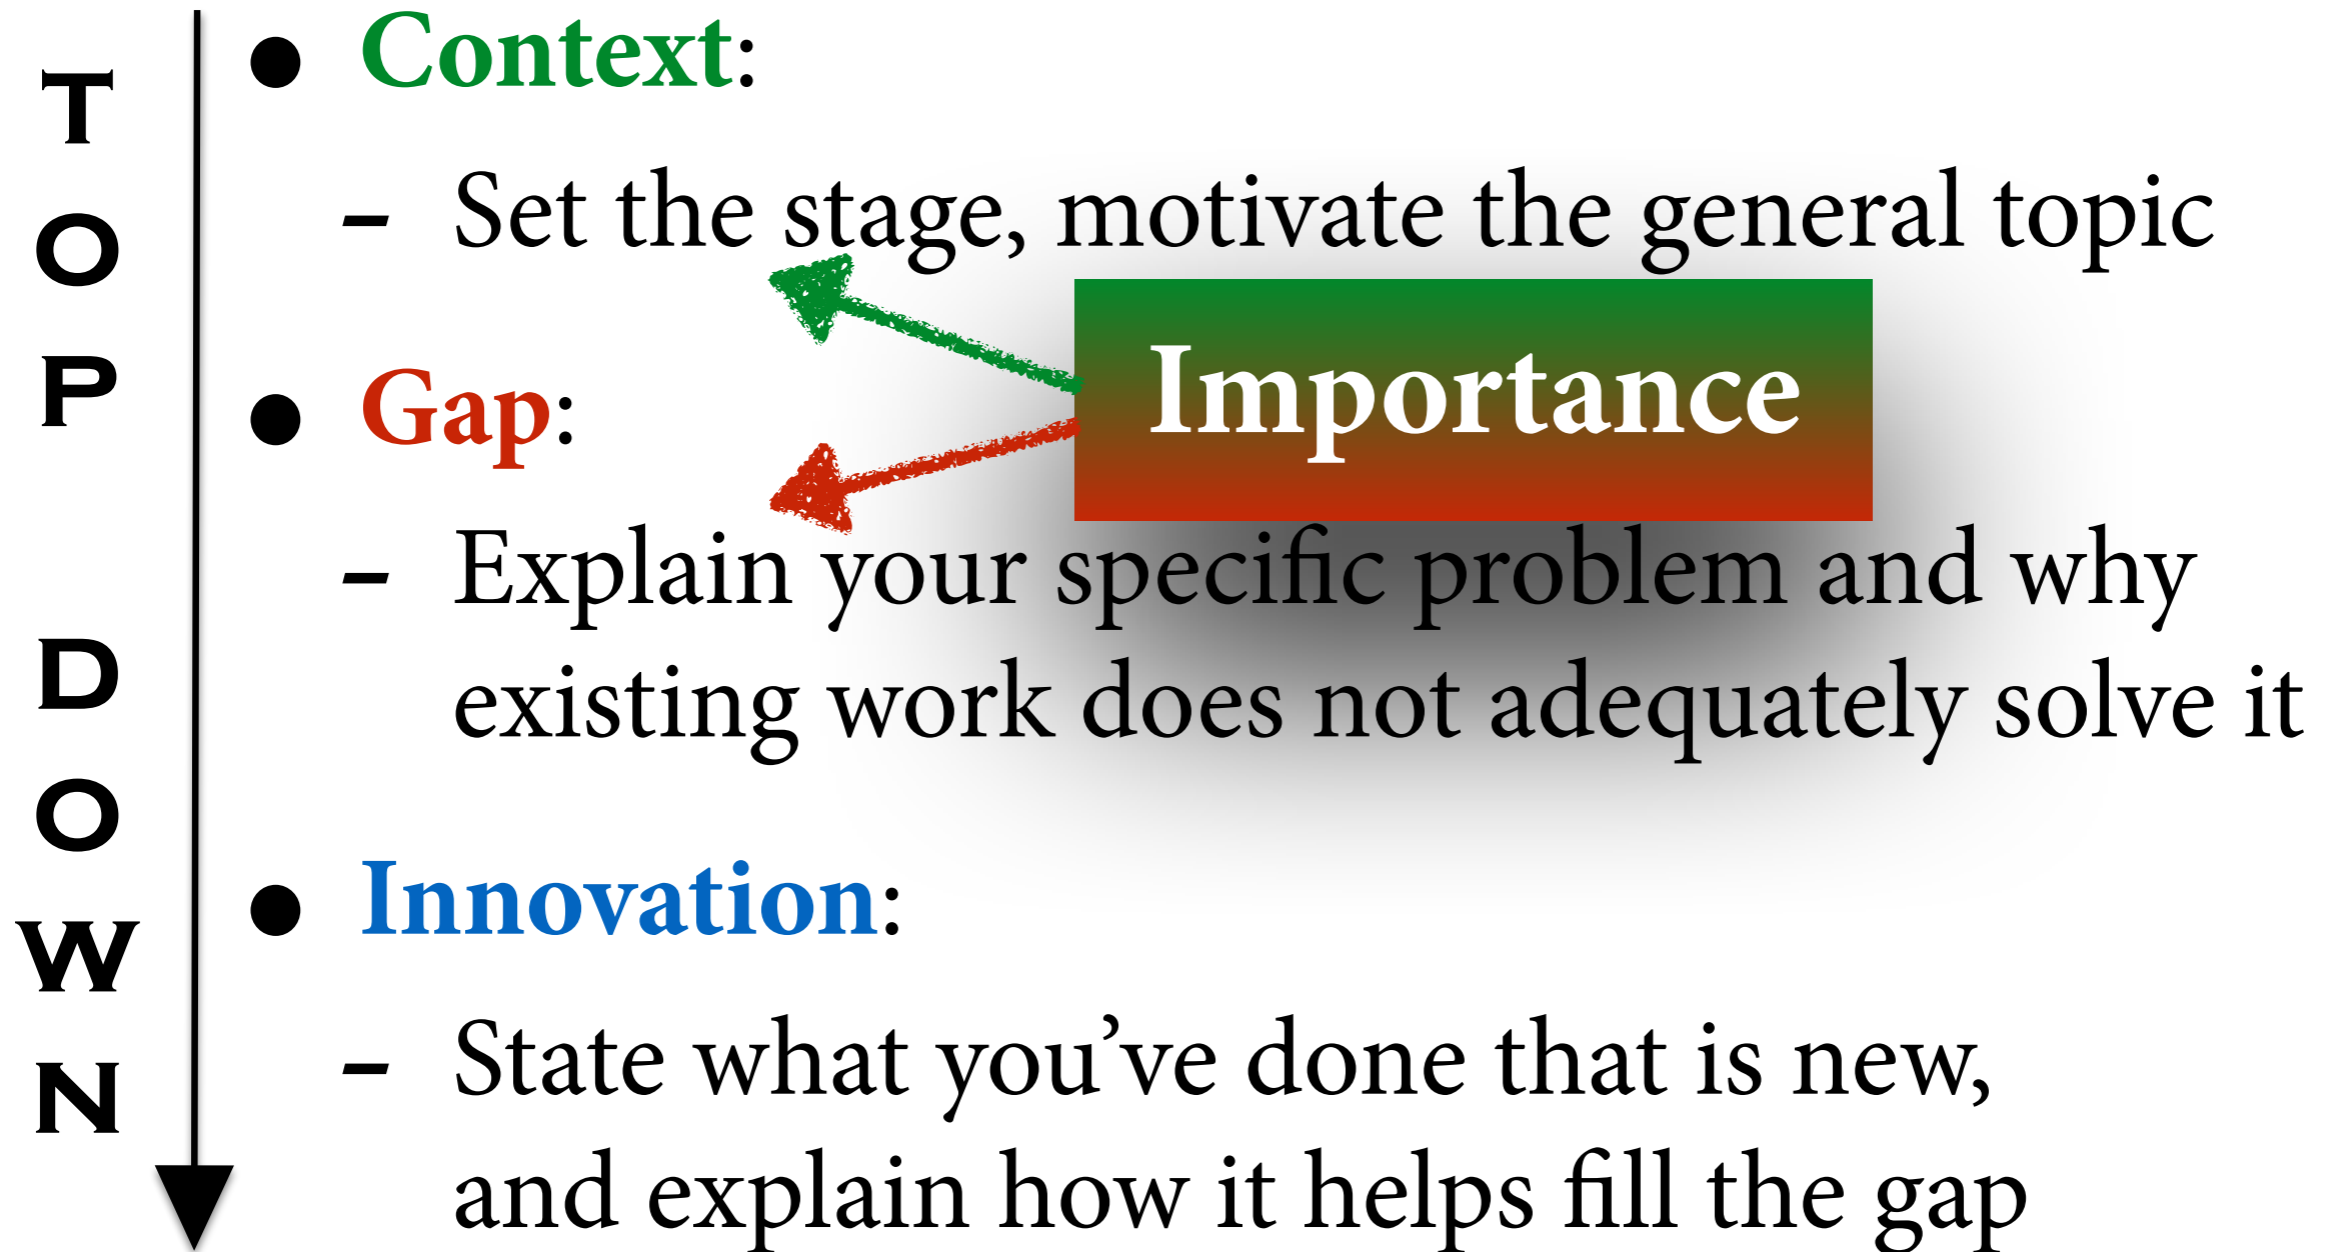
The CGI model for an abstract/intro

- **Context:**
 - Set the stage, motivate the general topic
- **Gap:**
 - Explain your specific problem and why existing work does not adequately solve it
- **Innovation:**
 - State what you've done that is new, and explain how it helps fill the gap

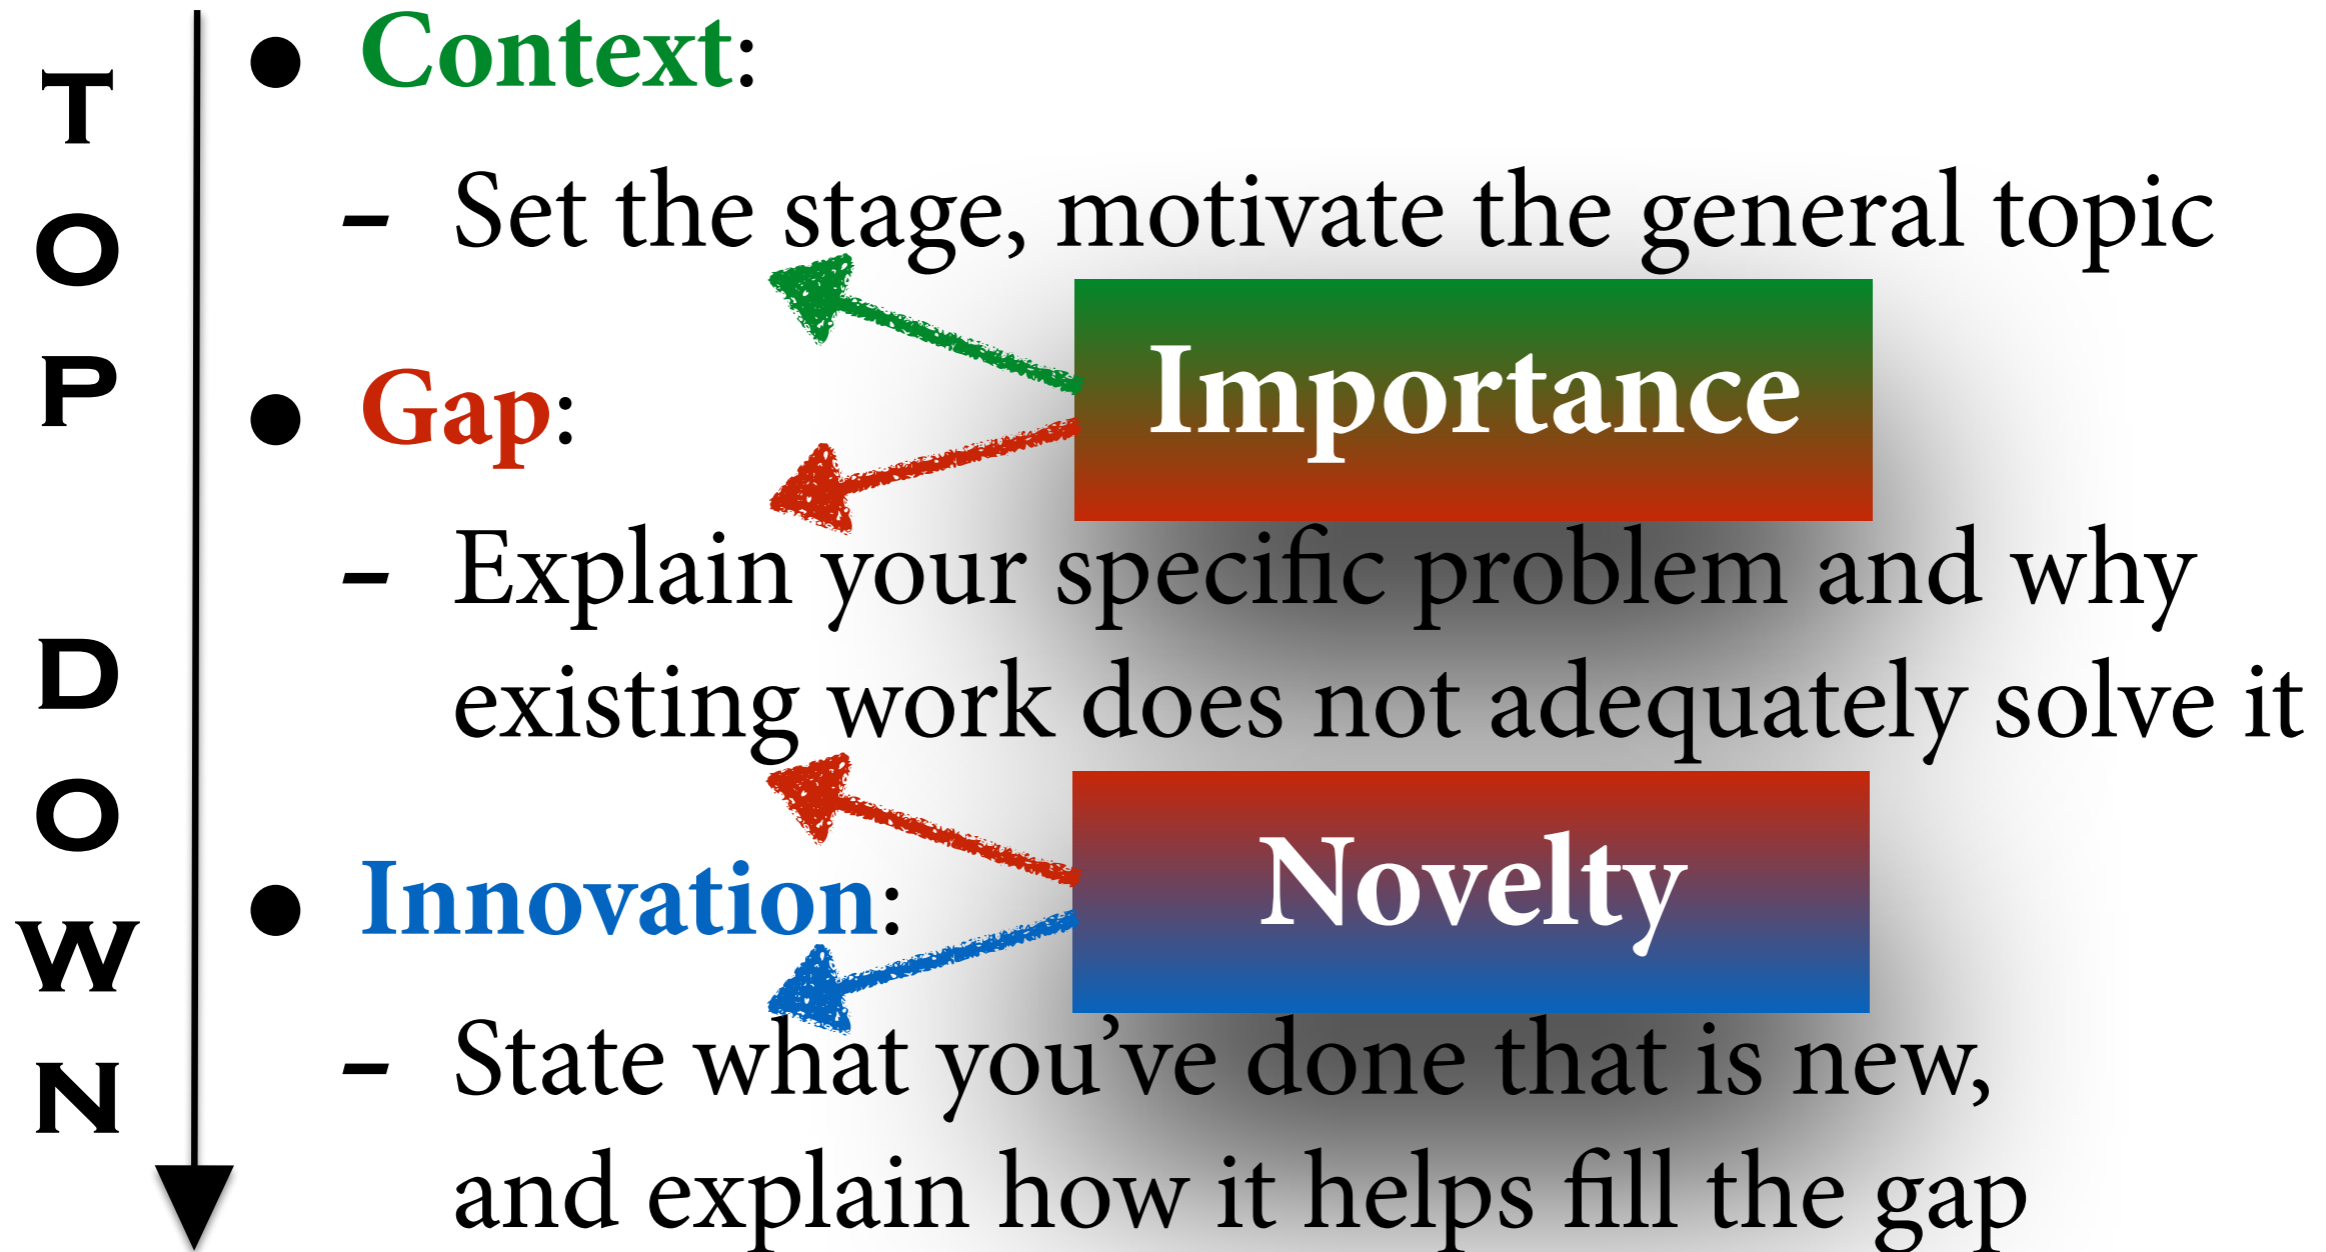
The CGI model for an abstract/intro

- T
O
P**
**D
O
W
N**
- 
- **Context:**
 - Set the stage, motivate the general topic
 - **Gap:**
 - Explain your specific problem and why existing work does not adequately solve it
 - **Innovation:**
 - State what you've done that is new, and explain how it helps fill the gap

The CGI model for an abstract/intro



The CGI model for an abstract/intro



An abstract for this talk

Context

Learning to write well is an essential part of becoming a successful researcher.

Gap

Learning to write well is an essential part of becoming a successful researcher. Unfortunately, many researchers find it very hard to write well because they do not know how to view their text from the perspective of the reader.

Innovation

Learning to write well is an essential part of becoming a successful researcher. Unfortunately, many researchers find it very hard to write well because they do not know how to view their text from the perspective of the reader. In this talk, we present a simple set of principles for good writing, based on an understanding of how readers process information. Unlike such platitudes as "Be clear" or "Omit needless words", our principles are *constructive*: one can easily check whether a piece of text satisfies them, and if it does not, the principles suggest concrete ways to improve it.

Introduction

- Like an expanded version of the abstract
- Alternative approach (SPJ): Eliminate **C**ontext
 - Start with a concrete example, e.g. “Consider this Haskell code...”
 - If this works, it can be effective, but I find it often doesn’t work
 - It assumes reader already knows context



A structure that works

- Abstract (1-2 paragraphs, 1000 readers)
- Intro (2-4 pages, 100 readers)
- **Key ideas** (4-6 pages, 50 readers)
- Technical meat (8-12 pages, 5 readers)
- Related work (1-3 pages, 100 readers)

“Key ideas” section



- Use **concrete illustrative examples** and high-level intuition
- Do **not** have to show the general solution (that's what the technical section is for)

Why have a “key ideas” section at all?



1. Forces you to have a **takeaway**, i.e. something **interesting**!
2. Many readers only care about the takeaway, not the technical details
3. For those who want the technical details, the key ideas are still useful as “scaffolding”

A structure that works

- Abstract (1-2 paragraphs, 1000 readers)
- Intro (2-4 pages, 100 readers)
- Key ideas (4-6 pages, 50 readers)
- Technical meat (8-12 pages, 5 readers)
- **Related work** (1-3 pages, 100 readers)

Related work

1. **It goes at the end** of the paper.
 - You can only properly compare to related work once you've explained your own.
2. **Give real comparisons**, not a “laundry list”!
 - Explain in detail how your work fills the **G**ap in a way that related work doesn't.

Summary of principles

- Flow via “old to new”
- Coherence via “one paragraph, one point”
- Name your baby, just in time, short subjects
- CGI model for abstract/intro
- Scaffold presentation with “key ideas” section
- Detailed related work section goes at the end



That's all Folks!