

The inspirehep package

Live INSPIRE-HEP citation counts, references, and plots

Lawrence Lee

<https://github.com/lawrenceleejr/inspirehep-latex>

Version 0.2 2026/08/29

Abstract

inspirehep puts live data from INSPIRE-HEP into a L^AT_EX document. Hand a command a paper's id and it typesets the title, the citation count, the full reference, or a plot of its citations over time; hand it a person's id and it gives their paper count, citation count, *h*-index, or either as a plot. The numbers live in a generated file, so the document itself always compiles — offline, on a stranger's machine, on Overleaf.

Every example below is executed exactly as shown; the grey panels are this manual's real output, from data fetched on 2026-08-29.

1 Quick start

Put `inspirehep.sty` next to your document (or anywhere on `TEXINPUTS`) and write:

```
\documentclass{article}
\usepackage[hidelinks]{hyperref}
\usepackage{inspirehep}
\begin{document}
Our review \inspirepub{1701002} has aged well.

I have written \inspirepapers{1071846} papers, cited
\inspirecitations[round=1000]{1071846} times.
\end{document}
```

Then either compile with shell escape and let the package fetch for itself,

```
pdflatex -shell-escape cv.tex % fetches, writes inspirehep-data.tex
pdflatex -shell-escape cv.tex % the numbers appear
```

or run the bundled helper once and compile normally:

```
python3 inspirehep-fetch.py cv.tex
pdflatex cv.tex
```

Either way the body renders as:

— you write —
Our review \inspirepub{1701002} has aged well.

I have written \inspirepapers{1071846} papers, cited
\inspirecitations[round=1000]{1071846} times.

you get

Our review **Collider Searches for Long-Lived Particles Beyond the Standard Model** [245 citations] has aged well.
I have written 1,467 papers, cited 207,000 times.

Two passes the first time, exactly like a cross-reference. After that the package only refetches when there is a reason: a record you have just added, or data older than `maxage` days.

2 Where the numbers come from

Both fetchers write the same generated file, `inspirehep-data.tex`: a flat list of declarations the package reads at `\begin{document}`.

```
%% GENERATED by inspirehep-fetch.py -- do not edit.
\inspiresetcites{1701002}{245}
\inspiresettitle{1701002}{Collider Searches for Long-Lived ...}
\inspiresetkey{1701002}{Lee:2018pag}
\inspiresetref{1701002}{L.~Lee, C.~Ohm, A.~Soffer and ...}
\inspiresetyears{2642414}{(2019,1) (2021,2) (2022,6) ...}
\inspiresetauthorstat{1071846}{citations}{207838}
\inspiresetfetched{2026-08-29}{739767}
```

Commit this file. Anyone — including Overleaf, including you offline — then compiles your document with correct numbers and no tooling at all. When the data is older than `maxage` days, every compile says so:

```
Package inspirehep Warning: The data dates from 2026-08-29 (about 130
(inspirehep)                days ago).
(inspirehep)                Refresh with -shell-escape or inspirehep-fetch.py
```

The build still succeeds; the warning is a reminder, not an error.

What each side needs. The self-fetch needs unrestricted shell escape *and* `python3` on the path — it delegates the network work to the helper, so there is exactly one implementation of the API handling. The helper alone needs only `python3`; it reads your sources to find every id you use, so in the common case it takes no arguments. Neither is needed to *typeset*.

Multi-file documents. Point either fetcher at the main file and it follows `\input`, `\include`, and `\subfile` from there, however deeply nested:

```
python3 inspirehep-fetch.py main.tex
Scanning 4 file(s): 3 record(s), 0 plot(s), 0 citation(s), 1 author(s)
```

Names resolve the way L^AT_EX resolves them — relative to the main document’s directory — and then, failing that, relative to the file doing the including, which is what `subfiles` and `import` do. A file that cannot be found is skipped (it may be generated, or inside a conditional), and a cycle of `\inputs` terminates rather than looping. Pass `--no-follow` to scan only the files you name.

Overleaf. Overleaf disables shell escape and its compile servers have no network, so nothing can fetch *there* — but your document still typesets correct numbers, because they are just an `\input` file. To keep them current from inside Overleaf:

1. Upload `inspirehep.sty` to the project.
2. Add `inspirehep-data.tex` with *New file* → *From external URL*, pointing at the raw copy in your repository (kept fresh by CI or your local builds).
3. Overleaf gives linked files a *Refresh* button: new numbers are one click, no recompile tricks.

Overleaf’s Git sync does not fetch submodules, so include the `.sty` as a plain file there.

3 Papers

A record can be named two ways, and every command takes either — they are told apart by shape, so nothing has to be declared:

	Looks like	Where to find it
recid	2642414	the number in the record’s URL, inspirehep.net/literature/2642414
texkey	Accettura:2023ked	the B _I B _T _E X key INSPIRE assigns

you write <pre>by recid: \inspirepub[cites=false]{1701002}</pre> <pre>by texkey: \inspirepub[cites=false]{Lee:2018pag}</pre>
you get <pre>by recid: Collider Searches for Long-Lived Particles Beyond the Standard Model</pre> <pre>by texkey: Collider Searches for Long-Lived Particles Beyond the Standard Model</pre>

Both link to the record by number, so a texkey costs nothing. A record with more than one texkey answers to any of them.

3.1 `\inspirepub`: one command, options for what it shows

<code>\inspirepub{2642414}</code>
you get
Towards a muon collider [407 citations]

The title came from INSPIRE, linked to the record, with the live citation count. Each piece is an option:

<code>\inspirepub[year]{1701002}</code>
you get
Collider Searches for Long-Lived Particles Beyond the Standard Model [245 citations] (2019)

<code>\inspirepub[cites=false]{2642414}</code>
you get
Towards a muon collider

<code>\inspirepub[link=false, cites=false]{2642414}</code>
you get
Towards a muon collider

Prefer your own wording? Your title, INSPIRE's count:

<code>\inspirepub[title={A muon collider design study}]{2642414}</code>
you get
A muon collider design study [407 citations]

And `ref` swaps the title for the full reference, exactly as INSPIRE formats it (see the next subsection):

<code>\inspirepub[ref, cites=false]{1701002}</code>
you get
L. Lee, C. Ohm, A. Soffer and T. T. Yu, “Collider Searches for Long-Lived Particles Beyond the Standard Model,” Prog. Part. Nucl. Phys. 106 (2019), 210-255 doi:10.1016/j.ppnp.2019.02.006 [arXiv:1810.12602 [hep-ph]].

3.2 The pieces, separately

Everything `\inspirepub` assembles is also a command of its own, for building your own layout:

you write	
<pre>\begin{tabular}{@{}l@{}} title & \inspiretitle{2642414} \\ count & \inspirecites{2642414} \\ year & \inspireyear{2642414} \\ key & \texttt{\inspirekey{2642414}} \end{tabular}</pre>	
you get	
title	Towards a muon collider
count	[407 citations]
year	(2023)
key	Accettura:2023ked

3.3 References

`\inspireref` gives the reference as INSPIRE itself renders it — authors, journal, errata, DOI, arXiv number — so it matches what everyone else in the field cites, and there is no author-list or journal-abbreviation logic in this package to get wrong.

you write	
<pre>\inspireref[cites=false]{1701002}</pre>	
you get	
L. Lee, C. Ohm, A. Soffer and T. T. Yu, “Collider Searches for Long-Lived Particles Beyond the Standard Model,” <i>Prog. Part. Nucl. Phys.</i> 106 (2019), 210-255 doi:10.1016/j.pnpnp.2019.02.006 [arXiv:1810.12602 [hep-ph]].	

By default it, too, carries the count (`cites=false` above turned it off).

INSPIRE records errata, which a CV usually does not want, so they are off by default. `errata` puts them back:

you write	
<pre>\inspireref[cites=false, errata]{1701002}</pre>	
you get	
L. Lee, C. Ohm, A. Soffer and T. T. Yu, “Collider Searches for Long-Lived Particles Beyond the Standard Model,” <i>Prog. Part. Nucl. Phys.</i> 106 (2019), 210-255 doi:10.1016/j.pnpnp.2019.02.006 [arXiv:1810.12602 [hep-ph]]. [erratum: <i>Prog. Part. Nucl. Phys.</i> 122 (2022), 103912]	

A collaboration paper is credited by INSPIRE to its first author and *et al.*; `collab` credits the collaboration instead, which is the usual convention in a publication list:

you write	
<pre>\inspireref[cites=false]{1298030}</pre>	
you get	
G. Aad <i>et al.</i> [ATLAS], “Light-quark and gluon jet discrimination in <i>pp</i> collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector,” <i>Eur. Phys. J. C</i> 74 (2014) no.8, 3023 doi:10.1140/epjc/s10052-014-3023-z [arXiv:1405.6583 [hep-ex]].	

you write <pre>\inspireref[cites=false, collab]{1298030}</pre>
you get ATLAS Collaboration, “Light-quark and gluon jet discrimination in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector,” Eur. Phys. J. C 74 (2014) no.8, 3023 doi:10.1140/epjc/s10052-014-3023-z [arXiv:1405.6583 [hep-ex]].

The wording is `\inspirecollabformat`, so `\renewcommand{\inspirecollabformat}[1]{#1}` gives a bare “ATLAS”. Both options work per call or document-wide — the reference is stored in its parts and assembled as it is typeset, so switching either one never means refetching.

The `style` option picks INSPIRE’s European numbering (`latex-eu`, the default: volume (year) page) or American (`latex-us`: volume, page (year)); that one does apply at fetch time, so set it as a package option.

3.4 A .bib file you never write

`\inspirecite{<id>}` is `\cite` by INSPIRE id. The fetcher collects INSPIRE’s own BibTeX entry — verbatim, from `?format=bibtex` — for every id you cite, into `inspirehep-refs.bib`:

```
\usepackage{inspirehep}
...
as shown in \inspirecite{1701002}, ...
...
\bibliographystyle{unsrt}
\bibliography{inspirehep-refs}
```

The generated file begins:

```
%% GENERATED by inspirehep-fetch.py -- do not edit.
@article{Lee:2018pag,
  author = "Lee, Lawrence and Ohm, Christian and Soffer, Abner and
    Yu, Tien-Tien",
  title = "{Collider Searches for Long-Lived Particles Beyond the
    Standard Model}",
  ...
```

You never handle a BibTeX key yourself — though `\inspirekey{<id>}` hands it to you if you want it (it is `Lee:2018pag` here), so `\inspirecite` composes with everything that expects `\cite`.

`\inspirecite` takes a texkey too, which is the one case where it works before anything has ever been fetched: a texkey *is* the citation key, so `\inspirecite{Lee:2018pag}` is simply `\cite{Lee:2018pag}`. The entry still has to reach your `.bib`, so a fetch is still what puts it there.

4 People

The id is the number in a profile URL, `inspirehep.net/authors/1071846`, or a BAI such as `J.Smith.1`. Every command takes the id it is asking about — there is deliberately no

package-wide author — so one document can discuss several people:

you write

```
\begin{tabular}{@{}lrrr@{}}
      & papers & citations & $h$ \\
L.~Lee & \inspirepapers{1071846}
      & \inspirecitations{1071846}
      & \inspirehindex{1071846} \\
T.~Holmes & \inspirepapers{1067690}
      & \inspirecitations{1067690}
      & \inspirehindex{1067690}
\end{tabular}
```

you get

	papers	citations	h
L. Lee	1,467	207,838	211
T. Holmes	1,222	147,191	189

4.1 Round numbers that stay true

A summary line usually wants “over N ”, and the rounding must go *down*: rounded up, the claim is false the moment it is typeset. `round=` rounds any figure down to a multiple of its value:

you write

```
Over \inspirepapers[round=100]{1071846} papers, with over
\inspirecitations[round=1000]{1071846} citations and an
$h$-index of \inspirehindex{1071846}.
```

you get

Over 1,400 papers, with over 207,000 citations and an h -index of 211.

`\inspireauthorstat{<id>}{<key>}` is the general form behind all three, with `<key>` one of papers, citations, hindex:

you write


```
\inspireauthorstat[round=1000]{1071846}{citations}
```

you get


207,000

5 Plots


Load the package with the `plots` option (it pulls in `pgfplots`, which is why it is opt-in — a CV that never plots should not pay for `TikZ`). `\inspireplot{<id>}` is the citation history of one paper:



`\inspireauthorplot{<id>}{papers}` or `{citations}` is a person's output or impact over time:



The default style is deliberately spare: one line, axis lines only where they carry information, no grid, box, marks, or legend. Three hooks change it — `\inspireplotwidth`, `\inspireplotheight`, and `\inspireplotstyle`, which takes any `pgfplots` axis keys:



Braces confine the change to one plot; at the top level it holds for all of them.

6 Numbers and separators

Figures are digit-grouped with a comma by default. `sep=` chooses another convention, per call or document-wide, and `sepstring=` takes anything else:

— you write —

```
\begin{tabular}{@{}ll@{}}
  \texttt{sep=comma}      & \inspirecitations{1071846} \\
  \texttt{sep=period}    & \inspirecitations[sep=period]{1071846} \\
  \texttt{sep=space}     & \inspirecitations[sep=space]{1071846} \\
  \texttt{sep=thin}      & \inspirecitations[sep=thin]{1071846} \\
  \texttt{sep=underscore}& \inspirecitations[sep=underscore]{1071846} \\
  \texttt{sep=none}      & \inspirecitations[sep=none]{1071846} \\
  \texttt{sepstring='}   & \inspirecitations[sepstring='{'}]{1071846}
\end{tabular}
```

you get

sep=comma	207,838
sep=period	207.838
sep=space	207 838
sep=thin	207838
sep=underscore	207_838
sep=none	207838
sepstring='	207'838

Set once for the whole document with `\usepackage[sep=thin]{inspirehep}` or anywhere with `\inspiresetup{sep=thin}`; either way a per-call option still wins for that one call, and a change made inside a group ends with the group.

7 When data is missing

An id the data file knows nothing about — brand new, mistyped, or never fetched — typesets a visible marker rather than something silently wrong, and flags itself so the next fetch asks about it:

— you write —

```
\inspirepub{999999999}
```

you get

[? title 999999999]

Fresh papers deserve better than advertising a zero: counts below `mincites` (default 1) print nothing at all.

— you write —

```
before\inspirecites{999999999}after
({\inspiresetup{mincites=1000}%
  also nothing: 152 < 1000\inspirecites{1298030}})
```

you get

beforeafter (also nothing: 152 < 1000)

Fallback values can be declared in the document, used only until the real figures arrive — handy for a repository’s first ever build, and they never override fetched data:

you write <pre> \inspiredefaultauthorstat{20000000}{papers}{42} \inspiredefaultauthorstat{1071846}{papers}{9} never fetched: \inspirepapers{20000000}; fetched wins: \inspirepapers{1071846} </pre>
you get <pre> never fetched: 42; fetched wins: 1,467 </pre>

`\inspiredefaultstat{<key>}{<n>}` does the same for the single person named by the helper's `-author`, who needs no id at the call site.

8 Options

Every option is both a package option and (where it makes sense per call) an optional argument; `\inspiresetup{...}` changes any of them mid-document.

Option	Default	Meaning
<code>cites</code>	<code>true</code>	show the citation count
<code>ref</code>	<code>false</code>	full reference instead of the title
<code>year</code>	<code>false</code>	append the publication year
<code>errata</code>	<code>false</code>	include any erratum in a reference
<code>collab</code>	<code>false</code>	credit a collaboration paper to the collaboration rather than to its first author
<code>link</code>	<code>true</code>	hyperlink the title (needs <code>hyperref</code> ; plain text without it)
<code>title</code>	—	override the fetched title
<code>round</code>	<code>1</code>	round a figure <i>down</i> to a multiple of this
<code>mincites</code>	<code>1</code>	counts below this print nothing
<code>sep</code>	<code>comma</code>	<code>comma</code> , <code>period</code> , <code>space</code> , <code>thin</code> , <code>underscore</code> , <code>none</code>
<code>sepstring</code>	—	any other separator text
<code>plots</code>	<code>false</code>	load <code>pgfplots</code> for the plotting commands
<code>style</code>	<code>latex-eu</code>	reference format, <code>latex-eu</code> or <code>latex-us</code> (applies at fetch time)
<code>data</code>	<code>inspirehep-data</code>	basename of the generated data file
<code>bib</code>	<code>inspirehep-refs</code>	basename of the generated <code>.bib</code>
<code>script</code>	<code>inspirehep-fetch.py</code>	where the self-fetch finds the helper (a submodule path, say)
<code>maxage</code>	<code>120</code>	days before the data is called stale
<code>fetch</code>	<code>auto</code>	<code>auto</code> , <code>on</code> (warn if impossible), <code>off</code>

9 House style

Five hooks control how things are set; redefine any of them:

— you write —

```
{\renewcommand{\inspiretitleformat}[1]{\textsc{#1}}%
\renewcommand{\inspirecitetext}[1]{cited #1 times}%
\renewcommand{\inspirecitesformat}[1]{ --- \emph{#1}}%
\inspirepub{1298030}}
```

you get

LIGHT-QUARK AND GLUON JET DISCRIMINATION IN pp COLLISIONS AT $\sqrt{s} = 7$ TeV
WITH THE ATLAS DETECTOR — *cited 152 times*

```
\renewcommand{\inspiretitleformat}[1]{\textbf{#1}}      % around the title
\renewcommand{\inspirecitetext}[1]{#1~citations}       % the words
\renewcommand{\inspirecitesformat}[1]{\small[#1]}      % attaching them
\renewcommand{\inspireyearformat}[1]{~(#1)}            % the year
\renewcommand{\inspirenumformat}[1]{...}               % digit grouping itself
\renewcommand{\inspirenumsep}{\,,}                    % what goes between groups
\renewcommand{\inspirecollabformat}[1]{#1~Collab.}     % ‘collab’ wording
\renewcommand{\inspiremissing}[1]{\textbf{[?~#1]}}     % the missing-data marker
```

`\inspirenumsep` is the separator on its own, for a character the `sep` option does not offer;
`\inspiremissing` is what an unknown id typesets — the next section describes when, and
this is how to restyle it.

A CV-flavoured example — the count as quiet grey metadata:

— you write —

```
{\renewcommand{\inspirecitesformat}[1]{%
\nobreakspace{\small\textcolor{black!45}{[#1]}}}%
\inspirepub{2642414}}
```

you get

Towards a muon collider [407 citations]

10 The data file, for the curious

The declarations a fetcher writes are ordinary commands, so you can also set or patch anything by hand — useful for records INSPIRE does not have:

<code>\inspiresetcites{<id>}{<n>}</code>	a record’s citation count
<code>\inspiresetttitle{<id>}{<title>}</code>	its title
<code>\inspiresetkey{<id>}{<key>}</code>	its BibTeX key
<code>\inspiresetrecid{<id>}{<recid>}</code>	the record number behind a texkey, so links still work
<code>\inspiresetrefauthors{<id>}{...}</code>	its author list
<code>\inspiresetrefcollab{<id>}{ATLAS}</code>	its collaboration, if any
<code>\inspiresetrefpub{<id>}{...}</code>	journal, DOI, arXiv number
<code>\inspiresetreferrata{<id>}{...}</code>	any erratum
<code>\inspiresetref{<id>}{<reference>}</code>	a whole reference, overriding the four above
<code>\inspiresetyears{<id>}{(y,n) (y,n) ...}</code>	its citations per year
<code>\inspiresetauthorstat{<id>}{<key>}{<n>}</code>	a person’s figure
<code>\inspiresetstat{<key>}{<n>}</code>	the same, for the one person named by <code>-author</code>
<code>\inspiresetauthoryears{<id>}{<kind>}{...}</code>	a person’s series
<code>\inspiresetfetched{<date>}{<dayno>}</code>	when it was all fetched

11 Requirements and limits

Typesetting needs a current L^AT_EX (xparse and l3keys2e from the kernel), plus pgfplots if and only if you pass `plots`. Refreshing the data needs `python3` (standard library only); the self-fetch additionally needs unrestricted shell escape. Counts are whatever INSPIRE reports, self-citations included. A texkey given to `\inspirecite` is used as-is when nothing has been fetched yet; every other command needs the data file either way.

12 Licence

LPPL 1.3c or later; maintained by Lawrence Lee. Issues and contributions: <https://github.com/lawrenceleejr/inspirehep-latex>.