

Today I learned.....

The 2026 FLOSS Research Roundup

Dr. Kaylea Champion – kaylea@uw.edu – [@kaylea@social.coop](https://social.coop)

<https://www.kayleachampion.com>

<https://communitydata.science>

University of Washington Bothell

SeaGL—October 24, 2026



What can research tell us about building, leading, and sustaining open source?

- Getting oriented – me and why I'm here

What can research tell us about building, leading, and sustaining open source?

- Getting oriented – me and why I'm here
- What's new? What have researchers learned lately?

A little introduction

- Dr. Kaylea Champion

A little introduction

- Dr. Kaylea Champion
- Assistant Professor in Computing and Software Systems

A little introduction

- Dr. Kaylea Champion
- Assistant Professor in Computing and Software Systems
- PhD from UW Seattle

A little introduction

- Dr. Kaylea Champion
- Assistant Professor in Computing and Software Systems
- PhD from UW Seattle
- I study free software communities...

A little introduction

- Dr. Kaylea Champion
- Assistant Professor in Computing and Software Systems
- PhD from UW Seattle
- I study free software communities...
...including what works, what fails, and how we can improve.

A little introduction

- Dr. Kaylea Champion
- Assistant Professor in Computing and Software Systems
- PhD from UW Seattle
- I study free software communities...
...including what works, what fails, and how we can improve.
- Official tagline: My work sits at the busy intersection of empirical software engineering, AI, and cybersecurity

A little introduction

- Dr. Kaylea Champion
- Assistant Professor in Computing and Software Systems
- PhD from UW Seattle
- I study free software communities...
...including what works, what fails, and how we can improve.
- Official tagline: My work sits at the busy intersection of empirical software engineering, AI, and cybersecurity
- Background in industry, working my way up from tech support and sysadmin...

A little introduction

- Dr. Kaylea Champion
- Assistant Professor in Computing and Software Systems
- PhD from UW Seattle
- I study free software communities...
...including what works, what fails, and how we can improve.
- Official tagline: My work sits at the busy intersection of empirical software engineering, AI, and cybersecurity
- Background in industry, working my way up from tech support and sysadmin...
...to an IT director leading teams of teams

A little introduction

- Dr. Kaylea Champion
- Assistant Professor in Computing and Software Systems
- PhD from UW Seattle
- I study free software communities...
...including what works, what fails, and how we can improve.
- Official tagline: My work sits at the busy intersection of empirical software engineering, AI, and cybersecurity
- Background in industry, working my way up from tech support and sysadmin...
...to an IT director leading teams of teams
First Linux: Slackware 1994

Why I'm here

- I believe in free software as a force for good in the world

Why I'm here

- I believe in free software as a force for good in the world
- My role in the free software community is as a researcher



Why I'm here

- I believe in free software as a force for good in the world
- My role in the free software community is as a researcher
- My work is enriched by my understanding of how people are getting stuff done day to day

Why I'm here

- I believe in free software as a force for good in the world
- My role in the free software community is as a researcher
- My work is enriched by my understanding of how people are getting stuff done day to day
- Your taxes helped to fund my education and now help to pay my salary, so thanks for that ;)





The world runs on open source software – including academic research.

How do we use open source? Both subject and object.

The world runs on open source software – including academic research.

How do we use open source? Both subject and object.

- as a tool (like everyone else)

The world runs on open source software – including academic research.

How do we use open source? Both subject and object.

- as a tool (like everyone else)
- as a setting and data source

The world runs on open source software – including academic research.

How do we use open source? Both subject and object.

- as a tool (like everyone else)
- as a setting and data source
- as important in its own right – the specific things we make matter

The world runs on open source software – including academic research.

How do we use open source? Both subject and object.

- as a tool (like everyone else)
- as a setting and data source
- as important in its own right – the specific things we make matter
- as part of something larger – society, technology, the world we build together

- Research owes a lot to open source!

- Research owes a lot to open source!
- I read literally hundreds of papers about open source and empirical software engineering each year!

- Research owes a lot to open source!
- I read literally hundreds of papers about open source and empirical software engineering each year!
- Putting this research into practice can be a challenge.

- Research owes a lot to open source!
- I read literally hundreds of papers about open source and empirical software engineering each year!
- Putting this research into practice can be a challenge.
- This talk is about translating just a few of those papers into terms I hope are helpful to you.

- Research owes a lot to open source!
- I read literally hundreds of papers about open source and empirical software engineering each year!
- Putting this research into practice can be a challenge.
- This talk is about translating just a few of those papers into terms I hope are helpful to you.
- You being here helps to bridge the gap between research and practice! Thank you!

- Research owes a lot to open source!
- I read literally hundreds of papers about open source and empirical software engineering each year!
- Putting this research into practice can be a challenge.
- This talk is about translating just a few of those papers into terms I hope are helpful to you.
- You being here helps to bridge the gap between research and practice! Thank you!
- Note: tons of citations – all at the end. All links to free / open downloads!

What can research tell us about building, leading, and sustaining open source?

- ~~Getting oriented — me and why I'm here~~
- What's new? What have researchers learned lately?

Key topics:

- Things that are AI

Key topics:

- Things that are AI
- Things that are not AI

Key topics:

- Insights on working together

Key topics:

- Insights on working together
- News from cybersecurity

Key topics:

- Insights on working together
- News from cybersecurity
- Challenges around maintaining and sustaining

Key topics:

- Insights on working together
- News from cybersecurity
- Challenges around maintaining and sustaining
- Organizing and governing

Key topics:

- Insights on working together
- News from cybersecurity
- Challenges around maintaining and sustaining
- Organizing and governing
- Graceful transitions

Key topics:

- Insights on working together
- News from cybersecurity
- Challenges around maintaining and sustaining
- Organizing and governing
- Graceful transitions
- ...but also AI. So much AI.

Working together: Newcomers

Students as newcomers (Casseo et al., 2026) – they find:

- Large scale potential from students

Working together: Newcomers

Students as newcomers (Casseo et al., 2026) – they find:

- Large scale potential from students
- Student-reported barriers:

Working together: Newcomers

Students as newcomers (Casseo et al., 2026) – they find:

- Large scale potential from students
- Student-reported barriers:
 - Lack of maintainer response

Working together: Newcomers

Students as newcomers (Casseo et al., 2026) – they find:

- Large scale potential from students
- Student-reported barriers:
 - Lack of maintainer response
 - Code complexity

Working together: Newcomers

Students as newcomers (Casseo et al., 2026) – they find:

- Large scale potential from students
- Student-reported barriers:
 - Lack of maintainer response
 - Code complexity
 - Lack of documentation

Working together: Newcomers

Students as newcomers (Casseo et al., 2026) – they find:

- Large scale potential from students
- Student-reported barriers:
 - Lack of maintainer response
 - Code complexity
 - Lack of documentation
- Having taught an open source class, my observation: Students often do not know their own skill level.

Working together: Newcomers

Students as newcomers (Casseo et al., 2026) – they find:

- Large scale potential from students
- Student-reported barriers:
 - Lack of maintainer response
 - Code complexity
 - Lack of documentation
- Having taught an open source class, my observation: Students often do not know their own skill level.
- With AI? My observation: “good for newcomers” tagging seems like a casualty of AI.

Working together: Newcomers

Students as newcomers (Cassee et al., 2026) – they find:

- Large scale potential from students
- Student-reported barriers:
 - Lack of maintainer response
 - Code complexity
 - Lack of documentation
- Having taught an open source class, my observation: Students often do not know their own skill level.
- With AI? My observation: “good for newcomers” tagging seems like a casualty of AI.
- If you want to interact with new humans and consider funneling them to a place your community hangs out, with some onboarding bots and automation, and make the access route very clear.

Working together: Newcomers

Students as newcomers (Casseo et al., 2026) – they find:

- Large scale potential from students
- Student-reported barriers:
 - Lack of maintainer response
 - Code complexity
 - Lack of documentation
- Having taught an open source class, my observation: Students often do not know their own skill level.
- With AI? My observation: “good for newcomers” tagging seems like a casualty of AI.
- If you want to interact with new humans and consider funneling them to a place your community hangs out, with some onboarding bots and automation, and make the access route very clear.
- If you’re interested in how the students experienced an open source class and how to improve, consider their talk at 4:30 today!

Peng et al. (2026) study Fedora, Debian, Ubuntu, openEuler – they find:

- Distributions are doing duplicative work:

Peng et al. (2026) study Fedora, Debian, Ubuntu, openEuler – they find:

- Distributions are doing duplicative work:
 - Large distributions (even those in different families, e.g. Debian, Fedora) often distribute the same packages

Peng et al. (2026) study Fedora, Debian, Ubuntu, openEuler – they find:

- Distributions are doing duplicative work:
 - Large distributions (even those in different families, e.g. Debian, Fedora) often distribute the same packages
 - Distributions sometimes share patches, but these can take a long time to propagate – a median of 530 days for Fedora to Debian!

Peng et al. (2026) study Fedora, Debian, Ubuntu, openEuler – they find:

- Distributions are doing duplicative work:
 - Large distributions (even those in different families, e.g. Debian, Fedora) often distribute the same packages
 - Distributions sometimes share patches, but these can take a long time to propagate – a median of 530 days for Fedora to Debian!
 - More patch sharing is possible (they did it, with just a `git apply --check` + PR).

Peng et al. (2026) study Fedora, Debian, Ubuntu, openEuler – they find:

- Distributions are doing duplicative work:
 - Large distributions (even those in different families, e.g. Debian, Fedora) often distribute the same packages
 - Distributions sometimes share patches, but these can take a long time to propagate – a median of 530 days for Fedora to Debian!
 - More patch sharing is possible (they did it, with just a `git apply --check` + PR).
 - Maintainers may be unaware of reusable patches

Peng et al. (2026) study Fedora, Debian, Ubuntu, openEuler – they find:

- Distributions are doing duplicative work:
 - Large distributions (even those in different families, e.g. Debian, Fedora) often distribute the same packages
 - Distributions sometimes share patches, but these can take a long time to propagate – a median of 530 days for Fedora to Debian!
 - More patch sharing is possible (they did it, with just a `git apply --check` + PR).
 - Maintainers may be unaware of reusable patches
- My take: Opportunity to streamline across our broader community

Yu et al. (2026) examine 1.6M developers over 14 years (World Of Code dataset):

- Women tend to work with other women at rates higher than their prevalence**

**Note: Gender inferred by name, greater than .7 probability cutoff (otherwise, *unknown*).

Yu et al. (2026) examine 1.6M developers over 14 years (World Of Code dataset):

- Women tend to work with other women at rates higher than their prevalence**
- Men tend to work with women at rates higher than their prevalence too!

**Note: Gender inferred by name, greater than .7 probability cutoff (otherwise, *unknown*).

Yu et al. (2026) examine 1.6M developers over 14 years (World Of Code dataset):

- Women tend to work with other women at rates higher than their prevalence**
- Men tend to work with women at rates higher than their prevalence too!
- Women who collaborate with other women are 21% less likely to drop out

**Note: Gender inferred by name, greater than .7 probability cutoff (otherwise, *unknown*).

Yu et al. (2026) examine 1.6M developers over 14 years (World Of Code dataset):

- Women tend to work with other women at rates higher than their prevalence**
- Men tend to work with women at rates higher than their prevalence too!
- Women who collaborate with other women are 21% less likely to drop out
- “Diversity as Infrastructure, Not Decoration”

**Note: Gender inferred by name, greater than .7 probability cutoff (otherwise, *unknown*).

- PHP as the latest example

- PHP as the latest example
- Continuing challenges around AI slop, but substantial improvements in bug bounty triage and AI security scanning tools

- PHP as the latest example
- Continuing challenges around AI slop, but substantial improvements in bug bounty triage and AI security scanning tools
- Discovery, but make it (even) faster

- PHP as the latest example
- Continuing challenges around AI slop, but substantial improvements in bug bounty triage and AI security scanning tools
- Discovery, but make it (even) faster
- Repair (possibly faster)

- PHP as the latest example
- Continuing challenges around AI slop, but substantial improvements in bug bounty triage and AI security scanning tools
- Discovery, but make it (even) faster
- Repair (possibly faster)
- Validation...handling of disclosure....patching...

- PHP as the latest example
- Continuing challenges around AI slop, but substantial improvements in bug bounty triage and AI security scanning tools
- Discovery, but make it (even) faster
- Repair (possibly faster)
- Validation...handling of disclosure....patching...

These steps were already strained and backlogged (my work in progress)

- PHP as the latest example
- Continuing challenges around AI slop, but substantial improvements in bug bounty triage and AI security scanning tools
- Discovery, but make it (even) faster
- Repair (possibly faster)
- Validation...handling of disclosure....patching...
These steps were already strained and backlogged (my work in progress)
- Useful metaphor: A dam with an outlet and a spillway, in a heavy rain.



[The Oroville embankment dam. By California Department of Water Resources. Public Domain, via Wikipedia.]



[Claerwen Dam Overflowing. Charlie Marshall. via flickr, CC BY-NC-SA 4.0]

Better engineering, maintaining, and sustaining

Shen et al. (2026) on forks and PRs

- An abstracted view of a workflow we might take for granted – the fork and the pull request

Better engineering, maintaining, and sustaining

Shen et al. (2026) on forks and PRs

- An abstracted view of a workflow we might take for granted – the fork and the pull request
- Examine the extent to which commits end up merged across forks

Better engineering, maintaining, and sustaining

Shen et al. (2026) on forks and PRs

- An abstracted view of a workflow we might take for granted – the fork and the pull request
- Examine the extent to which commits end up merged across forks
 - New metric: “convergence entropy”

Better engineering, maintaining, and sustaining

Shen et al. (2026) on forks and PRs

- An abstracted view of a workflow we might take for granted – the fork and the pull request
- Examine the extent to which commits end up merged across forks
 - New metric: “convergence entropy”
 - The extent to which a project’s forks are converging or diverging from one another

Better engineering, maintaining, and sustaining

Shen et al. (2026) on forks and PRs

- An abstracted view of a workflow we might take for granted – the fork and the pull request
- Examine the extent to which commits end up merged across forks
 - New metric: “convergence entropy”
 - The extent to which a project’s forks are converging or diverging from one another
 - Across a slew of forks – are we all headed in the same direction?

Better engineering, maintaining, and sustaining

Shen et al. (2026) on forks and PRs

- An abstracted view of a workflow we might take for granted – the fork and the pull request
- Examine the extent to which commits end up merged across forks
 - New metric: “convergence entropy”
 - The extent to which a project’s forks are converging or diverging from one another
 - Across a slew of forks – are we all headed in the same direction?
- In non-company projects: Higher entropy associated with slower issue resolution

Better engineering, maintaining, and sustaining

Shen et al. (2026) on forks and PRs

- An abstracted view of a workflow we might take for granted – the fork and the pull request
- Examine the extent to which commits end up merged across forks
 - New metric: “convergence entropy”
 - The extent to which a project’s forks are converging or diverging from one another
 - Across a slew of forks – are we all headed in the same direction?
- In non-company projects: Higher entropy associated with slower issue resolution
- In company projects: Higher entropy is associated with lower commit volume

Better engineering, maintaining, and sustaining

Shen et al. (2026) on forks and PRs

- An abstracted view of a workflow we might take for granted – the fork and the pull request
- Examine the extent to which commits end up merged across forks
 - New metric: “convergence entropy”
 - The extent to which a project’s forks are converging or diverging from one another
 - Across a slew of forks – are we all headed in the same direction?
- In non-company projects: Higher entropy associated with slower issue resolution
- In company projects: Higher entropy is associated with lower commit volume
- Benefits of convergence are weaker for large / mature projects

Better engineering, maintaining, and sustaining

Shen et al. (2026) on forks and PRs

- An abstracted view of a workflow we might take for granted – the fork and the pull request
- Examine the extent to which commits end up merged across forks
 - New metric: “convergence entropy”
 - The extent to which a project’s forks are converging or diverging from one another
 - Across a slew of forks – are we all headed in the same direction?
- In non-company projects: Higher entropy associated with slower issue resolution
- In company projects: Higher entropy is associated with lower commit volume
- Benefits of convergence are weaker for large / mature projects
- My takeaway: worth thinking about how much you want to try to cultivate convergence versus entropy

Better engineering, maintaining, and sustaining

Shen et al. (2026) on forks and PRs

- An abstracted view of a workflow we might take for granted – the fork and the pull request
- Examine the extent to which commits end up merged across forks
 - New metric: “convergence entropy”
 - The extent to which a project’s forks are converging or diverging from one another
 - Across a slew of forks – are we all headed in the same direction?
- In non-company projects: Higher entropy associated with slower issue resolution
- In company projects: Higher entropy is associated with lower commit volume
- Benefits of convergence are weaker for large / mature projects
- My takeaway: worth thinking about how much you want to try to cultivate convergence versus entropy
- Time to harvest from forks?

Organizing and Governing: Advice and Practice

From Gaughan, Shaw & Gergle (2026)

- Can a community FLOSS project operate in partnership with a large organization without being squashed by it?

Organizing and Governing: Advice and Practice

From Gaughan, Shaw & Gergle (2026)

- Can a community FLOSS project operate in partnership with a large organization without being squashed by it?
 - A study of Mediawiki and Wikimedia says yes

Organizing and Governing: Advice and Practice

From Gaughan, Shaw & Gergle (2026)

- Can a community FLOSS project operate in partnership with a large organization without being squashed by it?
 - A study of Mediawiki and Wikimedia says yes
 - Comparison of workload (higher for the paid staff in the partner org) and how decisionmaking and discussions happen (equal across community and partner)

Organizing and Governing: Advice and Practice

From Gaughan, Shaw & Gergle (2026)

- Can a community FLOSS project operate in partnership with a large organization without being squashed by it?
 - A study of Mediawiki and Wikimedia says yes
 - Comparison of workload (higher for the paid staff in the partner org) and how decisionmaking and discussions happen (equal across community and partner)
- My Takeaway:

Organizing and Governing: Advice and Practice

From Gaughan, Shaw & Gergle (2026)

- Can a community FLOSS project operate in partnership with a large organization without being squashed by it?
 - A study of Mediawiki and Wikimedia says yes
 - Comparison of workload (higher for the paid staff in the partner org) and how decisionmaking and discussions happen (equal across community and partner)
- My Takeaway:
 - In FLOSS, we're accustomed to commits and status going hand in hand.

Organizing and Governing: Advice and Practice

From Gaughan, Shaw & Gergle (2026)

- Can a community FLOSS project operate in partnership with a large organization without being squashed by it?
 - A study of Mediawiki and Wikimedia says yes
 - Comparison of workload (higher for the paid staff in the partner org) and how decisionmaking and discussions happen (equal across community and partner)
- My Takeaway:
 - In FLOSS, we're accustomed to commits and status going hand in hand.
 - But lately we've also been trying to remind ourselves that commit count isn't everything (documentation, support, design)

Organizing and Governing: Advice and Practice

From Gaughan, Shaw & Gergle (2026)

- Can a community FLOSS project operate in partnership with a large organization without being squashed by it?
 - A study of Mediawiki and Wikimedia says yes
 - Comparison of workload (higher for the paid staff in the partner org) and how decisionmaking and discussions happen (equal across community and partner)
- My Takeaway:
 - In FLOSS, we're accustomed to commits and status going hand in hand.
 - But lately we've also been trying to remind ourselves that commit count isn't everything (documentation, support, design)
 - Time to listen to ourselves a bit better!

Organizing and Governing: Advice and Practice

From Gaughan, Shaw & Gergle (2026)

- Can a community FLOSS project operate in partnership with a large organization without being squashed by it?
 - A study of Mediawiki and Wikimedia says yes
 - Comparison of workload (higher for the paid staff in the partner org) and how decisionmaking and discussions happen (equal across community and partner)
- My Takeaway:
 - In FLOSS, we're accustomed to commits and status going hand in hand.
 - But lately we've also been trying to remind ourselves that commit count isn't everything (documentation, support, design)
 - Time to listen to ourselves a bit better!
 - Maybe even labor isn't everything

Organizing and Governing: Advice and Practice

From Gaughan, Shaw & Gergle (2026)

- Can a community FLOSS project operate in partnership with a large organization without being squashed by it?
 - A study of Mediawiki and Wikimedia says yes
 - Comparison of workload (higher for the paid staff in the partner org) and how decisionmaking and discussions happen (equal across community and partner)
- My Takeaway:
 - In FLOSS, we're accustomed to commits and status going hand in hand.
 - But lately we've also been trying to remind ourselves that commit count isn't everything (documentation, support, design)
 - Time to listen to ourselves a bit better!
 - Maybe even labor isn't everything
 - Openness and "the long tail" contributes diverse perspectives, lends a critical eye, broadens use cases, shifts priorities

Organizing and Governing: Advice and Practice

From Gaughan, Shaw & Gergle (2026)

- Can a community FLOSS project operate in partnership with a large organization without being squashed by it?
 - A study of Mediawiki and Wikimedia says yes
 - Comparison of workload (higher for the paid staff in the partner org) and how decisionmaking and discussions happen (equal across community and partner)
- My Takeaway:
 - In FLOSS, we're accustomed to commits and status going hand in hand.
 - But lately we've also been trying to remind ourselves that commit count isn't everything (documentation, support, design)
 - Time to listen to ourselves a bit better!
 - Maybe even labor isn't everything
 - Openness and "the long tail" contributes diverse perspectives, lends a critical eye, broadens use cases, shifts priorities
 - Community engagement sustains both legitimacy and quality in FLOSS

Organizing and Governing: Advice and Practice

From Gaughan, Shaw & Gergle (2026)

- Can a community FLOSS project operate in partnership with a large organization without being squashed by it?
 - A study of Mediawiki and Wikimedia says yes
 - Comparison of workload (higher for the paid staff in the partner org) and how decisionmaking and discussions happen (equal across community and partner)
- My Takeaway:
 - In FLOSS, we're accustomed to commits and status going hand in hand.
 - But lately we've also been trying to remind ourselves that commit count isn't everything (documentation, support, design)
 - Time to listen to ourselves a bit better!
 - Maybe even labor isn't everything
 - Openness and "the long tail" contributes diverse perspectives, lends a critical eye, broadens use cases, shifts priorities
 - Community engagement sustains both legitimacy and quality in FLOSS
 - Attitude of mutual respect is key

The impact of codes of conduct (Sun et al. (2026))

- We have a lot of data now – we can look retrospectively and draw large samples

The impact of codes of conduct (Sun et al. (2026))

- We have a lot of data now – we can look retrospectively and draw large samples
- Adoption is a signal of welcoming – projects with a CoC draw newcomers short- and long-term

The impact of codes of conduct (Sun et al. (2026))

- We have a lot of data now – we can look retrospectively and draw large samples
- Adoption is a signal of welcoming – projects with a CoC draw newcomers short- and long-term
- At the point of adoption, may lead to short-term disengagement from existing contributors

The impact of codes of conduct (Sun et al. (2026))

- We have a lot of data now – we can look retrospectively and draw large samples
- Adoption is a signal of welcoming – projects with a CoC draw newcomers short- and long-term
- At the point of adoption, may lead to short-term disengagement from existing contributors
- Their recommendations:

The impact of codes of conduct (Sun et al. (2026))

- We have a lot of data now – we can look retrospectively and draw large samples
- Adoption is a signal of welcoming – projects with a CoC draw newcomers short- and long-term
- At the point of adoption, may lead to short-term disengagement from existing contributors
- Their recommendations:
 - “Adopt CoCs as a meaningful signal, then treat them as living documents”

The impact of codes of conduct (Sun et al. (2026))

- We have a lot of data now – we can look retrospectively and draw large samples
- Adoption is a signal of welcoming – projects with a CoC draw newcomers short- and long-term
- At the point of adoption, may lead to short-term disengagement from existing contributors
- Their recommendations:
 - “Adopt CoCs as a meaningful signal, then treat them as living documents”
 - “Leverage diverse update strategies aligned with community needs.”

The impact of codes of conduct (Sun et al. (2026))

- We have a lot of data now – we can look retrospectively and draw large samples
- Adoption is a signal of welcoming – projects with a CoC draw newcomers short- and long-term
- At the point of adoption, may lead to short-term disengagement from existing contributors
- Their recommendations:
 - “Adopt CoCs as a meaningful signal, then treat them as living documents”
 - “Leverage diverse update strategies aligned with community needs.”
 - “Promote transparency and inclusiveness in CoC governance.”

The impact of codes of conduct (Sun et al. (2026))

- We have a lot of data now – we can look retrospectively and draw large samples
- Adoption is a signal of welcoming – projects with a CoC draw newcomers short- and long-term
- At the point of adoption, may lead to short-term disengagement from existing contributors
- Their recommendations:
 - “Adopt CoCs as a meaningful signal, then treat them as living documents”
 - “Leverage diverse update strategies aligned with community needs.”
 - “Promote transparency and inclusiveness in CoC governance.”
 - “Track and assess CoC influence systematically.” (ex: tools in the CHAOSS project)

Graceful Transitions

Miller et al. (2026) continuing a thread of work on dependency Abandonment –

- They interviewed open source maintainers to understand when it matters that a dependency is abandoned

Miller et al. (2026) continuing a thread of work on dependency Abandonment –

- They interviewed open source maintainers to understand when it matters that a dependency is abandoned
- 4 concerns:

Miller et al. (2026) continuing a thread of work on dependency Abandonment –

- They interviewed open source maintainers to understand when it matters that a dependency is abandoned
- 4 concerns:
 - Importance of functionality

Miller et al. (2026) continuing a thread of work on dependency Abandonment –

- They interviewed open source maintainers to understand when it matters that a dependency is abandoned
- 4 concerns:
 - Importance of functionality
 - Depth of integration

Miller et al. (2026) continuing a thread of work on dependency Abandonment –

- They interviewed open source maintainers to understand when it matters that a dependency is abandoned
- 4 concerns:
 - Importance of functionality
 - Depth of integration
 - Availability of alternatives

Miller et al. (2026) continuing a thread of work on dependency Abandonment –

- They interviewed open source maintainers to understand when it matters that a dependency is abandoned
- 4 concerns:
 - Importance of functionality
 - Depth of integration
 - Availability of alternatives
 - External environmental pressure (will they be forced to act?)

Graceful Transitions

Miller et al. (2026) continuing a thread of work on dependency Abandonment –

- They interviewed open source maintainers to understand when it matters that a dependency is abandoned
- 4 concerns:
 - Importance of functionality
 - Depth of integration
 - Availability of alternatives
 - External environmental pressure (will they be forced to act?)
- Used these factors as predictors in a classifier as to whether a given abandonment will matter to maintainers, with good success

Graceful Transitions

Miller et al. (2026) continuing a thread of work on dependency Abandonment –

- They interviewed open source maintainers to understand when it matters that a dependency is abandoned
- 4 concerns:
 - Importance of functionality
 - Depth of integration
 - Availability of alternatives
 - External environmental pressure (will they be forced to act?)
- Used these factors as predictors in a classifier as to whether a given abandonment will matter to maintainers, with good success
- Takeaways:

Graceful Transitions

Miller et al. (2026) continuing a thread of work on dependency Abandonment –

- They interviewed open source maintainers to understand when it matters that a dependency is abandoned
- 4 concerns:
 - Importance of functionality
 - Depth of integration
 - Availability of alternatives
 - External environmental pressure (will they be forced to act?)
- Used these factors as predictors in a classifier as to whether a given abandonment will matter to maintainers, with good success
- Takeaways:
 - organizations can adopt tools and strategies like this to prioritize work in complex environments

Graceful Transitions

Miller et al. (2026) continuing a thread of work on dependency Abandonment –

- They interviewed open source maintainers to understand when it matters that a dependency is abandoned
- 4 concerns:
 - Importance of functionality
 - Depth of integration
 - Availability of alternatives
 - External environmental pressure (will they be forced to act?)
- Used these factors as predictors in a classifier as to whether a given abandonment will matter to maintainers, with good success
- Takeaways:
 - organizations can adopt tools and strategies like this to prioritize work in complex environments
 - ecosystems might build their own versions of this prediction tool

What was my take in 2025?

- Overall record for AI at improving efficiency is mixed

What was my take in 2025?

- Overall record for AI at improving efficiency is mixed
- Few companies reporting profit

What was my take in 2025?

- Overall record for AI at improving efficiency is mixed
- Few companies reporting profit
- Hype

What was my take in 2025?

- Overall record for AI at improving efficiency is mixed
- Few companies reporting profit
- Hype
- Concerns about impact on our brains

What was my take in 2025?

- Overall record for AI at improving efficiency is mixed
- Few companies reporting profit
- Hype
- Concerns about impact on our brains
- Open models are behind, with their innovations getting harvested by commercial models, and still making progress

What was my take in 2025?

- Overall record for AI at improving efficiency is mixed
- Few companies reporting profit
- Hype
- Concerns about impact on our brains
- Open models are behind, with their innovations getting harvested by commercial models, and still making progress
- Many questions about the impact on OSS projects

What (if anything....) has changed?

- More evidence about effective practices (skill.md, SDD) and business models

What (if anything....) has changed?

- More evidence about effective practices (skill.md, SDD) and business models
- Consumer backlash about data centers

What (if anything....) has changed?

- More evidence about effective practices (skill.md, SDD) and business models
- Consumer backlash about data centers
- Projects are increasingly *deciding to decide*

What (if anything....) has changed?

- More evidence about effective practices (skill.md, SDD) and business models
- Consumer backlash about data centers
- Projects are increasingly *deciding to decide*
 - Two high-contrast examples: FreeDOS, Debian

What (if anything....) has changed?

- More evidence about effective practices (skill.md, SDD) and business models
- Consumer backlash about data centers
- Projects are increasingly *deciding to decide*
 - Two high-contrast examples: FreeDOS, Debian
 - Study results from Hora, Robbes, and Zacchiroli:

What (if anything....) has changed?

- More evidence about effective practices (skill.md, SDD) and business models
- Consumer backlash about data centers
- Projects are increasingly *deciding to decide*
 - Two high-contrast examples: FreeDOS, Debian
 - Study results from Hora, Robbes, and Zacchiroli:
 - Study of 2000 most popular GitHub repositories

What (if anything....) has changed?

- More evidence about effective practices (skill.md, SDD) and business models
- Consumer backlash about data centers
- Projects are increasingly *deciding to decide*
 - Two high-contrast examples: FreeDOS, Debian
 - Study results from Hora, Robbes, and Zacchiroli:
 - Study of 2000 most popular GitHub repositories
 - Policies on usage (83% permit) & disclosure (49% require)

What (if anything....) has changed?

- More evidence about effective practices (skill.md, SDD) and business models
- Consumer backlash about data centers
- Projects are increasingly *deciding to decide*
 - Two high-contrast examples: FreeDOS, Debian
 - Study results from Hora, Robbes, and Zacchioli:
 - Study of 2000 most popular GitHub repositories
 - Policies on usage (83% permit) & disclosure (49% require)
 - Projects are taking on countermeasures against slop

What (if anything....) has changed?

- More evidence about effective practices (skill.md, SDD) and business models
- Consumer backlash about data centers
- Projects are increasingly *deciding to decide*
 - Two high-contrast examples: FreeDOS, Debian
 - Study results from Hora, Robbes, and Zacchiroli:
 - Study of 2000 most popular GitHub repositories
 - Policies on usage (83% permit) & disclosure (49% require)
 - Projects are taking on countermeasures against slop
 - Policies are still a work in progress (half of them changed during the study period).

What (if anything....) has changed?

- Continued investigation and development of open source AI models

What (if anything....) has changed?

- Continued investigation and development of open source AI models
 - Not developing in the same way OSS developed (less collaborative, less open)
(Large-scale comparison of contributions and interactions in GitHub versus Hugging Face, Ye & Zhou 2026)

What (if anything....) has changed?

- Continued investigation and development of open source AI models
 - Not developing in the same way OSS developed (less collaborative, less open)
(Large-scale comparison of contributions and interactions in GitHub versus Hugging Face, Ye & Zhou 2026)
- Many open questions about the impact on OSS projects (Min et al., 2026)

What (if anything....) has changed?

- Continued investigation and development of open source AI models
 - Not developing in the same way OSS developed (less collaborative, less open)
(Large-scale comparison of contributions and interactions in GitHub versus Hugging Face, Ye & Zhou 2026)
- Many open questions about the impact on OSS projects (Min et al., 2026)
- My work: Change in 'slop' discourse over the last year, with divergence in views between maintainers and contributors remains (Chatterjee et al., in progress)

What (if anything....) has changed?

- Continued investigation and development of open source AI models
 - Not developing in the same way OSS developed (less collaborative, less open)
(Large-scale comparison of contributions and interactions in GitHub versus Hugging Face, Ye & Zhou 2026)
- Many open questions about the impact on OSS projects (Min et al., 2026)
- My work: Change in 'slop' discourse over the last year, with divergence in views between maintainers and contributors remains (Chatterjee et al., in progress)
- Lack of careful AI policy can accelerate declines of open content systems (like Stack Exchange, see Wu et al., 2026)

What (if anything....) has changed?

- Continued investigation and development of open source AI models
 - Not developing in the same way OSS developed (less collaborative, less open)
(Large-scale comparison of contributions and interactions in GitHub versus Hugging Face, Ye & Zhou 2026)
- Many open questions about the impact on OSS projects (Min et al., 2026)
- My work: Change in 'slop' discourse over the last year, with divergence in views between maintainers and contributors remains (Chatterjee et al., in progress)
- Lack of careful AI policy can accelerate declines of open content systems (like Stack Exchange, see Wu et al., 2026)
- Even substantially AI-savvy communities (like Wikipedia) are experiencing disruption. (Strickland, 2026)

Your call to action

- Developing our governance response to AI is vital — and for more than open source.

Your call to action

- Developing our governance response to AI is vital — and for more than open source.
 - Laws and standards are often carved out of existing distributed practices

Your call to action

- Developing our governance response to AI is vital — and for more than open source.
 - Laws and standards are often carved out of existing distributed practices
 - Consider the standards being adopted by top projects and choose one that makes sense, or customize it to meet your needs

Your call to action

- Developing our governance response to AI is vital — and for more than open source.
 - Laws and standards are often carved out of existing distributed practices
 - Consider the standards being adopted by top projects and choose one that makes sense, or customize it to meet your needs
 - Use the Creative Commons “CC Signals” project to indicate how you want your community’s data to be used

Your call to action

- Developing our governance response to AI is vital — and for more than open source.
 - Laws and standards are often carved out of existing distributed practices
 - Consider the standards being adopted by top projects and choose one that makes sense, or customize it to meet your needs
 - Use the Creative Commons “CC Signals” project to indicate how you want your community’s data to be used
- Engaging with the open source AI landscape is critical for open source software

Your call to action

- Developing our governance response to AI is vital — and for more than open source.
 - Laws and standards are often carved out of existing distributed practices
 - Consider the standards being adopted by top projects and choose one that makes sense, or customize it to meet your needs
 - Use the Creative Commons “CC Signals” project to indicate how you want your community’s data to be used
- Engaging with the open source AI landscape is critical for open source software
 - We face a tipping point with respect to intellectual property

Your call to action

- Developing our governance response to AI is vital — and for more than open source.
 - Laws and standards are often carved out of existing distributed practices
 - Consider the standards being adopted by top projects and choose one that makes sense, or customize it to meet your needs
 - Use the Creative Commons “CC Signals” project to indicate how you want your community’s data to be used
- Engaging with the open source AI landscape is critical for open source software
 - We face a tipping point with respect to intellectual property
 - Our labor is training private models instead of being used to build a commons

Your call to action

- Developing our governance response to AI is vital — and for more than open source.
 - Laws and standards are often carved out of existing distributed practices
 - Consider the standards being adopted by top projects and choose one that makes sense, or customize it to meet your needs
 - Use the Creative Commons “CC Signals” project to indicate how you want your community’s data to be used
- Engaging with the open source AI landscape is critical for open source software
 - We face a tipping point with respect to intellectual property
 - Our labor is training private models instead of being used to build a commons
 - OSS trains AI, AI writes OSS....

Today you learned....

- Insights on working together

Today you learned....

- Insights on working together
 - Respond personally to your newcomers if you want them to join you

Today you learned....

- Insights on working together
 - Respond personally to your newcomers if you want them to join you
 - “Diversity as Infrastructure, Not Decoration”

Today you learned....

- Insights on working together
 - Respond personally to your newcomers if you want them to join you
 - “Diversity as Infrastructure, Not Decoration”
- News from cybersecurity

Today you learned....

- Insights on working together
 - Respond personally to your newcomers if you want them to join you
 - “Diversity as Infrastructure, Not Decoration”
- News from cybersecurity
 - We need to streamline how updates reach broad deployment

Today you learned....

- Insights on working together
 - Respond personally to your newcomers if you want them to join you
 - “Diversity as Infrastructure, Not Decoration”
- News from cybersecurity
 - We need to streamline how updates reach broad deployment
- Better engineering, maintaining, and sustaining

Today you learned....

- Insights on working together
 - Respond personally to your newcomers if you want them to join you
 - “Diversity as Infrastructure, Not Decoration”
- News from cybersecurity
 - We need to streamline how updates reach broad deployment
- Better engineering, maintaining, and sustaining
 - Other distros are patching the same packages – could share effort

Today you learned....

- Insights on working together
 - Respond personally to your newcomers if you want them to join you
 - “Diversity as Infrastructure, Not Decoration”
- News from cybersecurity
 - We need to streamline how updates reach broad deployment
- Better engineering, maintaining, and sustaining
 - Other distros are patching the same packages – could share effort
- Governing: Advice and Practice

Today you learned....

- Insights on working together
 - Respond personally to your newcomers if you want them to join you
 - “Diversity as Infrastructure, Not Decoration”
- News from cybersecurity
 - We need to streamline how updates reach broad deployment
- Better engineering, maintaining, and sustaining
 - Other distros are patching the same packages – could share effort
- Governing: Advice and Practice
 - Wikimedia as an example of how an organization can work with a community without squishing it

Today you learned....

- Insights on working together
 - Respond personally to your newcomers if you want them to join you
 - “Diversity as Infrastructure, Not Decoration”
- News from cybersecurity
 - We need to streamline how updates reach broad deployment
- Better engineering, maintaining, and sustaining
 - Other distros are patching the same packages – could share effort
- Governing: Advice and Practice
 - Wikimedia as an example of how an organization can work with a community without squishing it
- Graceful Transitions

Today you learned....

- Insights on working together
 - Respond personally to your newcomers if you want them to join you
 - “Diversity as Infrastructure, Not Decoration”
- News from cybersecurity
 - We need to streamline how updates reach broad deployment
- Better engineering, maintaining, and sustaining
 - Other distros are patching the same packages – could share effort
- Governing: Advice and Practice
 - Wikimedia as an example of how an organization can work with a community without squishing it
- Graceful Transitions
 - Does your project have a strategy for abandoned dependencies?

Today you learned....

- Insights on working together
 - Respond personally to your newcomers if you want them to join you
 - “Diversity as Infrastructure, Not Decoration”
- News from cybersecurity
 - We need to streamline how updates reach broad deployment
- Better engineering, maintaining, and sustaining
 - Other distros are patching the same packages – could share effort
- Governing: Advice and Practice
 - Wikimedia as an example of how an organization can work with a community without squishing it
- Graceful Transitions
 - Does your project have a strategy for abandoned dependencies?
- ...and also AI. So much AI.

Today you learned....

- Insights on working together
 - Respond personally to your newcomers if you want them to join you
 - “Diversity as Infrastructure, Not Decoration”
- News from cybersecurity
 - We need to streamline how updates reach broad deployment
- Better engineering, maintaining, and sustaining
 - Other distros are patching the same packages – could share effort
- Governing: Advice and Practice
 - Wikimedia as an example of how an organization can work with a community without squishing it
- Graceful Transitions
 - Does your project have a strategy for abandoned dependencies?
- ...and also AI. So much AI.

Does your project have an AI policy?

This has been your 2026 research roundup! Thank you!

What questions do you have?

What topic(s) do you think need more research in open source?

I am seeking:


- collaborators and partners for research!
- projects for the students in my open source class!
- guest speakers for my classes on open source, software engineering, and technical management!

`kaylea@uw.edu`—`@kaylea@social.coop`

`https://kayleachampion.com`

`@communitydata@social.coop`

I gratefully acknowledge support from the Sloan Foundation through the Ford/Sloan Digital Infrastructure Initiative, Sloan Award 2018-11356 as well as the National Science Foundation (Grant IIS-2045055). Important portions of my work have been conducted using the Hyak supercomputer at the University of Washington as well as research computing resources at Northwestern University.



Benjamin Mako Hill, Christian Pentzold, and Aaron Shaw. 2026. Peer Production. MIT Press

Jiayi Sun, Hongbo Fang, Junming Zhang, Jiakai Shi, Ruitao Lai, Anita Ihu- man, Richard Littauer, and Shurui Zhou. 2026. Beyond Adoption: Examining the Evolution and Impact of Codes of Conduct on Open Source Communi- ties. In 2026 IEEE/ACM 48th International Conference on Software Engineering (ICSE '26), April 12–18, 2026, Rio de Janeiro, Brazil. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3744916.3773268>

Jian Peng, Jiaxin Zhu, Yuwei Zhang, Wei Chen, Guoquan Wu, Wei Wang, and Jun Wei. 2026. Toward Efficient Package Maintenance: An Empirical Study of Patch Sharing across Four Linux Distributions. In 2026 IEEE/ACM 48th International Conference on Software Engineering (ICSE '26), April 12– 18, 2026, Rio de Janeiro, Brazil. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3744916.3787836>

Andre Hora, Romain Robbes, and Stefano Zacchiroli. 2026. "We Permit the Use of AI, but [...]": The Landscape of AI Policies in Popular Open Source Projects <https://arxiv.org/abs/2609.07542v1>

Credits, cont'd

Hengzhi Ye and Minghui Zhou. 2026. From OSS to Open Source AI: an Exploratory Study of Collaborative Development Paradigm Divergence. *Proc. ACM Hum.-Comput. Interact.* 10, 6, Article CSCW090 (October 2026), 30 pages. <https://doi.org/10.1145/3816938>

Nathan Cassee, Sankarsh Ravi, Italo Santos, Igor Steinmacher, and Alexander Serebrenik. 2026. What barriers do students experience when trying to contribute to Open Source Software projects?. In 2026 IEEE/ACM 48th International Conference on Software Engineering (ICSE-SEET '26), April 12–18, 2026, Rio de Janeiro, Brazil. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3786580.3786971>

Matthew Gaughan, Aaron Shaw, and Darren Gergle. 2026. Linguistic Similarity Within Centralized FLOSS Development. In Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems (CHI EA '26), April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3772363.3798729>

Eli Min, Hana Frluckaj, Matthew Gaughan, Kaylea Champion, Kevin Crowston, Benjamin Mako Hill, and James Howison. 2026. Beyond Productivity: AI, Community, and Sustainability in Open Source Software. In Companion of the Computer-Supported Cooperative Work and Social Computing (CSCW Companion '26), October 10–14, 2026, Salt Lake City, UT, USA. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3785651.3818817>

Eliza Strickland. 2026. How Wikipedia's Editors Keep the Online Encyclopedia Human-Written. *IEEE Spectrum*. <https://spectrum.ieee.org/wikipedia-wikiproject-ai-writing-cleanup>

Yiwei Wu, Leah Ajmani, Nathan TeBlunthuis, and Hanlin Li. 2026. AI Didn't Start the Fire: Examining the Stack Exchange Moderator and Contributor Strike. *Proc. ACM Hum.-Comput. Interact.* 10, 2, Article CSCW004 (April 2026), 27 pages. <https://doi.org/10.1145/3788040>

Courtney Miller, Hao He, Weigen Chen, Elizabeth Lin, Chenyang Yang, Bogdan Vasilescu, and Christian Kästner. 2026. Designing Abandabot: When Does Open Source Dependency Abandonment Matter?. In 2026 IEEE/ACM 48th International Conference on Software Engineering (ICSE '26), April 12– 18, 2026, Rio de Janeiro, Brazil. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3744916.3764519>

Yang Shen, Tao Wang, Xunhui Zhang, Yang Zhang, Cheng Yang, Bo Ding, and Huaimin Wang. 2026. Divergence or Convergence? A Deep Insight into the Crowd Collaboration and its Productivity in Open Source Software basedCHI '26, April 13–17, 2026, Barcelona, Spain on Entropy. In Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26), April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3772318.3791405>

