

Illustrative notes for obsessing over publishing aesthetics

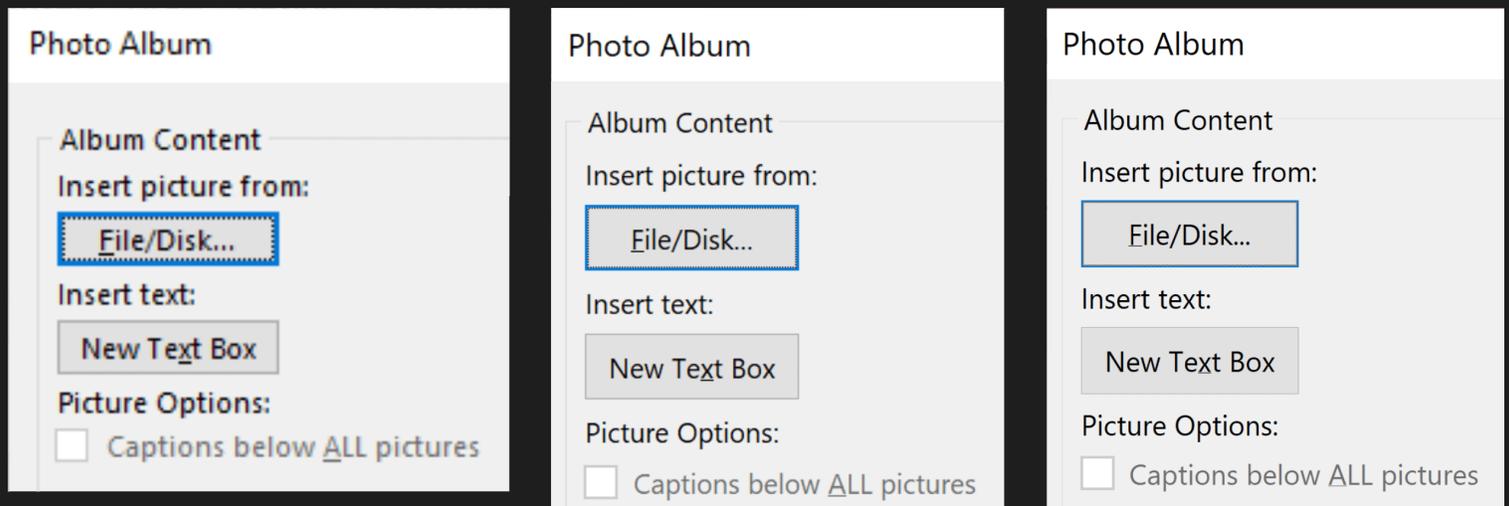
By [Jeff Huang](#) on 2022-06-27

These are my notes for making the most precise text and images for publication, mainly for PDF and print media. Unlike articles written for viewing on the web, publications are more rigid and should look good in higher densities like 600 DPI (dots-per-inch). I am only listing techniques that are missed by the vast majority of authors. View [this page saved as a PDF file](#) for a more accurate rendering of this page's examples, to see finer detail in the images.

Screenshots

1 Scaling up the display resolution before taking a screenshot

Screen resolutions are typically in 96 DPI for desktop at 100% scaling, so screenshots taken on a desktop don't have much detail, compared to a high-resolution print of 600 DPI. They'll appear pixelated when printed out, or when shown in a PDF file where it may be stretched to fit the width of the reader's screen. There's a trick to fix this: on your desktop operating system settings, adjust the Display scaling as high as you can. This scaling reduces the space available on screen, but as long as you can fit what you need in a screenshot, it will provide more pixels for detail when you take the screenshot. Below is an example of screenshots of the exact same interface taken at 100%, 200%, and 400% scaling, then placed in the same Word Document and saved as PDF to view.



PowerPoint's Insert Photo Album dialog screenshot taken in 100% (left), 200% (center), and 400% (right) scaling

2 Capturing vector-based screenshots of web pages

If you're taking a screenshot of a web page, you can instead save it to a PDF file so that it's vector-based instead of pixel-based. Raster graphics embedded in the web page, such as PNG and JPEG, will remain pixel-based but all the text, vector graphics, and CSS shapes like rounded buttons will keep their precision. This means it will look good at any zoom-level (any DPI). But in the ideal scenario, you would get a perfect representation of the website even when printed in a giant movie poster, or zoomed to the maximum level possible. Be sure to remove margins in the print dialog settings, then switch to landscape mode (since screens are wider than they are tall), and change the paper size to something like A3 to more closely resemble a typical 21-inch monitor's size. The example below is using the Save to PDF printer option in Firefox's Print dialog. The resulting PDF was converted to SVG to show on this page, and a small region is stretched into the clip on the right side to show how it looks at higher DPI.

Sort by

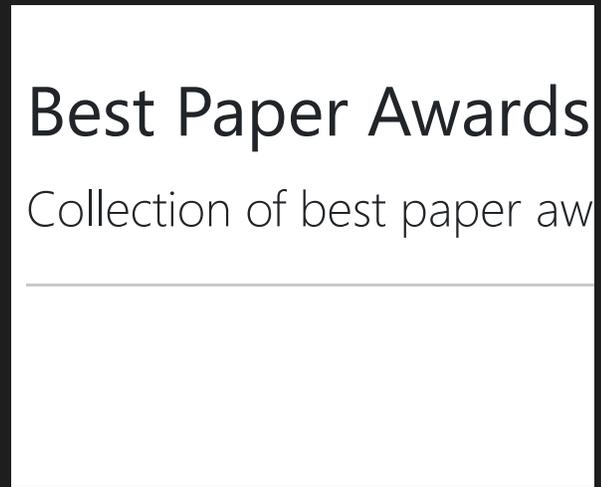
Institutions Ranked

Conferences

Years

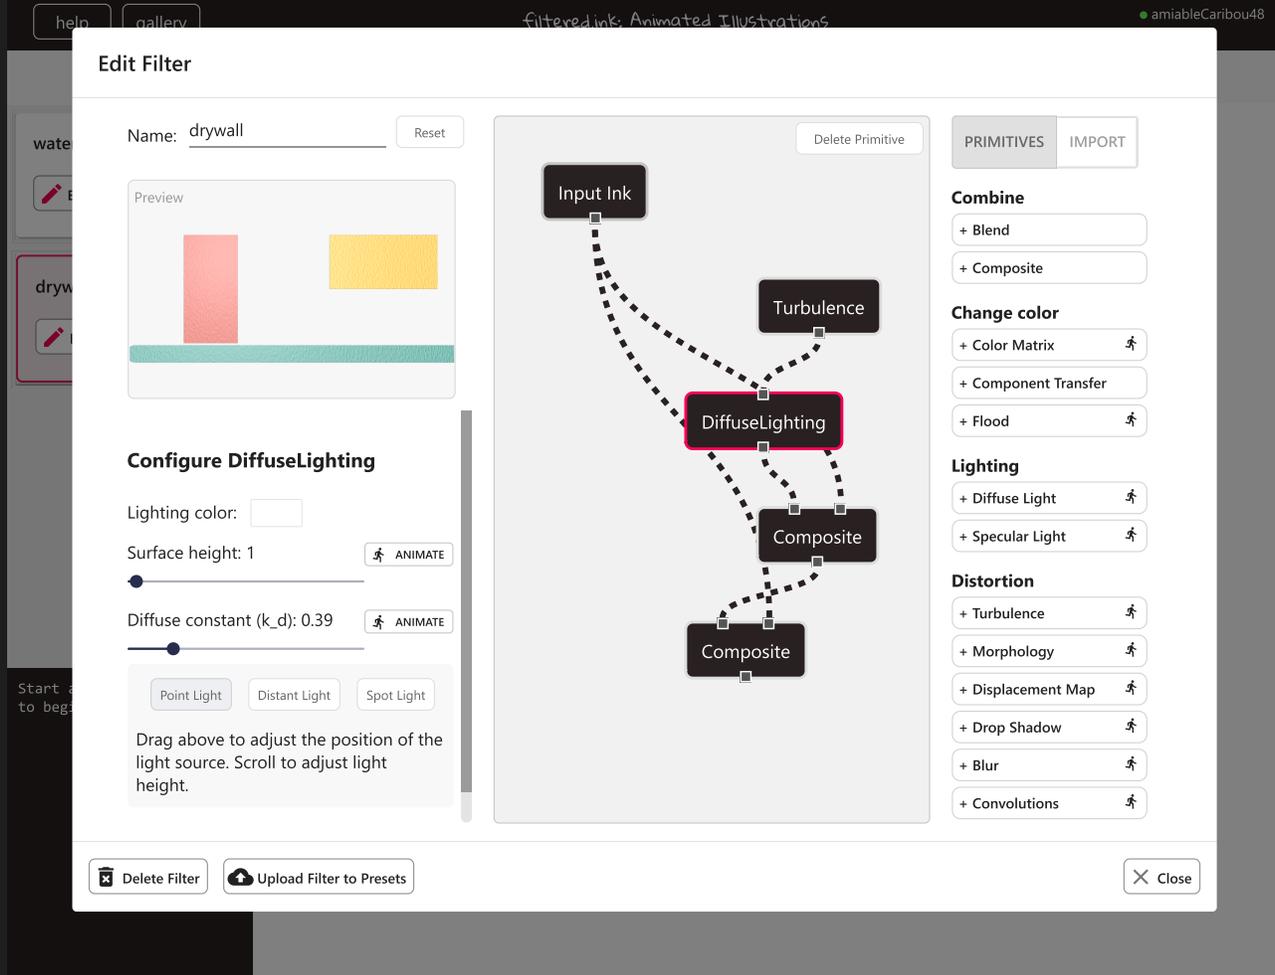
- 2021
- 2020
- 2019
- 2018
- 2017
- 2016
- 2015
- 2014
- 2013
- 2012
- 2011
- 2010
- 2009
- 2008
- 2007
- 2006
- 2005
- 2004
- 2003
- 2002
- 2001
- 2000
- 1999
- 1998
- 1997
- 1996

Best Paper Awards in Computer Science		
Collection of best paper awards for 30 computer science conferences since 1996		
<p>This is a collection of best paper awards from the main conferences in each computer science subfield, starting from 1996. This data was entered by hand from sources found online (many of them no longer available), so please email bestpaper@jeffhuang.com if you notice any errors or omissions. The page is maintained annually by Jeff Huang, thanks to contributions from Minghui Ray Zhang (2017, 2018), Will Gierke (2019, 2020), Alice Marbach (2019, 2020), AllenAI's Dawn Howell (2019, 2020), Long Do (2021), and Shaun Wallace (2021).</p> <p>Caveats: Note that some conferences do not have such an award (e.g. SIGGRAPH, CAVI, "Distinguished paper award" and "outstanding paper award" are included but not "best student paper" or "best 10-year old paper" (e.g. PODPL until 2018); at this point, it is unlikely that new additional conferences will be added due to the ongoing time commitment to prepare updates (about 8 hours a year); only each author's first affiliation is listed due to how the data was originally stored.</p>		
2021		
AAAI Artificial Intelligence	Informer: Beyond Efficient Transformer for Long Sequence Time-Series Forecasting	Haozi Zhou, Beihang University, et al.
	Exploration-Exploitation in Multi-Agent Learning: Catastrophe Theory Meets Game Theory	Stefanos Leonardos & Georgios Pilouras, Singapore University of Technology and Design
ACL Natural Language Processing	Vocabulary Learning via Optimal Transport for Neural Machine Translation	Jingjing Xu, Bytelligence, et al.
CHI Human-Computer Interaction	What Do Hackathons Do? Understanding Participation in Hackathons Through Program Theory Analysis	Jeanette Falk, Aarhus University, et al.
	Fake Synthesis-Powered Visualization Authoring	Chenglong Wang, University of Washington, et al.
	"Can I Not Be Social on a Sunday?": Understanding Technology-Mediated Pathways to Mental Health Support	Sachin R. Pendse, Georgia Institute of Technology, et al.
	CapContact: Super-resolution Contact Areas from Capacitive Touchscreens	Paul Streek & Christian Holz, ETH Zurich
	RadialNet: Efficient Gesture Recognition Technique Utilizing a Miniaturized Radial Sensor	Eij Wiyashi, Google, et al.
	Screen Recognition: Creating Accessibility Metadata for Mobile Applications from Pixels	Xiaoyi Zhang, Apple, et al.
	LGBTQ Persons' Pregnancy Loss Disclosures to Known Ties on Social Media: Disclosure Decisions and Ideal Disclosure Environments	Cassidy Pyle, University of Michigan, et al.
	Increasing Electrical Muscle Stimulation's Dexterity by means of Back of the Hand Actuation	Atsushi Nakahashi, University of Chicago, et al.
	Designing Menstrual Technologies with Adolescents	Marie Louise Isal Sandergaard, Royal Institute of Technology, et al.
	The Show Must Go On: A conceptual model of conducting synchronous participatory design with children online	Kung Jin Lee, University of Washington, et al.
	Building for "We": Safety Settings for Couples with Memory Concerns	Nora McDonald & Helena M. Mentis, University of Maryland, Baltimore County
	Oh, Sissy! A Fabrication Pipeline to Magically Connect Conventional and 3D-Printed Electronics	Martin Schmalz, TU Darmstadt, et al.
	Getting Ourselves Together: Data-censored participatory design research & epistemic burden	Jennifer Piene, University of California, Los Angeles, et al.
	"Everyone wants to do the model work, not the data work": Data Cascades in High-Stakes AI	Nithya Sambasivan, Google, et al.
	Resisting the Medicalisation of Menopause: Reclaiming the Body through Design	Mariela Coll-Felice, Royal Institute of Technology, et al.
	UScientist: Designing for People-Powered Research in Museums	Mmachi Godfrey Obiorah, Northwestern University, et al.
	"Why lose control?" A Study of Freedom's Perspectives with Gig Economy Platforms	Juan Carlos Alvarez de la Higuera, Northumbria University, et al.
	"That courage to encourage": Participation and Aspirations in Chat-based Peer Support for Youth Living with HIV	Nawona Kansala, University of Washington, et al.
	The Ethics of Multiplayer Game Design and Community Management: Industry Perspectives and Challenges	Lucy A. Sparrow, University of Melbourne, et al.
	Coupling Simulation and Hardware for Interactive Circuit Debugging	Evan Strainick, Stanford University, et al.
	Impact of Task on Attentional Tunneling in Handheld Augmented Reality	Brandon Victor Szymek, University of Melbourne, et al.
	X Gnomics: Facilitating the Creation of Ergonomic 3D Interfaces	José Marcelo Evangelista Beto, Aarhus University, et al.
	Heuristic Evaluation of Conversational Agents	Raina Langevin, University of Washington, et al.
	Characterizing the Effectiveness of Anti-harassment Feedback and Sanctions in Chat	Amber Crockett, University of Washington, et al.



A web page saved using Firefox's Save as PDF (left), with a portion zoomed in to show the fine detail (right)

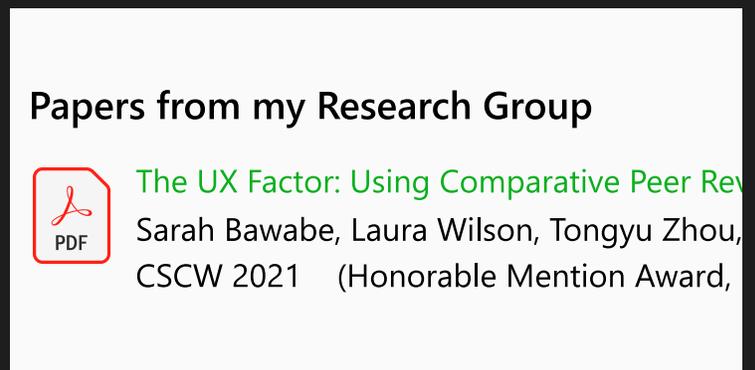
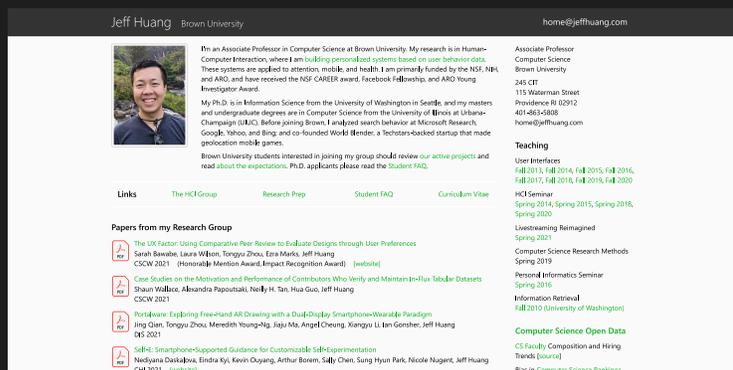
This is an example vector screenshot taken of the [filtered.ink](#) interface (a web application for editing SVG filters) to show the same technique for capturing vector screenshots is available in Chrome.



A web application's interface screenshots using Chrome's Save as PDF print option; interface elements are perfectly clear at any resolution

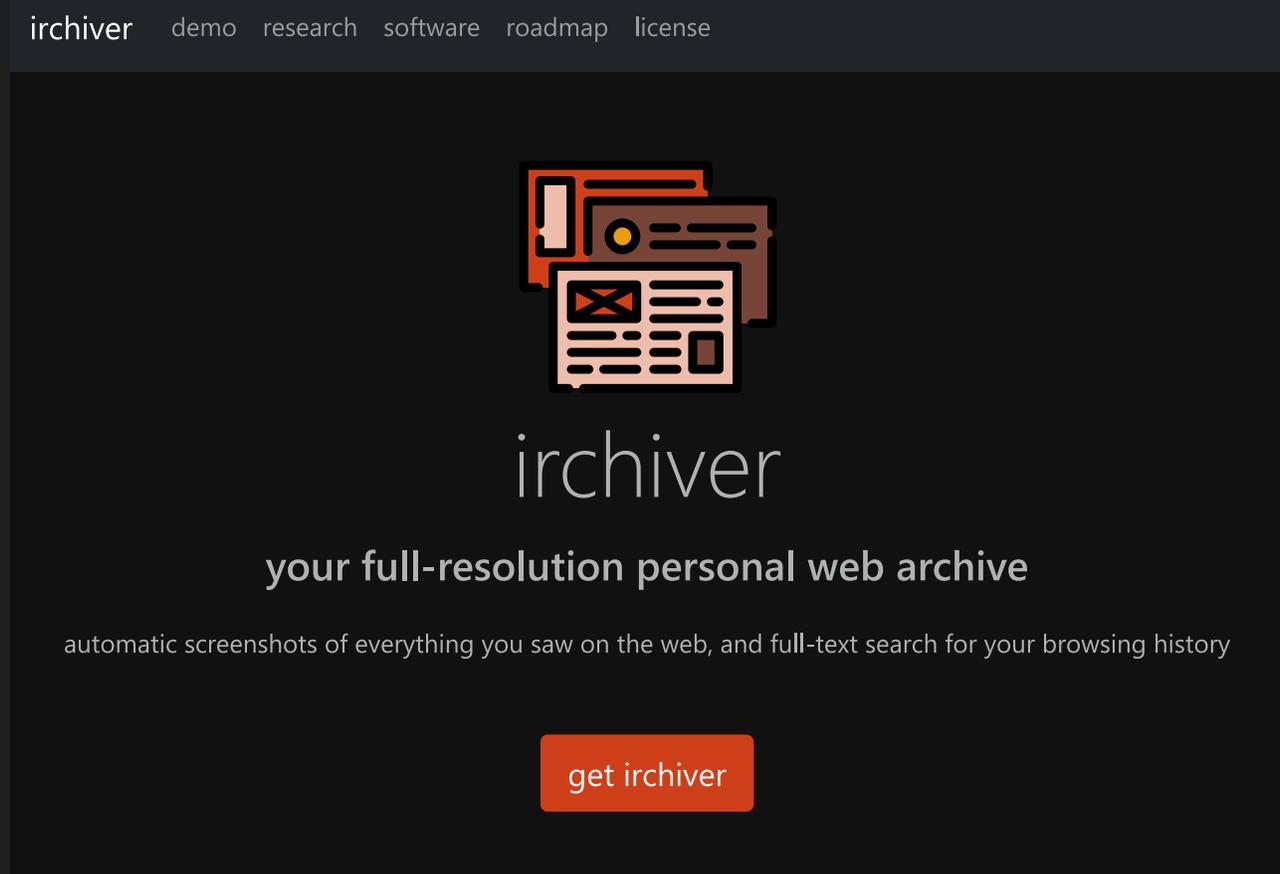
3 Emulating screen media before saving as PDF

Many websites won't look right when saved as PDF using the previous Save as PDF trick. There are a few more settings that may be helpful for getting that screenshot just right. First, pages often look very different because of the print media stylesheet being used instead of the screen stylesheet. You can force Chrome to use the screen stylesheet by going to Rendering > Emulate CSS media type and selecting 'screen'. The other thing is to play around with the scaling in the print dialog. Both these were used to get my homepage saved as PDF, and you can see even the PDF icon (which is an SVG image) is pixel-perfect in the blown-up clip on the right.



A website saved by first emulating screen media before Save as PDF (left), with a portion zoomed in to show perfect detail (right)

Here is a vector screenshot of a mobile-sized viewport for irchiver.com.



A website printed to mobile size to get a screenshot of how it appears on a mobile device

Photos

4 [Word] Ensuring JPEGs are not recompressed when loaded or exported

JPEGs are already a lossy format, so photos saved in JPEG are delicate; any editing or even resaving them will recompress the image, which decreases the quality of the image. This is because JPEGs are approximated with equations rather than having every pixel specified in the file, so when you approximate something that was already an approximation, quality degrades further, sometimes revealing compression artifacts like tiny worms near the sharp contrasts in the photo.

When using JPEGs in publications, LaTeX leaves the original JPEG intact but Microsoft Word's default settings can inadvertently recompress your JPEGs (and Google Docs always recompresses JPEGs unfortunately). Microsoft Word will compress high-resolution JPEGs in the document, so that they're

downsampled when they're saved. This can be disabled in the application under File > Options > Advanced > Image Size and Quality > Do not compress images in file.

However, what cannot be disabled in the Windows version of Microsoft Word is the downsampling that occurs when saving as PDF. So on Windows, even if the JPEG in the Word document is the original image, that image will be inexplicably recompressed when saving as PDF. Fortunately, the MacOS version of Word can preserve the exact original JPEG when saving to PDF, so to avoid recompression with Microsoft Word documents, switch to the MacOS version of Word for the final saving to PDF, making sure you have "Best for printing" selected in the Save dialog. Below is the same DOCX file containing a high-resolution JPEG saved as PDF on Windows and then on MacOS, then zoomed in at 400% to emphasize the difference.



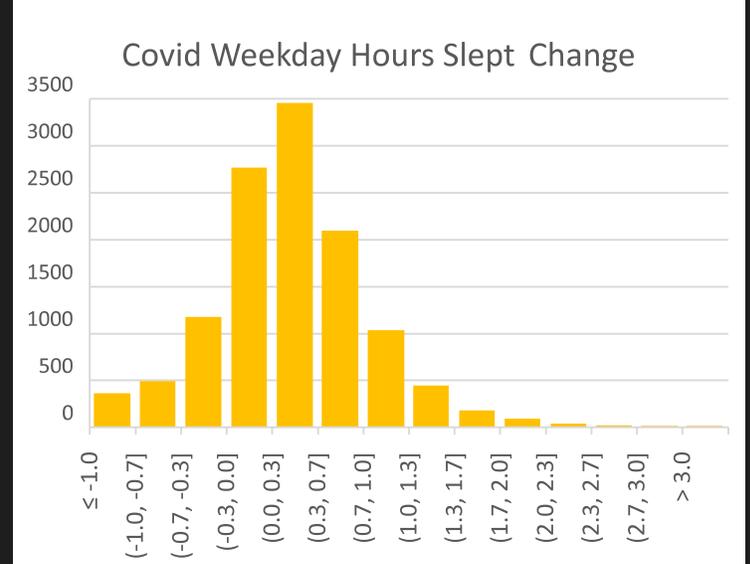
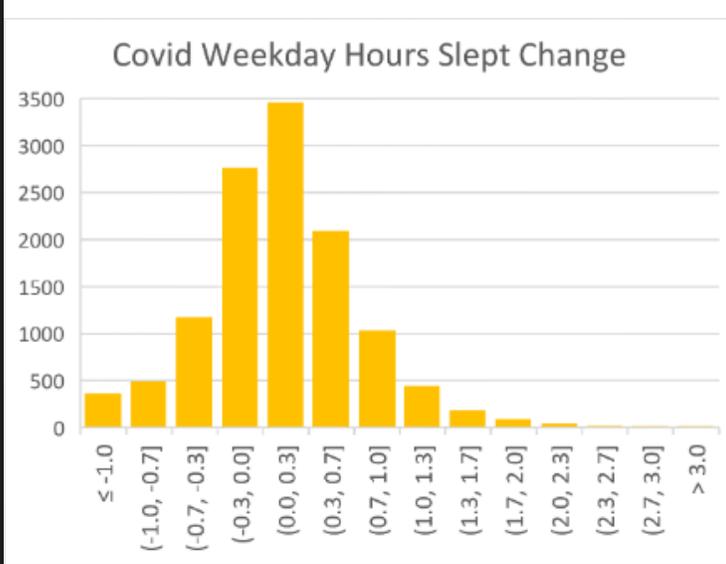
A photo zoomed in at 400% to show the difference in quality when saved in Microsoft Word on Windows (left) and MacOS (right)

Graphics

5 Getting illustrations and vector graphics from the source

Vector graphics are composed of lines and shapes, embedding instructions that say things like "draw a line from the left of the image 2/3s from the bottom to the top right." For many types of graphics, it is a more precise and shorter way of describing the image compared to raster (pixel-based) images, which specifies the color of each pixel. Charts, diagrams, logos, wireframes, calendars, and tables should be vector graphics. Converting image formats may take a few steps, so avoid accidentally converting the image into a pixel representation at any step along the way (such as taking a screenshot or saving as PNG), as rasterizing the vectors is irreversible.

Always save directly from the application to a vector graphic file type (SVG, PDF). If the application doesn't have an obvious way to export to one of those file type, then try printing to PDF. A common mistake is saving charts as raster graphics instead of vector graphics (a raster chart in the future on the left, below). Instead, it should be exported or printed to PDF. This may be counterintuitive, but even Excel charts can be printed to PDF, by first selecting the chart (the figure on the right, below).

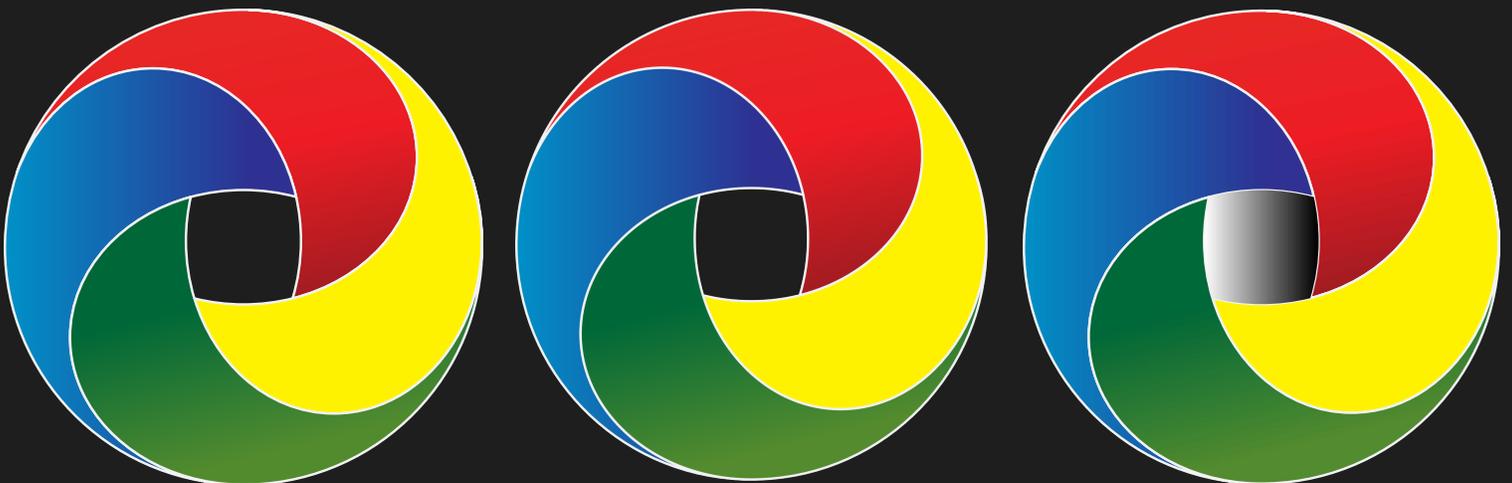


An Excel chart saved as a PNG screenshot (left) compared to printing to PDF (right)

Most chart-generating libraries can output to SVG directly, which have similar capabilities as PDF, but they should appear the same. Note that for compatibility, you'll want to outline the fonts when generating SVG files, because SVGs can't embed fonts. If you leave the fonts in the SVG file, the text will look differently depending on what device you are using to view it, which is undesirable.

6 [Word] Inserting PDF graphics into Microsoft Word on MacOS

While Microsoft Word on both Windows and MacOS can have vector graphics by inserting SVG files, they both seem a little inconsistent when saving documents with SVGs to PDF. I've noticed this especially for SVGs containing gradients or more complex features. In the below example SVG graphic, which is the "Vector-based example.svg" from Wikipedia, the left image shows what the original looks like. The center image shows the output from Microsoft Word on Windows, which if you can look closely, you'll notice it has rasterized the graphic when it saved it to PDF; this may be noticeable only when zoomed closely. The right image shows the output from Microsoft Word on MacOS, which retains the vector graphic format, but somehow added a gradient in the center and overlays subtle lines on top of each colored section.



An example SVG image file (left) saved from in Microsoft Word document in Windows (center) or MacOS (right)

However, Microsoft Word on MacOS can actually insert PDF files, so you can convert an SVG file to PDF and insert it, and when the document is saved as PDF, it will render quite like the original. Consider it the ultimate backup option when SVGs are not rendering correctly when you save your documents in Microsoft Word to PDF.

7 Trimming whitespace from images

Both raster and vector graphics can have extra (invisible) whitespace on the edges, which adds padding to the image when it's included in document. This prevents the image from aligning with the text, and makes it unnecessarily smaller while using up the same space. Removing this extra space is called "trimming", but for vector graphics it can be a bit trickier. Using a tool like Adobe Illustrator, you can adjust the artboard bounds to fit perfectly around the visible parts of the graphic using "Fit to Artboard Bounds". Sometimes there will be invisible elements preventing you from shrinking the artboards neatly around the visible graphic; so use the white arrow (Direct Selection) to draw selection boxes around invisible bits in that white space and delete them (sometimes pressing Delete twice to delete the entire invisible element instead of just a node), then try "Fit to Artboard Bounds" again so that the image fills the artboard exactly.

In the example below, the figure embedded in the paper on the left had a typical margin, so doesn't quite fill the column; the figure embedded in the paper on the right has had its whitespace trimmed so is a bit easier to read and aligns better with the edges.

e this reflects wild.

Completion
pCoach and these designs, If-experiment days in [14]), as designed to as 3 days per a set length of ad to tracking ther a shorter

nciple to set a lt was shown, periments for d this flexible required mini- d for a graph o of the three ached a result was too long.

nights tracked f participants east 4 nights, fourteen per- heir sleep for

That way, people will be able to learn something quantitative about themselves and make informed decisions about their behavior change choices. For the purposes of this paper, we also conducted a traditional t-test analysis on the users' data, which revealed that none of them would have reached a statistically significant result at the end of their experiments ($p < 0.05$).

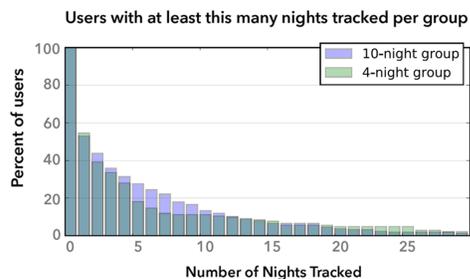


Figure 6: The overall number of nights tracked followed a similar trend in both groups of users. Around the fifth night, however, there is a dip in the percentage of users from the 4-night group that tracked their sleep (which is when those users saw a result in the app).

Reasons for Users Ending the Experiment
Once users reached the required minimum number of days in

e this reflects wild.

Completion
pCoach and these designs, If-experiment days in [14]), as designed to as 3 days per a set length of ad to tracking ther a shorter

nciple to set a lt was shown, periments for d this flexible required mini- d for a graph o of the three ached a result was too long.

nights tracked f participants east 4 nights, fourteen per- heir sleep for

That way, people will be able to learn something quantitative about themselves and make informed decisions about their behavior change choices. For the purposes of this paper, we also conducted a traditional t-test analysis on the users' data, which revealed that none of them would have reached a statistically significant result at the end of their experiments ($p < 0.05$).

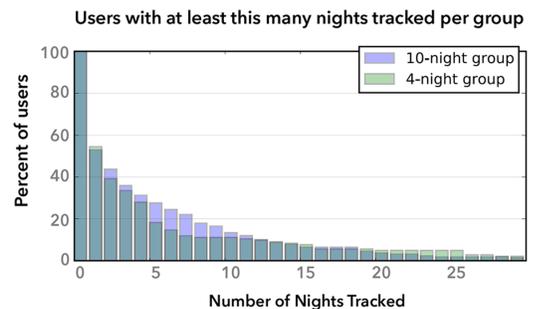


Figure 6: The overall number of nights tracked followed a similar trend in both groups of users. Around the fifth night, however, there is a dip in the percentage of users from the 4-night group that tracked their sleep (which is when those users saw a result in the app).

Reasons for Users Ending the Experiment
Once users reached the required minimum number of days in

An image with typical margins (left) compared to the same image with margins trimmed (right) embedded into a paper

8 [Word] Kerning, ligatures, and hyphenation

Microsoft Word does not enable kerning, ligatures, or hyphenation by default. Kerning adjusts the whitespace between characters in the text. For example, there is a large gap between "T" and "o" when they are next to each other; enabling kerning moves the "o" closer to the "T" and lessens the gap. Word has had this feature for some time but it has to be enabled manually, probably for backwards-compatibility reasons. To enable it, select all the text in your document, go to Font > Advanced, tick the Kerning for fonts checkbox, and put "1" for Points and above. You should notice a tiny shift in your text when you apply it. The example below shows the standard Word spacing, and the blue tint shows the spacing with kerning enabled. Notice the extra space provided before the question mark, but reduced space before the period.



Panegyry sew. Wharf?

Text in Microsoft Word shown by default (gray) overlaid on the same text when kerning is manually enabled (blue)

Word 2010 and up supports ligatures for most fonts, but they are also not enabled by default. Ligatures squeeze two characters together when appropriate. For example, "f" and "i" placed next to each other don't look right because the hood of the "f" almost touches the dot of the "i". Set Ligatures to "Standard only" in the same place where you enabled kerning. However, if you notice that ligatures are still not showing up, such as with typefaces like Times New Roman, you will need to enable All Ligatures because the standard ligatures were placed in the Historical ligatures category for an unknown reason. The below example shows that Garamond does not have any ligatures regardless of the setting, Calibri should be set to standard ligatures, and Times should be set to using all ligatures to get the same effect as standard ligatures.

Garamond	first waffle (off), first waffle (standard), first waffle (all)
Calibri	first waffle (off), first waffle (standard), first waffle (all)
Times New Roman	first waffle (off), first waffle (standard), first waffle (all)

The words "first waffle" written in three fonts, with ligatures off (left), standard ligatures on (center), and all ligatures on (right)

Published text often uses justified alignment, which aligns the text to both margins. However, this will occasionally produce lines with awkward spacing between words to fit it in the line, especially in documents with 2 columns. This makes the text look uneven. Enabling hyphenation allows Microsoft Word to segment words using a hyphen, eliminating the worst cases of awkward word spacing. In the Word toolbar (Ribbon), set Page Layout > Hyphenation to Automatic. If you prefer fewer or more hyphens, you can adjust when they kick in under Hyphenation Options, and you can also specify whether you are okay if

there are multiple lines in a row with hyphenation (a common rule-of-thumb is to avoid 3 lines in a row of hyphenation).

What may not be expected in a country of eternal light? I may there discover the wondrous power which attracts the needle and may regulate a thousand celestial observations that require only this voyage to render their seeming eccentricities consistent for ever. I shall satiate my ardent curiosity with the sight of a part of the world never before visited, and may tread a land never before imprinted by the foot of man. These are my enticements, and they are sufficient to conquer all fear of danger or death

What may not be expected in a country of eternal light? I may there discover the wondrous power which attracts the needle and may regulate a thousand celestial observations that require only this voyage to render their seeming eccentricities consistent for ever. I shall satiate my ardent curiosity with the sight of a part of the world never before visited, and may tread a land never before imprinted by the foot of man. These are my enticements, and they are sufficient to conquer all fear of danger or death

A paragraph of justified text with hyphenation set to "None" (left) compared to with it set to "Automatic" (right)

9 [LaTeX] Enabling Microtype

LaTeX already handles most of the typography nicely, but can be even more slightly improved with the microtype package. It's subtle, but notice how the word spacing and align is a bit smoother when the package is included. Microsoft Word doesn't even have this ability, so this is going above and beyond.

Your introduction goes here! Simply start writing your document and use the Recompile button to view the updated PDF preview. Examples of commonly used commands and features are listed below, to help you get started.

Once you're familiar with the editor, you can find various project setting in the Overleaf menu, accessed via the button in the very top left of the editor. To view tutorials, user guides, and further documentation, please visit our help library, or head to our plans page to choose your plan.

Your introduction goes here! Simply start writing your document and use the Recompile button to view the updated PDF preview. Examples of commonly used commands and features are listed below, to help you get started.

Once you're familiar with the editor, you can find various project setting in the Overleaf menu, accessed via the button in the very top left of the editor. To view tutorials, user guides, and further documentation, please visit our help library, or head to our plans page to choose your plan.

*Paragraph from a LaTeX document by default (left) compared to with the **microtype** package included (right)*

10 Tweaking spacing before units

Some authors put a regular space between a number and its units, like 33 mm, while others omit the space like 33mm. Neither seems ideal, as the former feels like too much spacing and can split the number and units into separate lines, while the latter is technically incorrect according to most style guides. I prefer the thin non-breaking space in these situations, which shows a bit of gap but still makes clear that the number and unit are together. In LaTeX, this is done using the `\` command. In Word, you can insert a non-breaking space under Insert > Symbol > Special Characters > Nonbreaking Space, then changing its width in Font > Advanced > Spacing > Condensed.

standard spaces

But please note that 7.2 oz is actually about 204 g or 204,117 mg.

thin non-breaking spaces

But please note that 7.2 oz is actually about 204 g or 204,117 mg.

Regular spaces (top) versus thin non-breaking spaces (bottom), which are half the width and keep the number and unit together

11 Treating variable labels in equations as text

If you are using variable labels in your equations, you might notice that the spacing is off. Basically, a sequence of letters in equations are often treated as single-letter variables, so they are not kerned like words. So if you must use labels in your subscripts or as variable names, you have to use the `\textrm` command in LaTeX or toggle the Conversions > Text mode for them in the Microsoft Word ribbon. The first example below is without `\textrm`, and you can see there is an extra amount of space before the 's' in "hours", and the f's in "off". The second example uses `\textrm` on the labels, fixing this issue.

$$hours = \mathbb{E}[r(y)|x_{on} = x_{off}]$$

$$\text{hours} = \mathbb{E}[r(y)|x_{\text{on}} = x_{\text{off}}]$$

Variable labels written by default (top) versus specifically set in Text mode (bottom)

Punctuation

12 Selecting the right type of dash

Most keyboards only have one key that looks like a dash, so it commonly serves as multiple punctuation marks that look like a dash. But there are several types of dashes: the hyphen, which is used to separate compound words; there's the (en) dash, representing a range; and there's the em dash to indicate a substantial pause. Those are just the popular use cases, as each type of dash has more varied uses. So rather than use the same hyphen glyph to represent all three, select the appropriate dash under the Insert

> Symbol dialog. In LaTeX, the en dash and em dash can be created by using two hyphens and three hyphens respectively. In the first example below, hyphens are used regardless of the type of dash needed, and in the second example, each dash is represented by its own glyph.

Their bi-weekly hours--only 1-5 pm--were infrequent.

Their bi-weekly hours—only 1–5 pm—were infrequent.

Dashes written with hyphens (top) versus written with the appropriate glyph for en dashes and em dashes

13 Formatting multiple numbered references

Numbered references are inconvenient to handle manually because the numbers referring to each reference can change when new references are added. ACM and IEEE use this numbered reference format, e.g. [23] or [6,11,32], so an automatic cross-referencing setup is helpful. LaTeX is usually used with BibTeX so the `\cite` command handles this automatically, but you can get similar functionality using Word's cross-references. To cite a reference, go to Insert > Cross-reference in the Word toolbar (Ribbon). Make sure the "Reference type" is "Numbered item" and "Insert reference to" is "Paragraph number" before selecting the reference from the list. This should insert something like [23] which is linked to the actual reference. When you update your list of references, your citations will be automatically updated when you Select All (Ctrl+A) and press F9.

Multiple citations show up as [6][11][32] if you just append them together, which is not the preferred format for ACM. To tidy them up in Microsoft Word, right-click on the citation number and Toggle Field Codes and add "`\# 0`" after the reference, e.g. "REF _Ref261299636 \r \h" becomes "REF _Ref261299636 \# 0 \r \h". This removes the brackets around the citation number, and you can add your own brackets and commas to make it look like [6,11,32] while maintaining the reference link. To do the same in LaTeX, simply list all the references in the same `\cite` command, like `\cite{huang2012,papoutsaki2016,qian2021}`. List them in the order they appear in the reference list. See the below example.

Here I'm citing multiple references together [1, 6, 8, 12, 13], as opposed to separately [1][6][8][12][13].

Multiple numbered citations appended together without extra brackets (top), versus the default appearance (bottom)

14 Checking curly and straight quotes

LaTeX users may already be familiar with using a backtick ['] for a curly single quote facing forward, two backticks [``] for a curly double quote facing forward, and in the opposite direction, a single quotation mark ['] for a curly single quote facing backwards, and two single quotation marks ["] for a curly double quote facing backward. However, LaTeX source files are often still littered with double quotation marks ["]

which render straight double quotes, so they should be replaced before submitting. Directional quotes may also be present in the LaTeX source when text is copied from rich text applications to a LaTeX source file; in the past, these would be silently discarded, leading to missing quotation marks. However, with the recent default support for Unicode in LaTeX, directional quotes are now rendered correctly so don't pose the same issue.

Microsoft Word's default settings automatically convert all quotes into curly (directional) quotes. The only caveat is that this is incorrect for abbreviated years, e.g. '08 for 2008, which should use an apostrophe (a straight quote) rather than a directional quote.

“Curly quotes are ‘represented’ by backticks and single quotes,” but "straight double quotes" are sometimes typed by accident and should be replaced. Now I type, “Unicode curly quotes in the source,” which appear correctly in newer LaTeX engines.

Examples of curly quotes and straight quotes used together in the same paragraph

The final step—compiling

If you are compiling the final PDF file yourself, you can tidy up the resulting PDF file.

Images are probably taking up most of the space in the PDF file. Losslessly compress your PNG and SVG images to shrink the PDF file size without losing any quality, ideally testing a few different tools and configurations. There is a drastic range in performance between compression tools, so it's worth comparing a few. Two online compression tools that I've found to work well are compresspng.com for PNG and Vecta Nano for SVG.

You'll want to embed the fonts used in your PDF, so that regardless of what system someone reads it on, it will look the same. One serif and one sans-serif is probably enough for most articles, so this may be a good time to comb through any PDF images and the article itself to consolidate your font choices to save space and get a consistent look. Sometimes, the same font is embedded twice because it is referred to with two different names, like Arial and ArialMT. Review the list of fonts in the PDF properties to see what's included. Many Microsoft Word documents accidentally use a couple of extra fonts in one or two places, maybe from a copy+paste that drags in a similar looking font. [This macro script](#) can generate a list of fonts used in the document and if you see some unfamiliar font, find where it's being used and replace it. Not only does this cut down on the final PDF file size, but also reduces font dependencies.

I tend to get the best results with Save as PDF using the MacOS version of Microsoft Word, and with the default pdflatex for compiling LaTeX source files. The Adobe Acrobat tools seem to not get the links right, and rasterizes and compresses the graphics unnecessarily, sometimes even reducing quality while increasing the file size at the same time. I have also tried compiling to PostScript (.ps) and then using GhostScript to convert to PDF, but the rendering seems to be less crisp.