

Image viewers

and how they can be a mess

Přemysl Eric Janouch '26

Introduction


- I've made my own image viewer
- Does about what I need, how I need, where I need it to
- I use it regularly
- Not trying to “sell” it to anyone
- I want to rewrite it completely

In any case, it lives here <https://git.janouch.name/p/fiv>

Why

There's enough of them already, right?

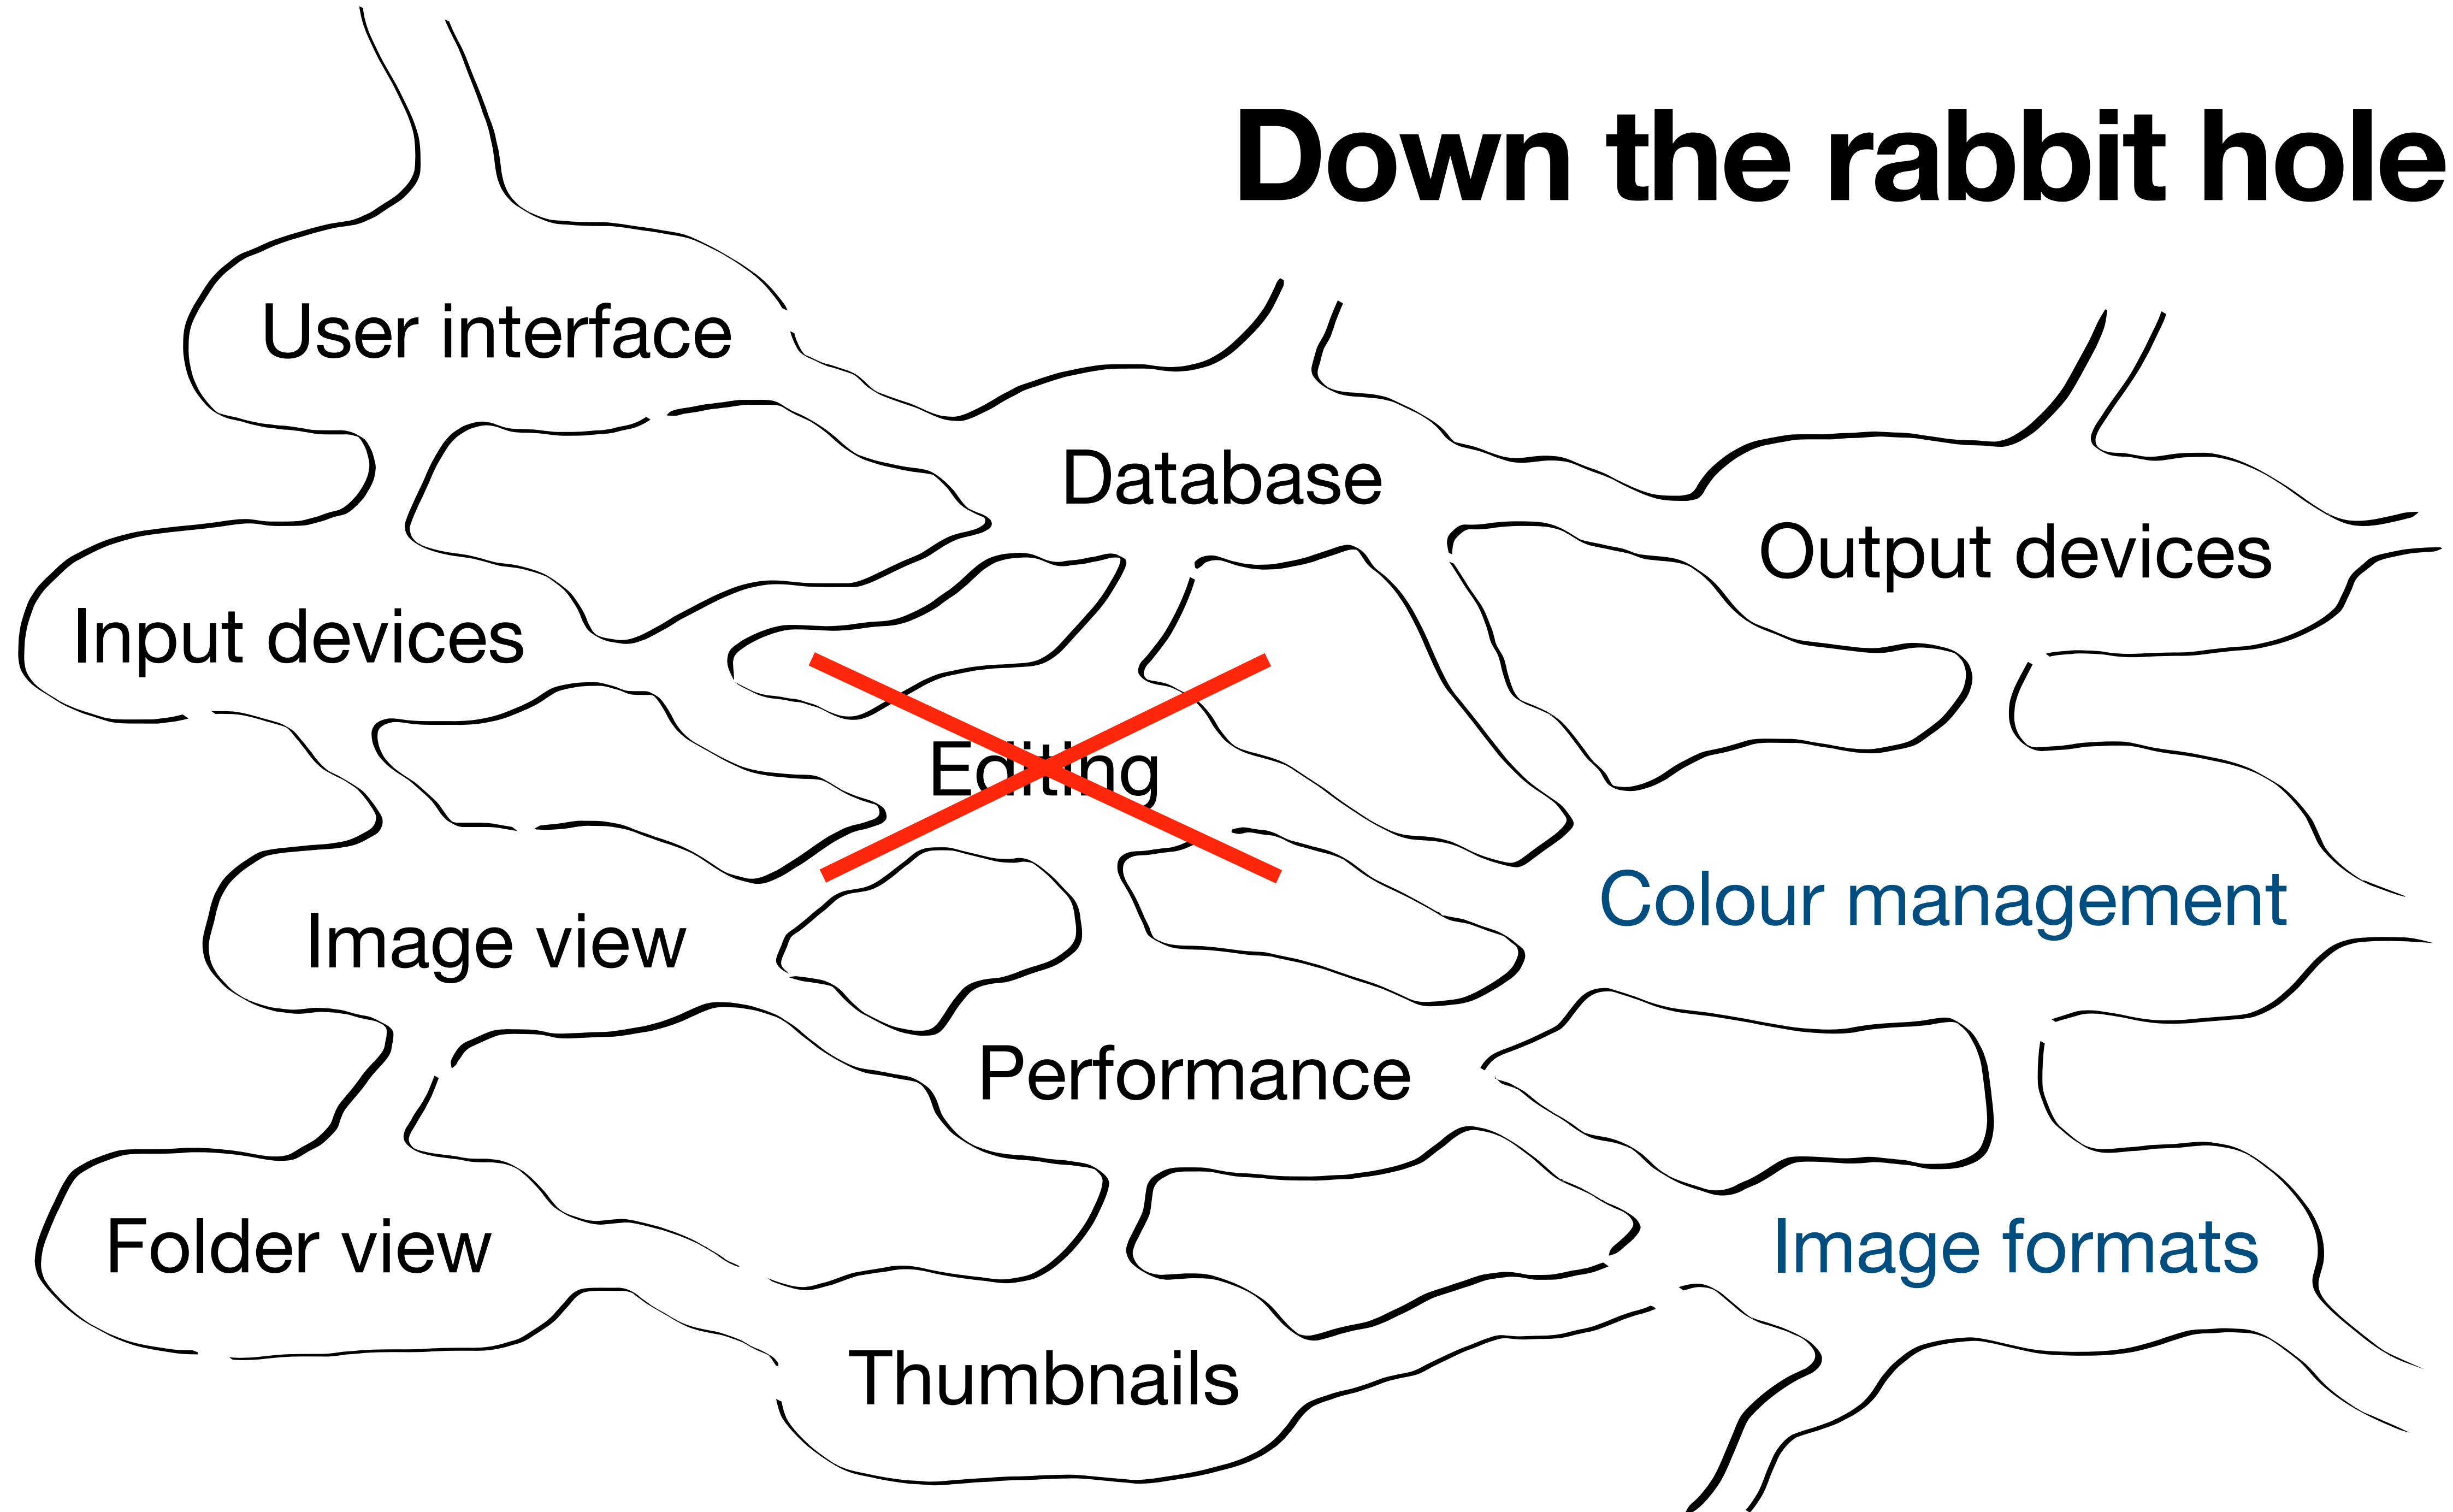
- Actually a complex problem
- Most of the existing ones do at least something wrong
- Replacing Nautilus + Eye of GNOME





<https://www.bbc.co.uk/programmes/p01f9nyx>

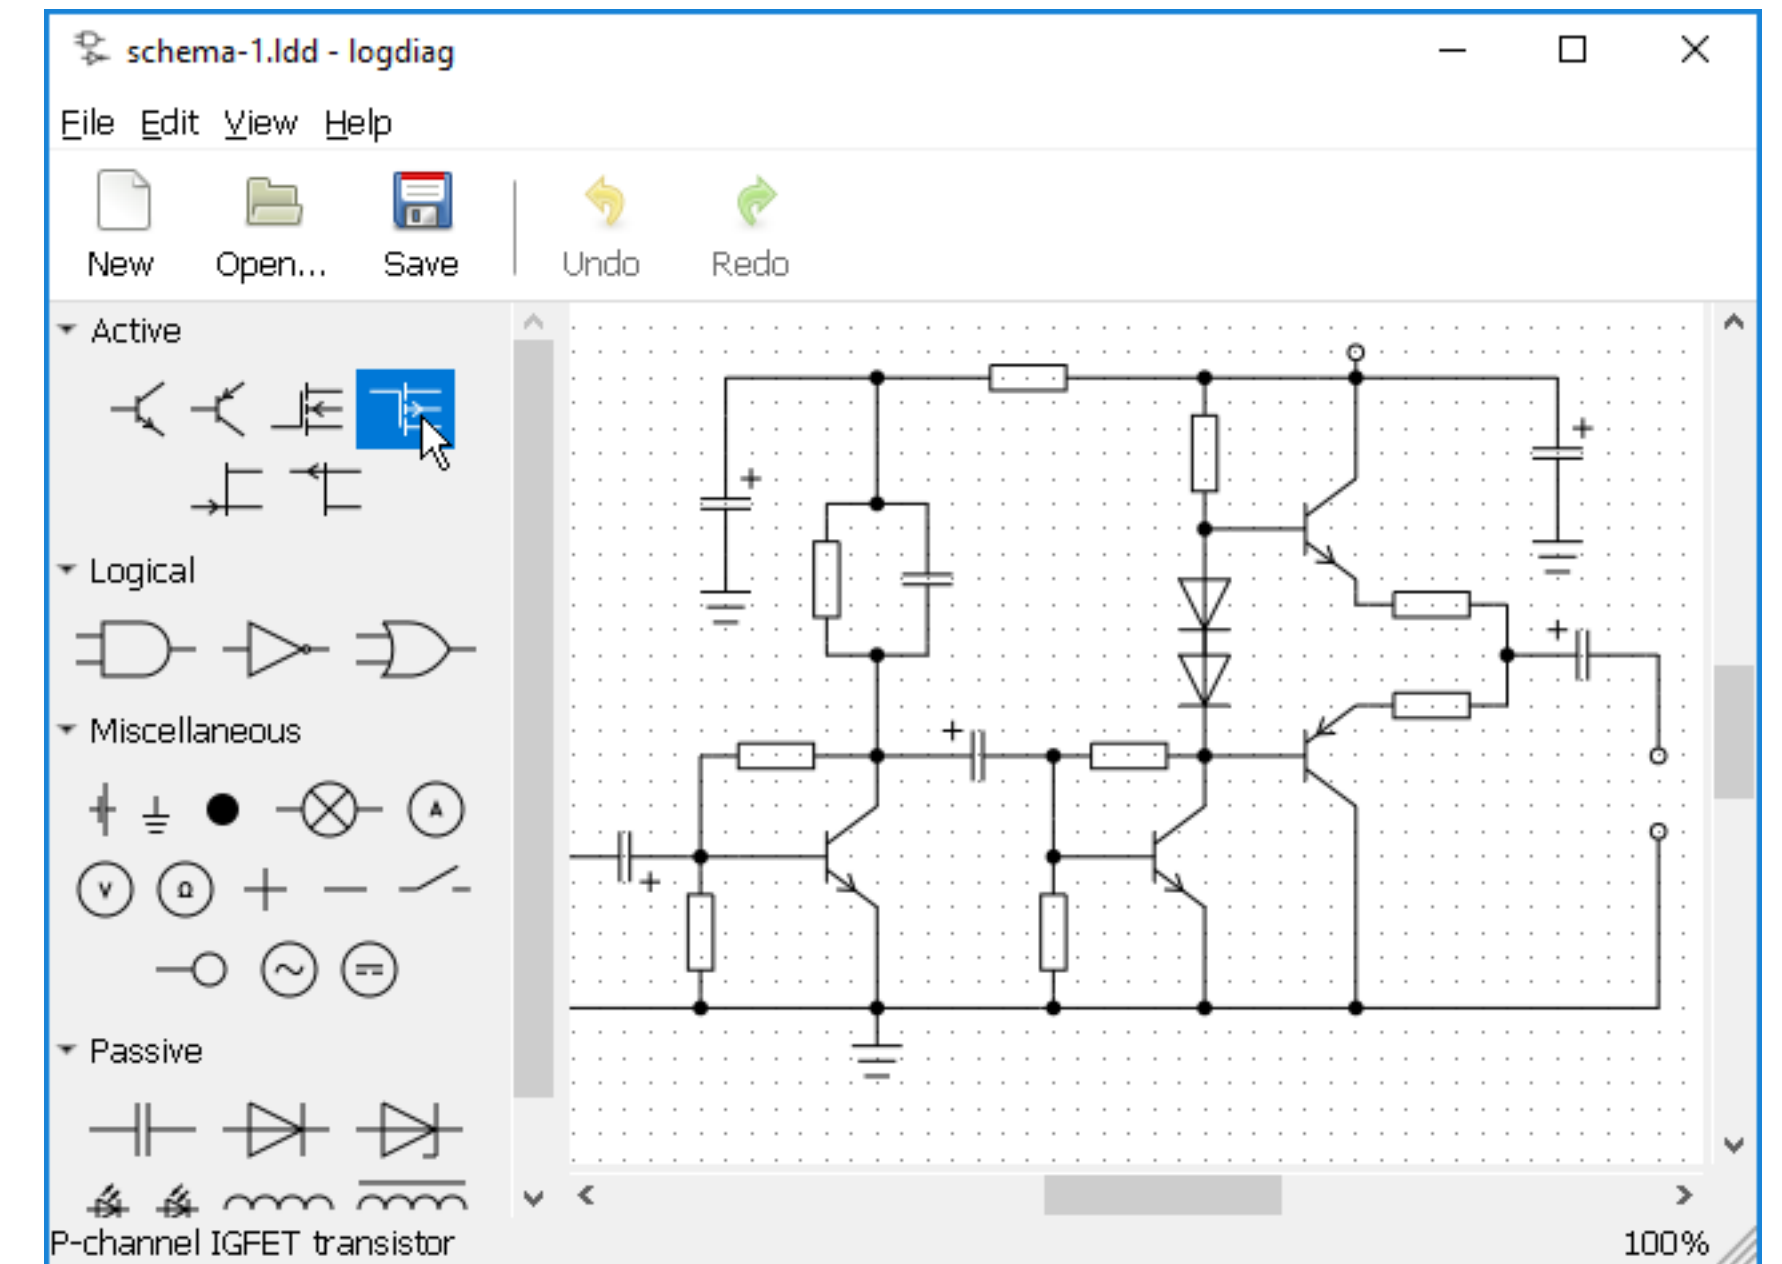
Down the rabbit hole



How I did mine

Not knowing any better

- In some way that I know I can
- GTK 3 GUI toolkit
- C + GObject



- Bad experience with Vala and other language bindings
- Meson build system
- All reasons to remake it from scratch*

* Second-system effect, from The Mythical Man-Month (1975)

On GTK

The good

- Relatively clean design
- Lots of freebies
 - GVfs: SFTP, MTP, trash://, ...
 - File open/save dialogs
 - Multiplatform support

On GTK

The not so good

- Sometimes opaque
- Custom widgets and styling inevitable
- macOS support still poor (try Inkscape)
- Opinionated
- Moving target
 - “The CADT model”

So where's the complexity?

Input devices

- Keyboard
- Mouse
- Touch: gestures
- 3Dconnexion SpaceMouse?

Image view

- Pages
- Animations
- Zoom
- CM
- Interpolation
- Transparency
- Transforms
- ...

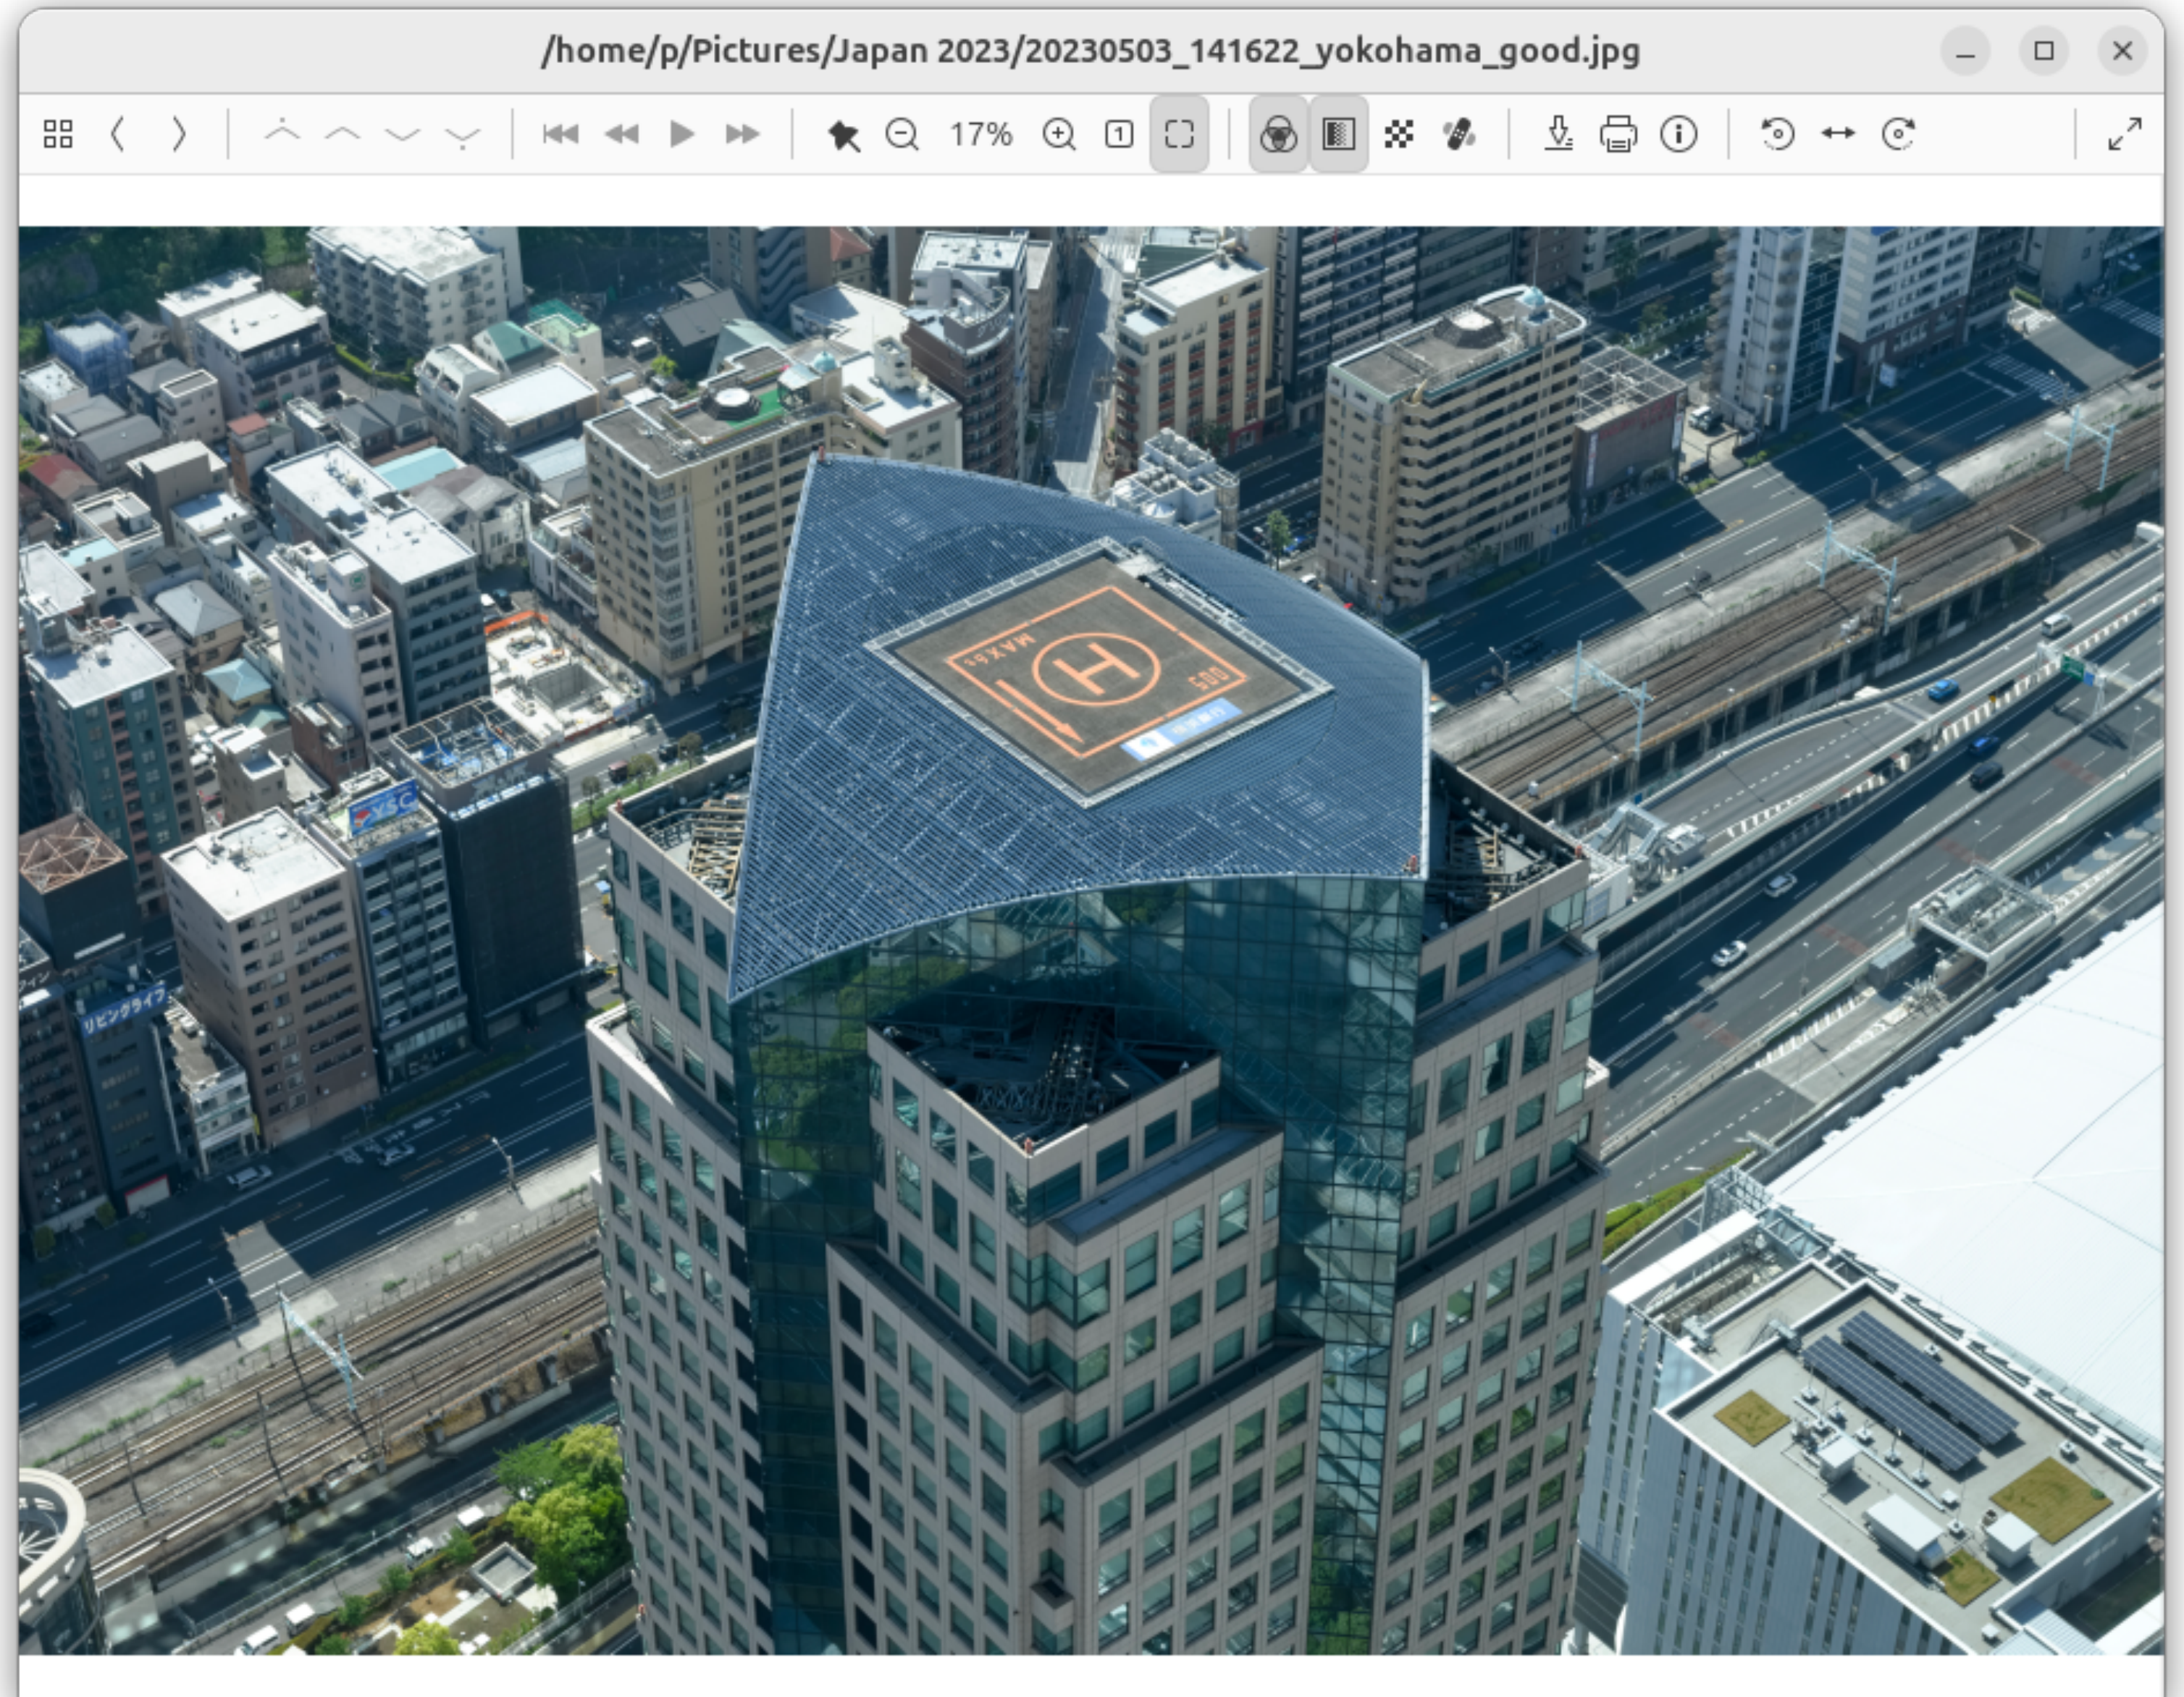
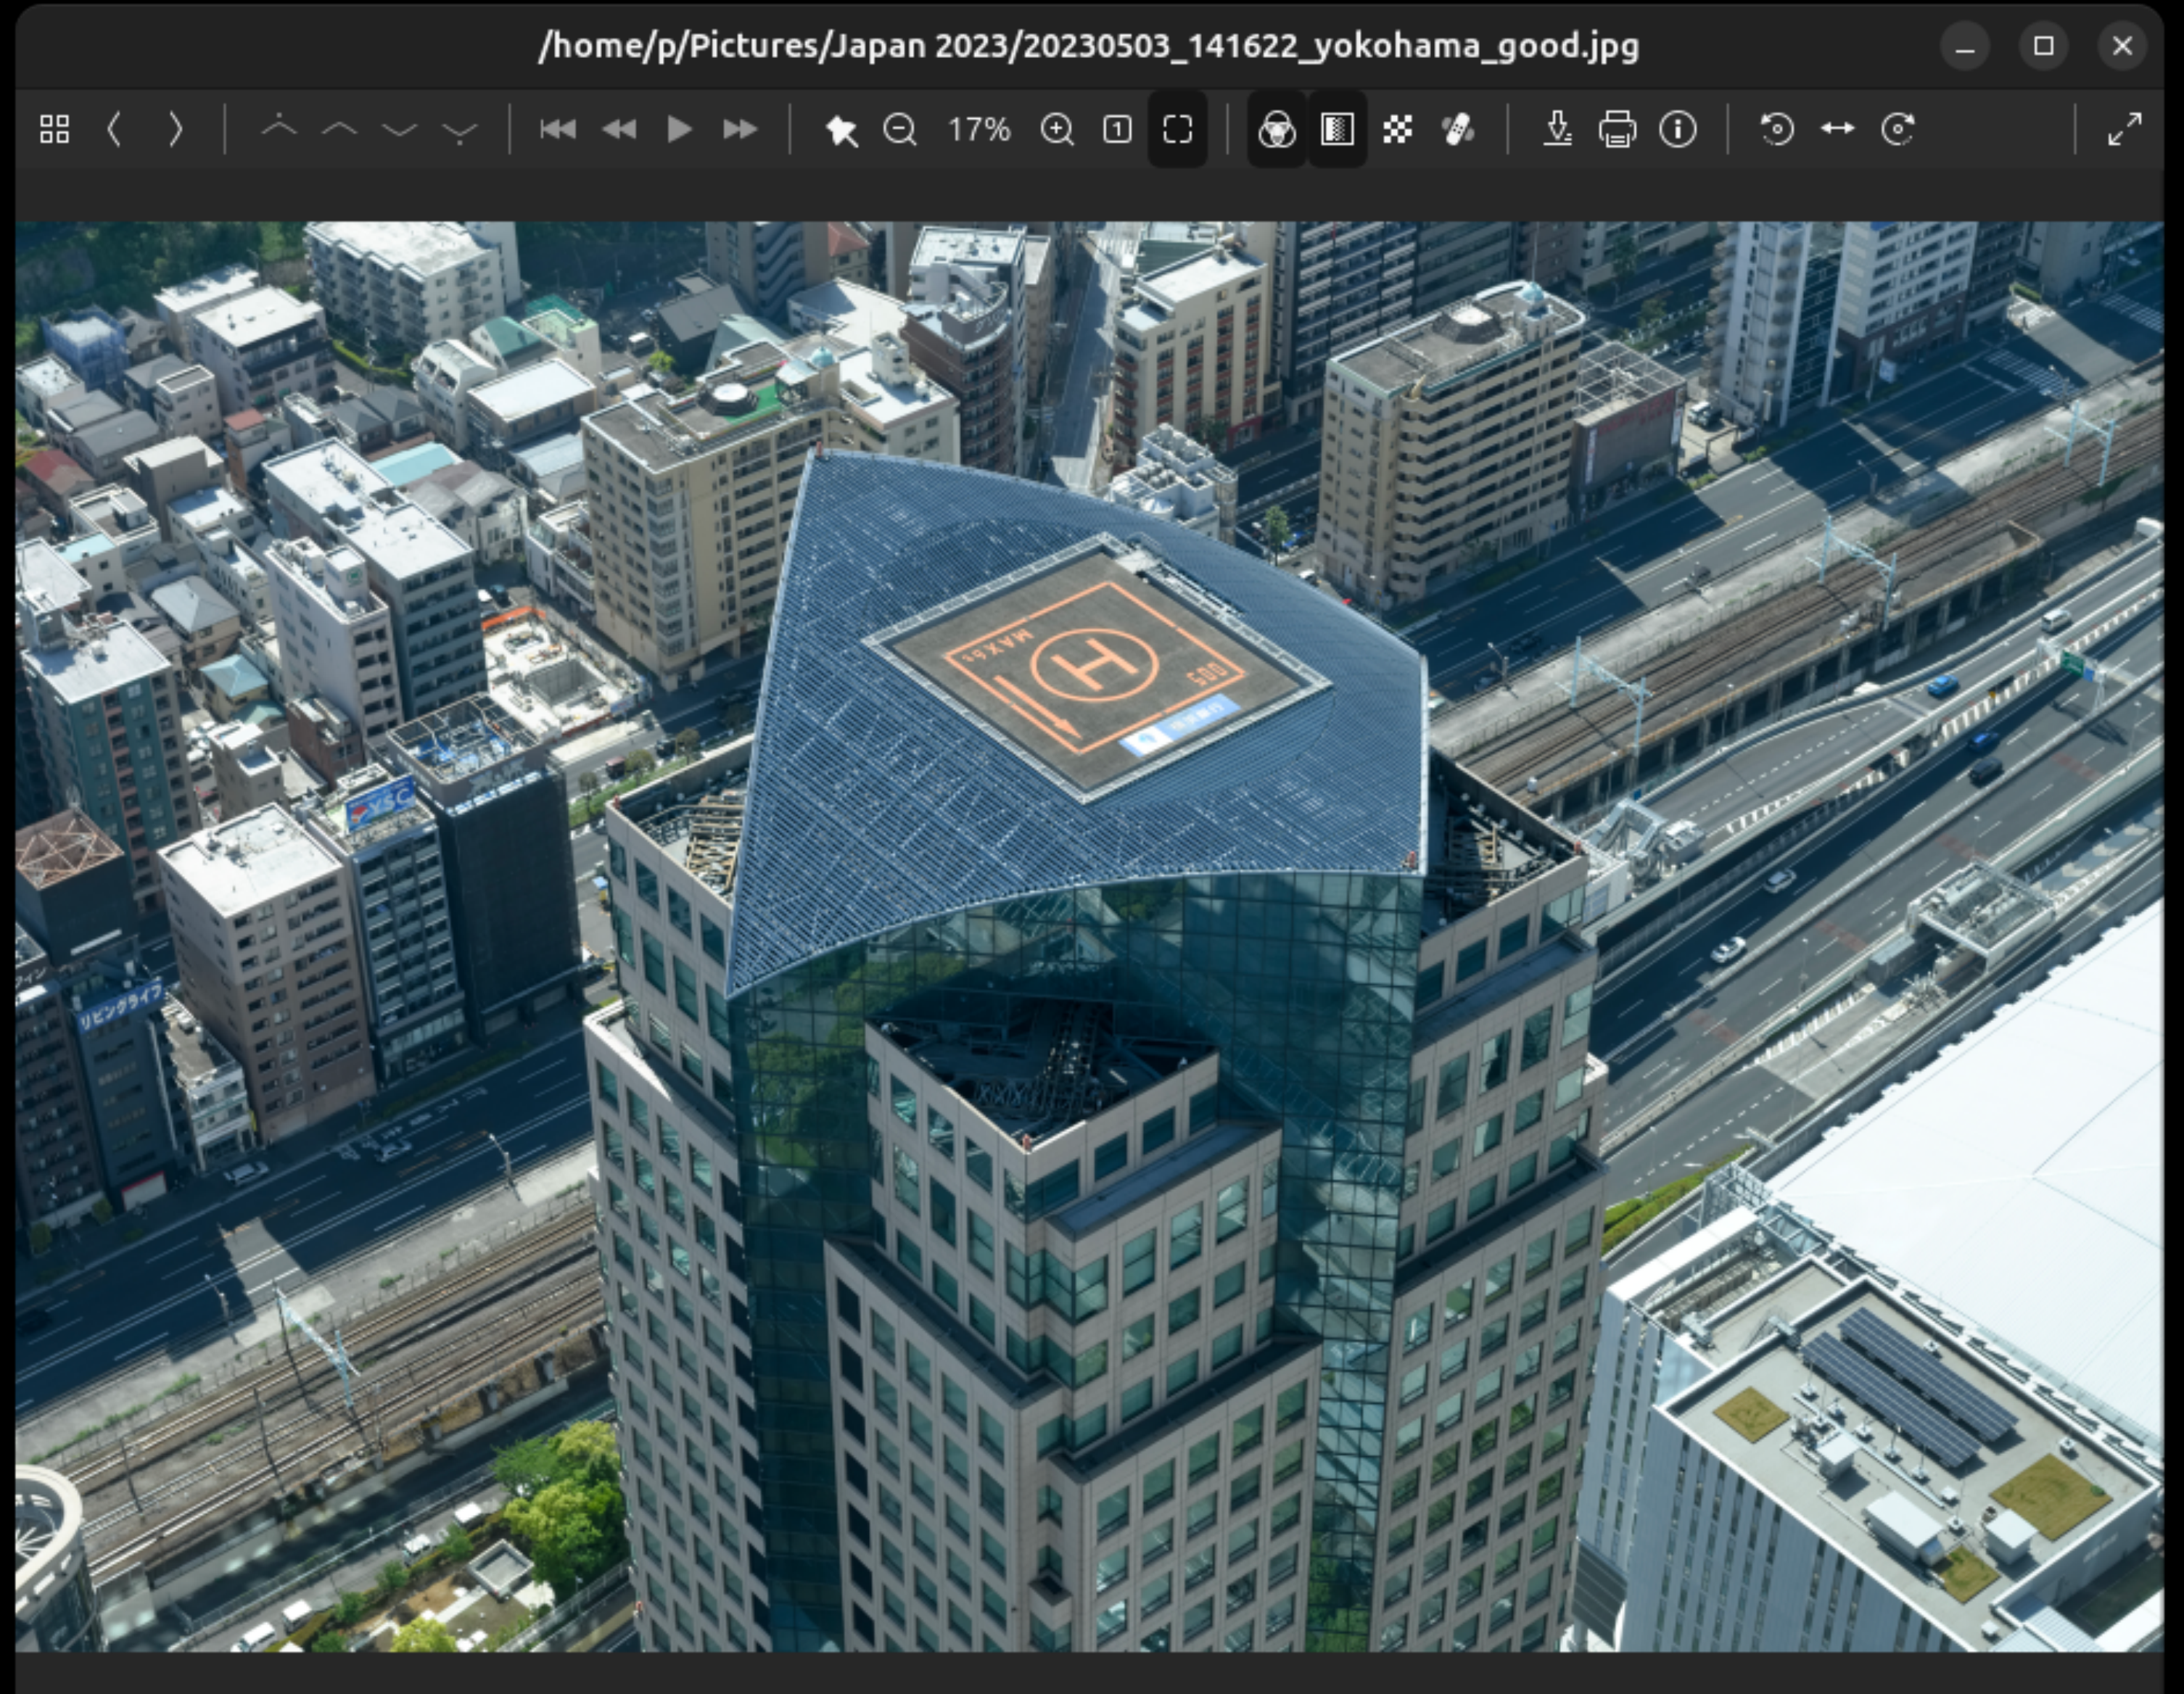



Image view

- Pages
- Animations
- Zoom
- CM
- Interpolation
- Transparency
- Transforms
- ...



Rendering



- Resampling

- Nearest neighbour? Bilinear? Bicubic? Lanczos (sinc)?


$$\begin{bmatrix} 0 & -1 & 0 \\ -1 & 5 & -1 \\ 0 & -1 & 0 \end{bmatrix}$$

- Always do this in linear space!

- Speed

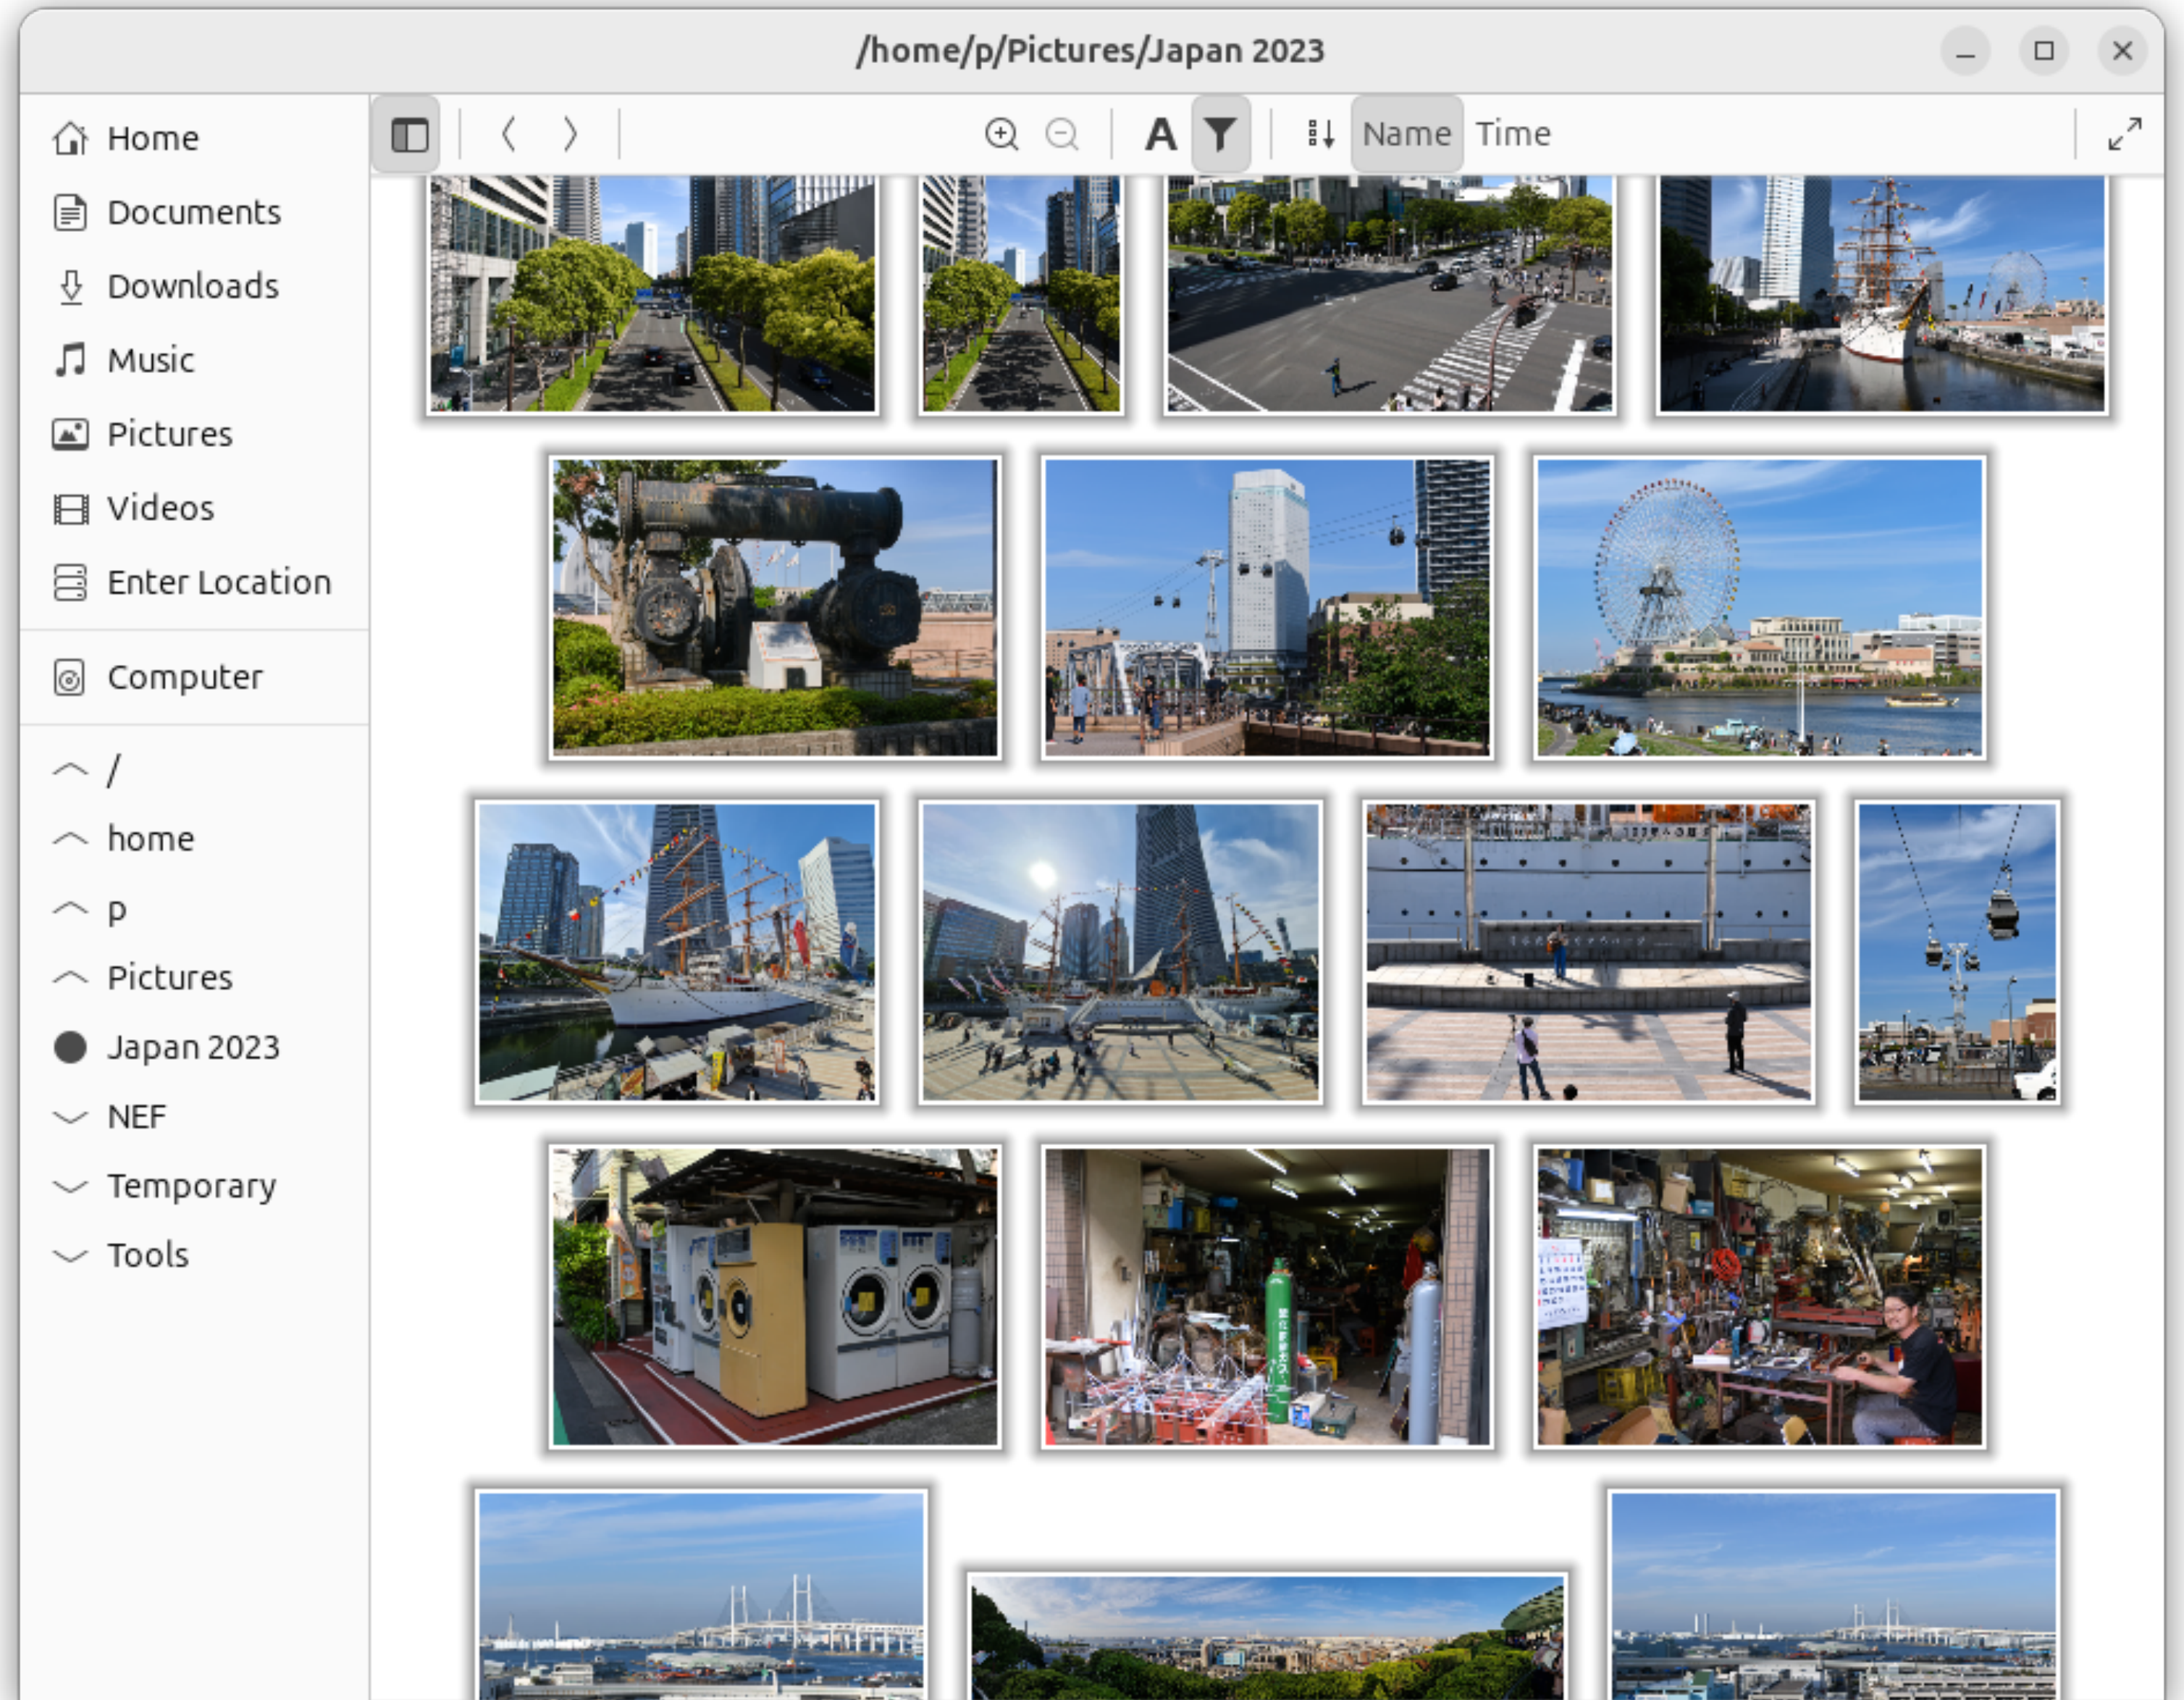
- OpenGL/Vulkan/Metal

- Lots of caching



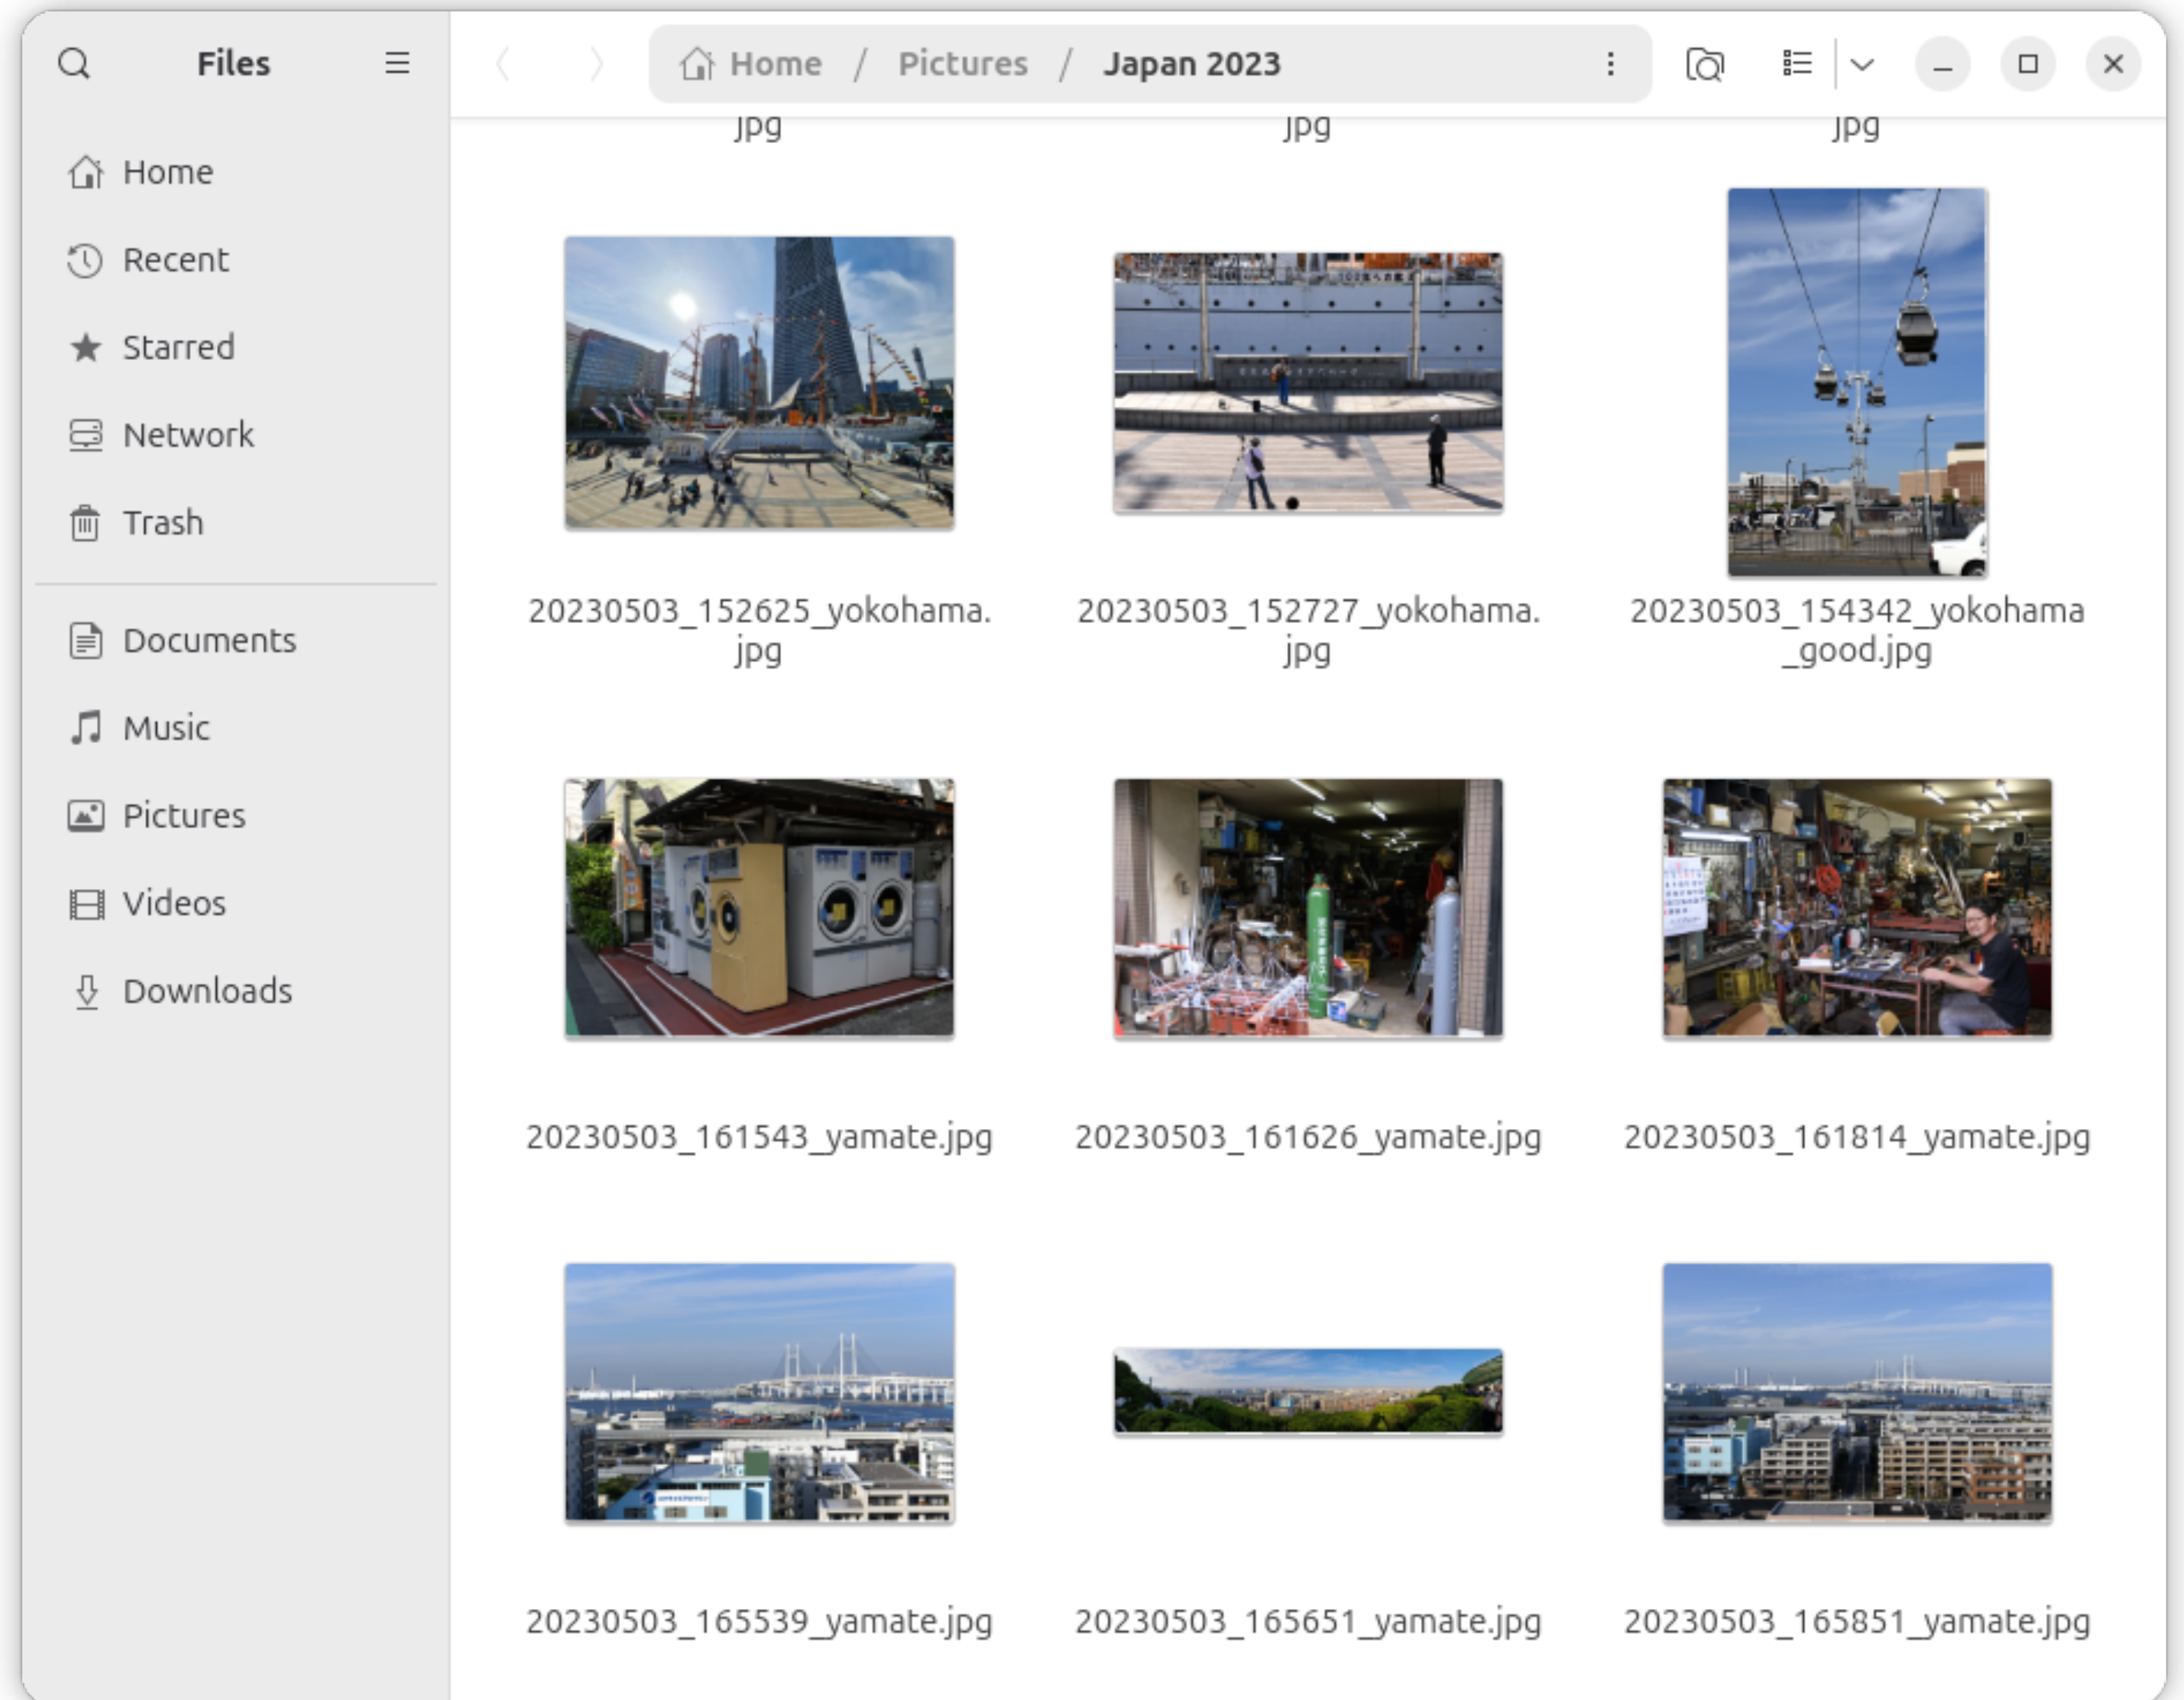
Galleries

- Make the most use of the resources you're given




Galleries

- So not like this




Tiling approaches


Fixed grid




“Flex”



Flex cropped



Squish to fill



Thumbnails

- Thumbnailing is computationally expensive
- Linux normally uses PNGs of grid dimensions, in multiple size steps (see thumbnail-spec)
- Switched to wide format WebP with colour space tagging
 - Enables optional lossy storage and animations



middle finger



sculpture



gnome statue



garden



toad hollow



ottmar horl



bik



A white garden gnome ...
alamy.com



Gnometastic Middle Fi...
amazon.com



Garden Gnome in Hammock - Gnome Statu...
ubuy.co.id



Gnome- Middle Finger – Pr...
presents.com · In stock



1pce 31cm Bobble Head Rude Finger ...
mattblatt.com.au · In stock

Database

How do you find anything?

- Indexing
 - Direct metadata: cryptographic hash, path, modification time, Exif, ...
 - Derived metadata: visual hash, AI classification/feature extraction
 - Either works for finding similar content
- Freeform tagging!
- Edge AI is already here
 - macOS (Photos ~ mediaanalysisd) and Windows 11 (Photos)

Database Storage

- Folder hierarchies and generally ontologies are perfectly futile
 - *[About Semantic Web] They formed committees to form working groups to write drafts of ontologies that carefully listed (in 100-page Word documents) all possible things in the universe and the various properties they could have, and they spent hours in Talmudic debates over whether a washing machine was a kitchen appliance or a household cleaning device. —Aaron Swartz*
- Apple was right

FEEL OLD?

Gen Z Kids Apparently Don't Understand How File Systems Work

"They see it like one bucket, and everything's in the bucket."

By **Dan Robitzski** / Published **Sep 24, 2021 2:54 PM EDT**




Image: Futurism

<https://futurism.com/the-byte/gen-z-kids-file-systems>

Database

Index vs. content vs. metadata

- Not wanting to give up the filesystem has its cost
 - Attach metadata to file? Hard to transfer.
 - Store metadata in a sibling file? Need to keep in sync.
 - Include metadata within file? Not always possible, “ugly”.
 - Indexing by path? Can’t just move it.
- How do you declare an image’s identity?
 - What if it changes very slightly, do you keep extra metadata?

There's more

- I've made another "image viewer"
- Indexes most images on my hard drive by content
 - Careful symlink hell
- Anime AI classification models work surprisingly well for photos
 - ML-Danbooru CAFormer
 - Alternative: freeform description and full-text search (Qwen 3.5)

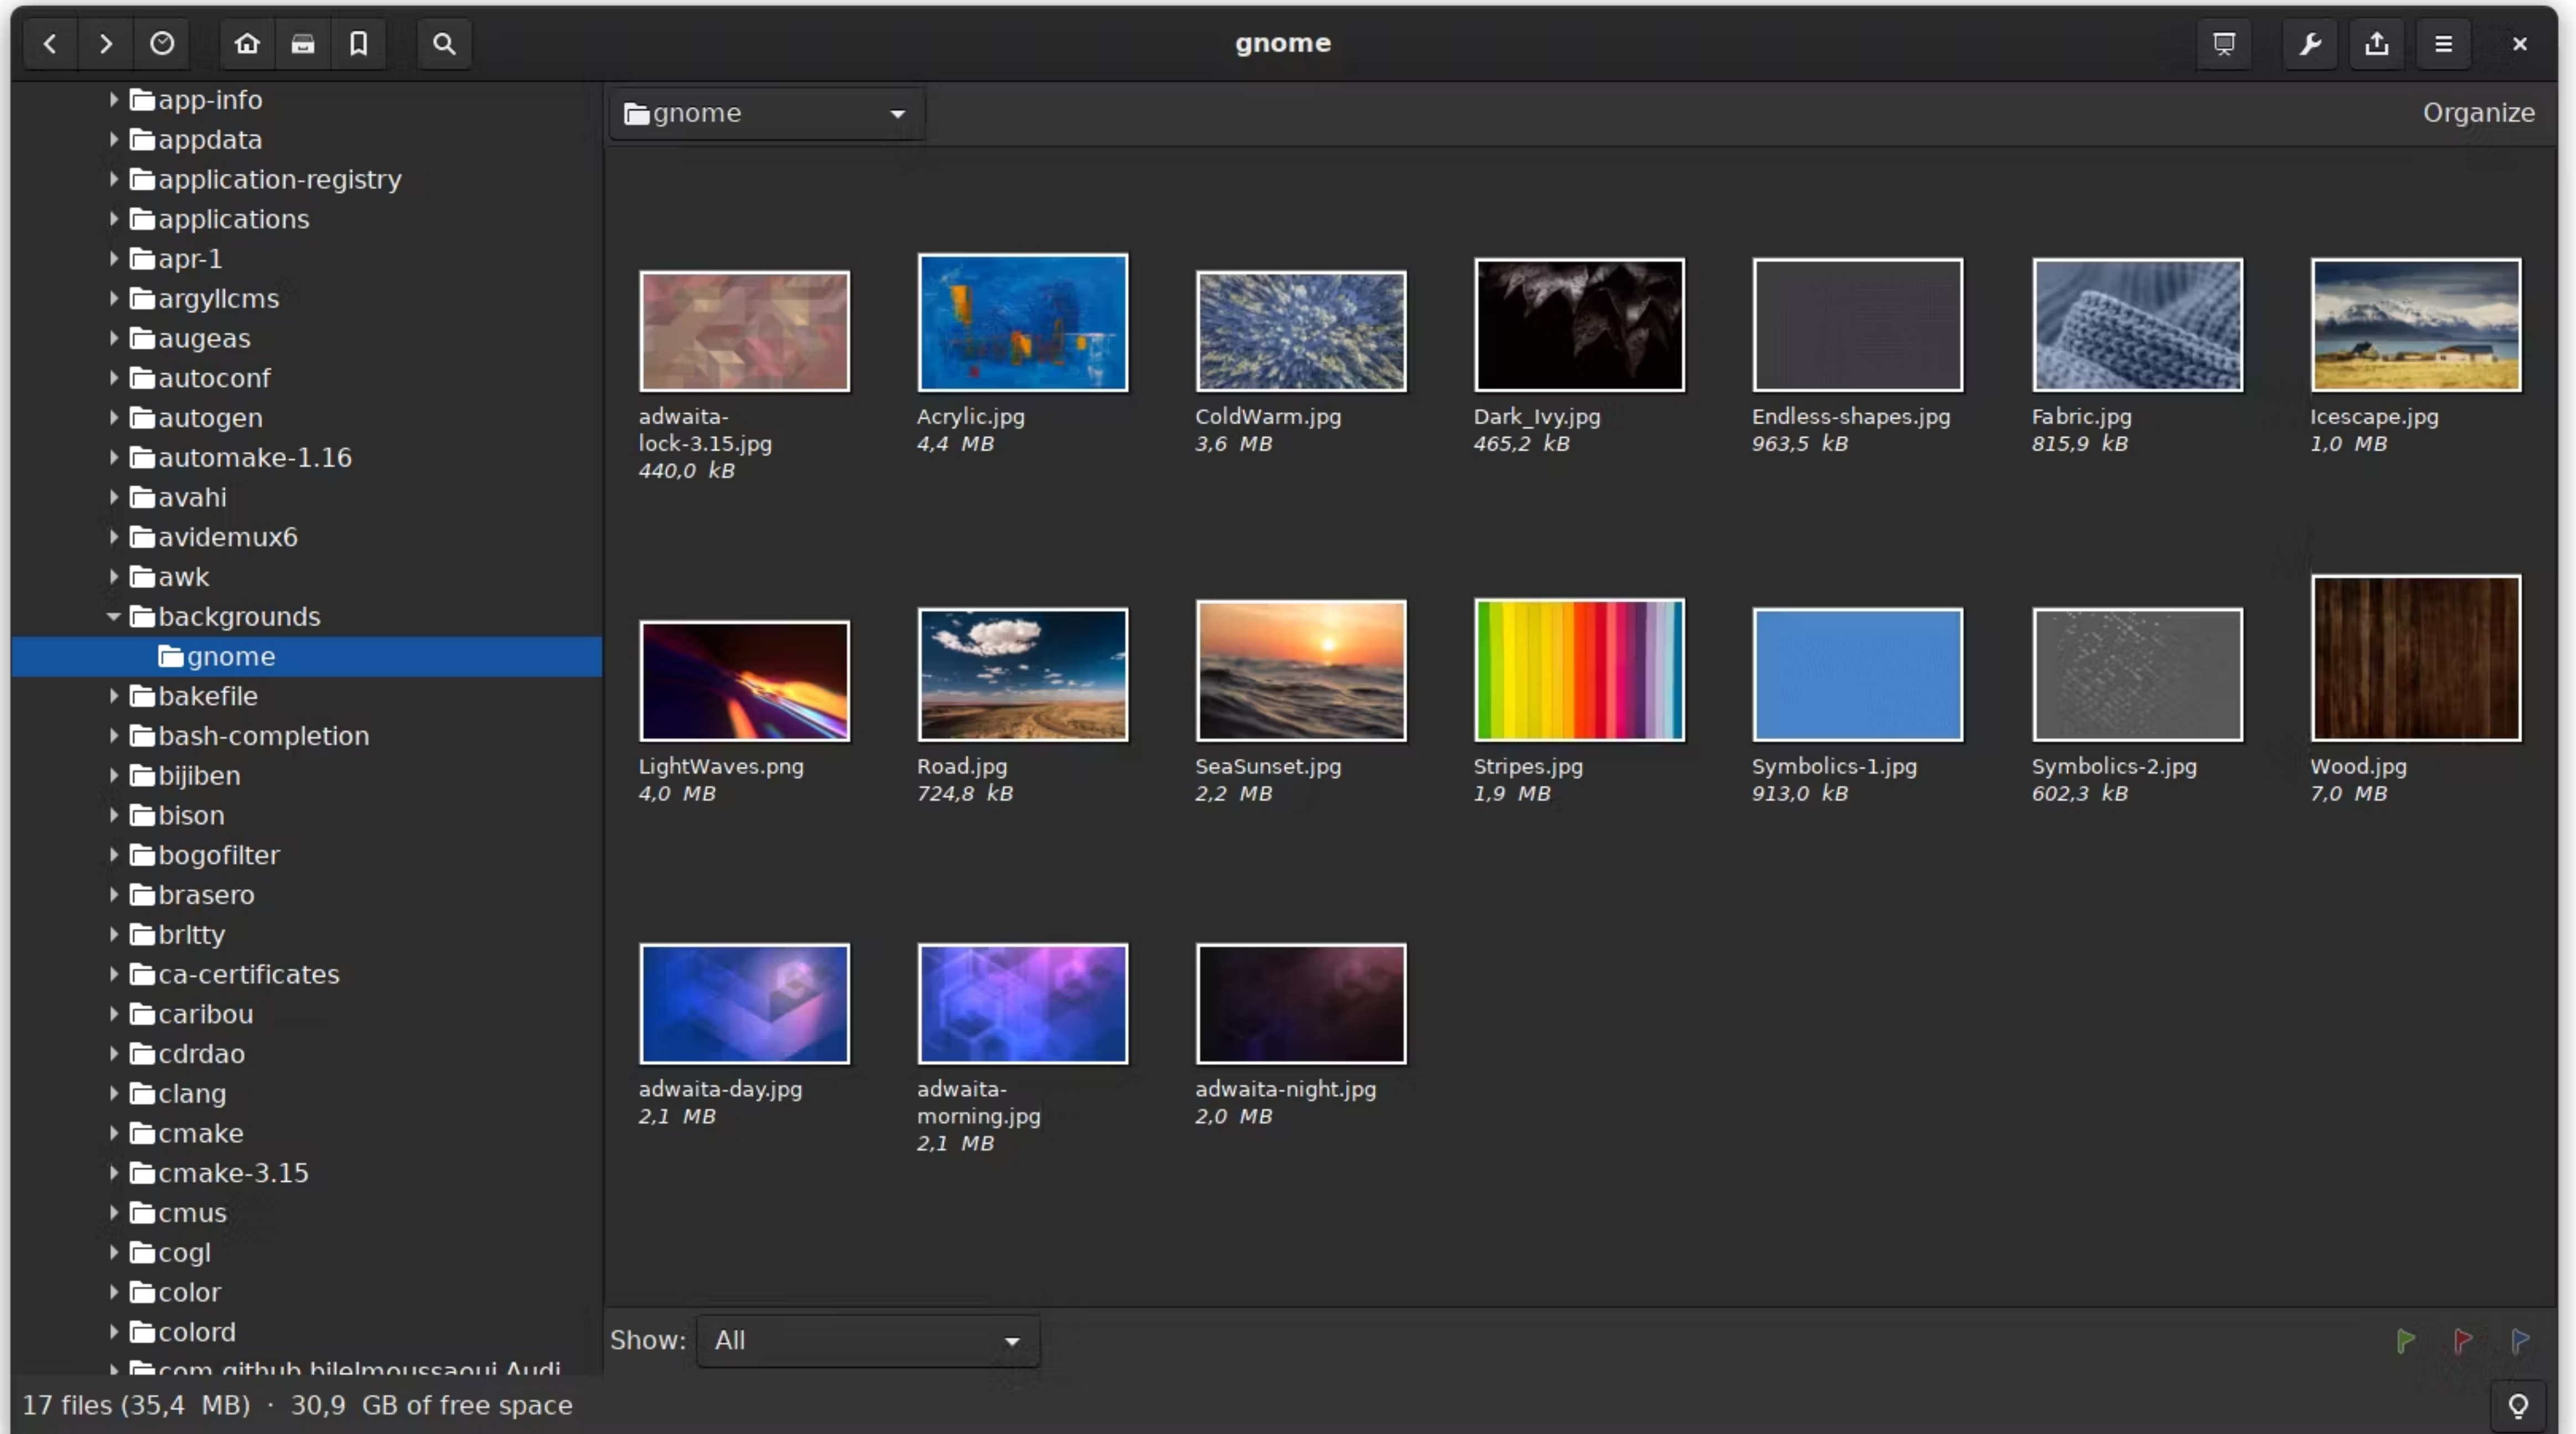
And this lives here <https://git.janouch.name/p/gallery>

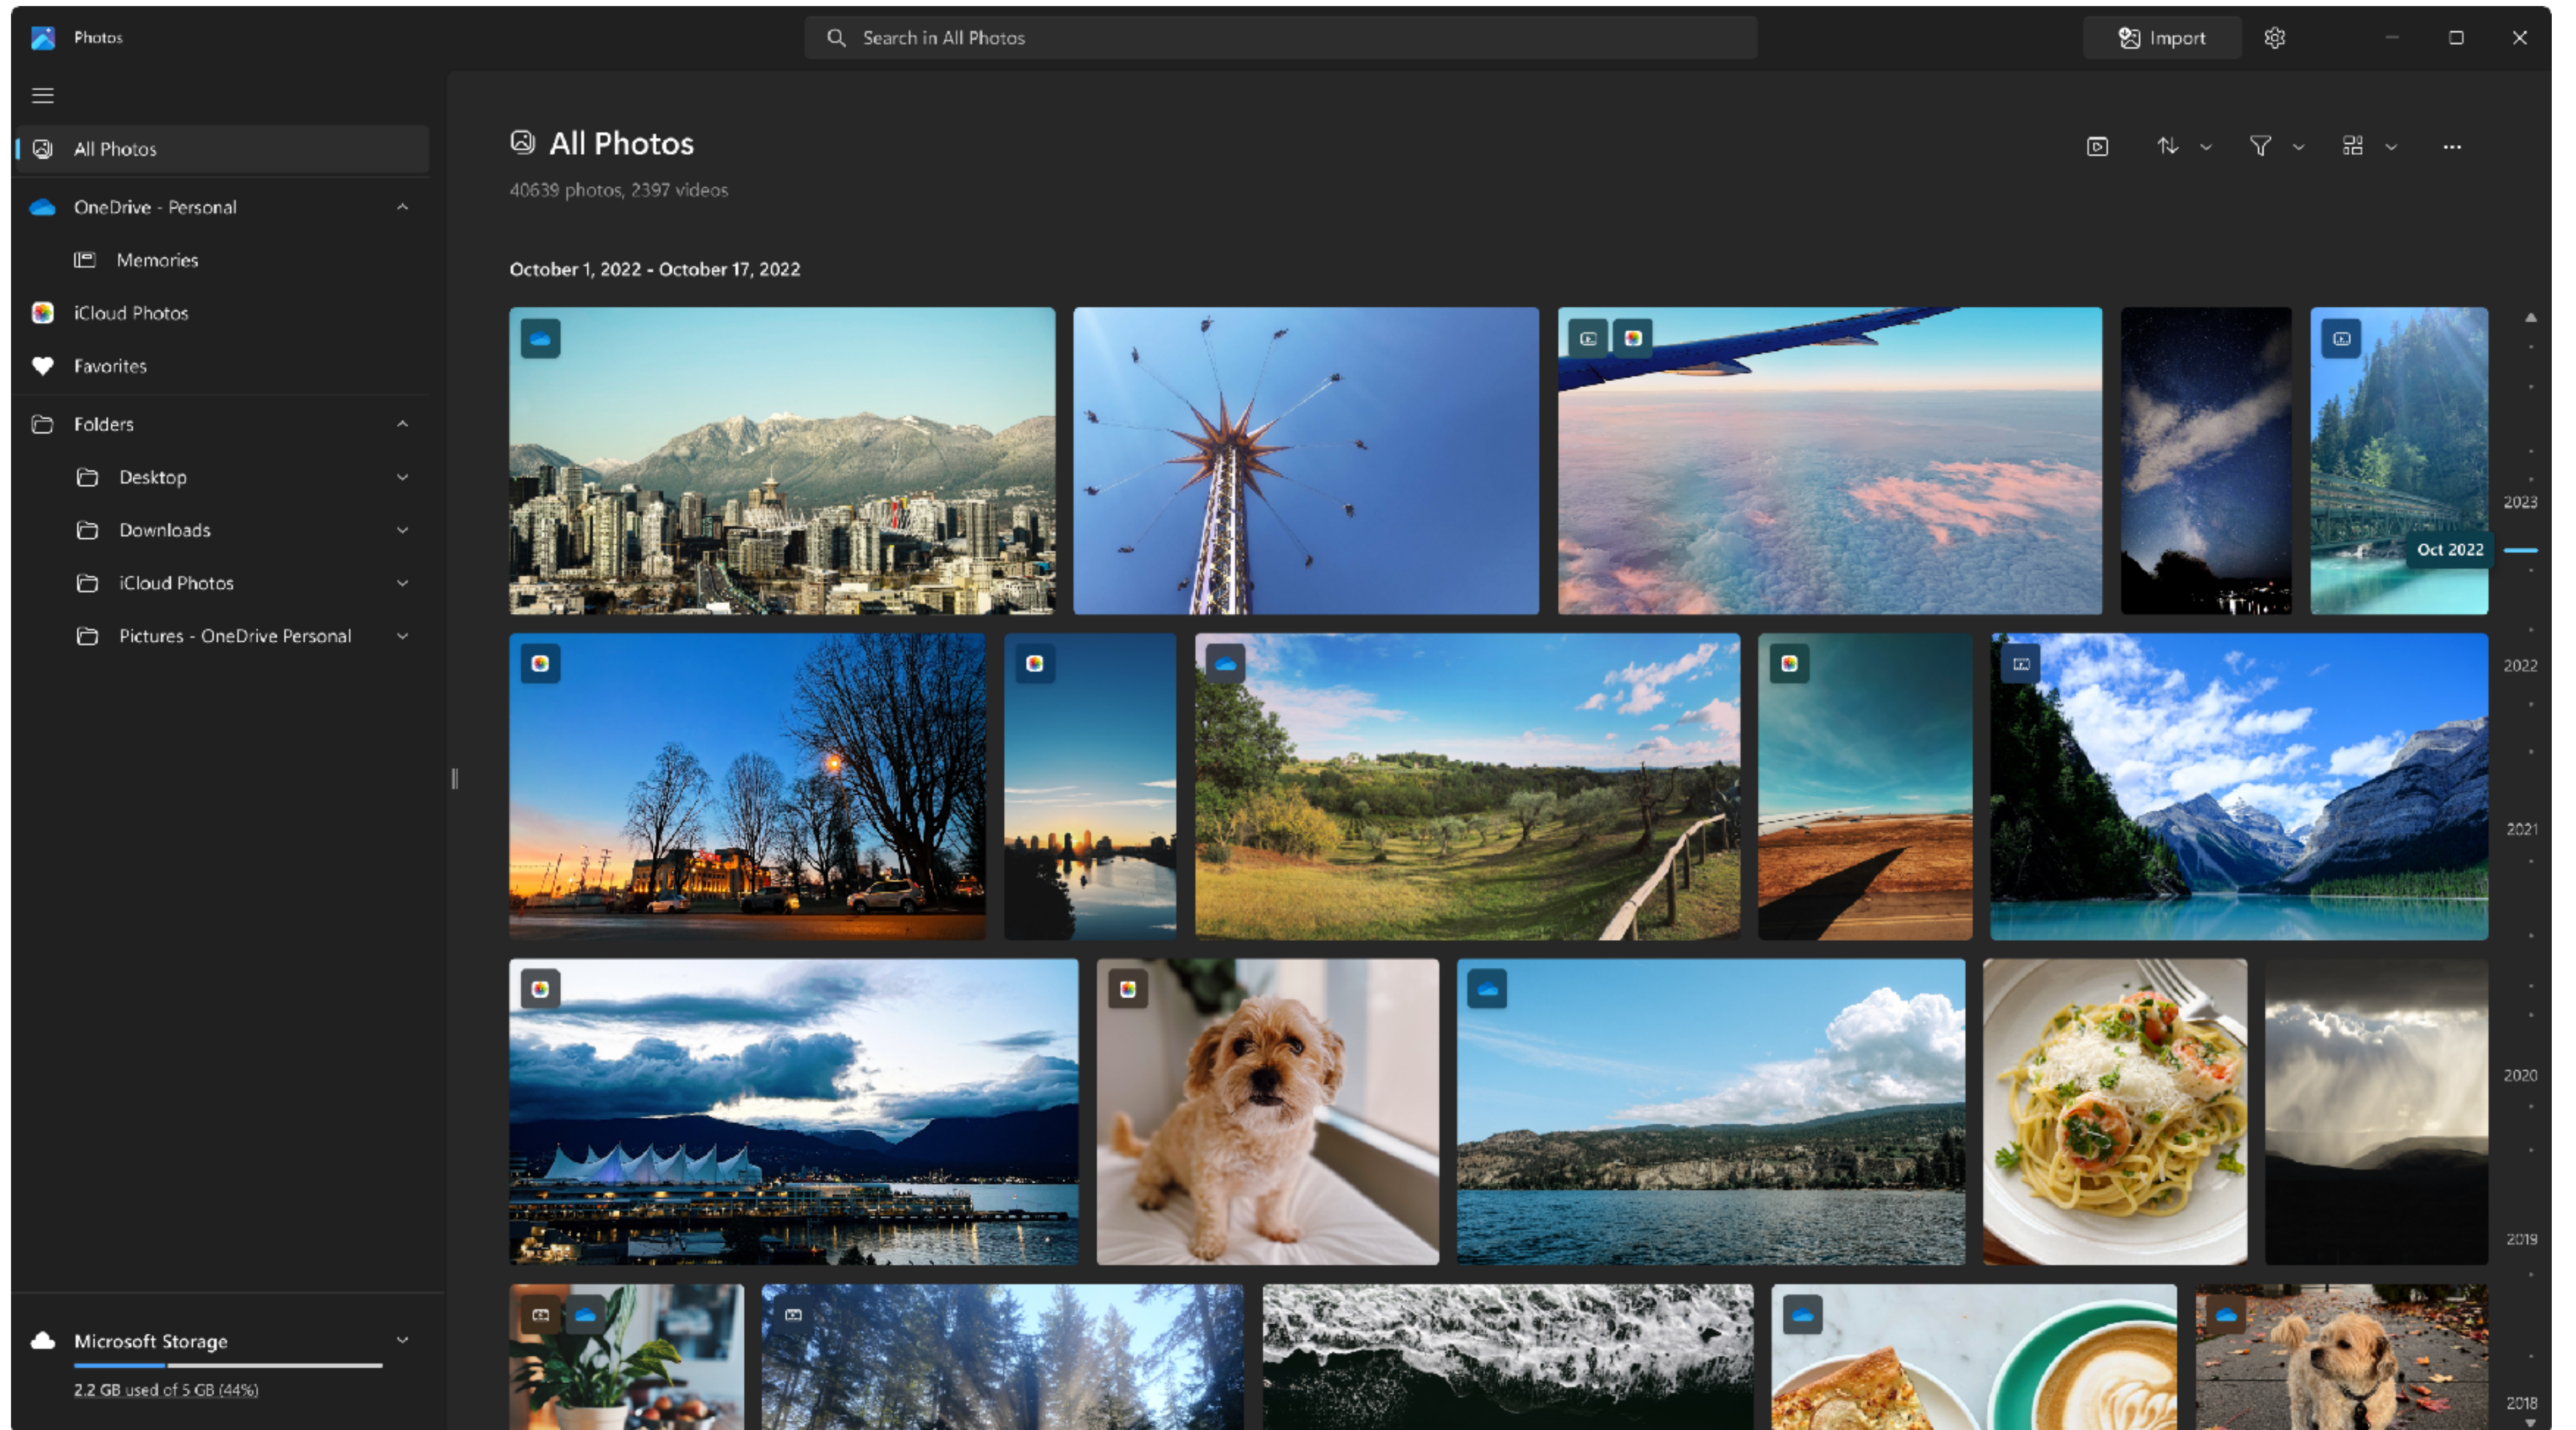
Similar image search

- Identical files
- Content hashes: dHash works better than it has any right to
- AI classification: proximity in N-dimensional space

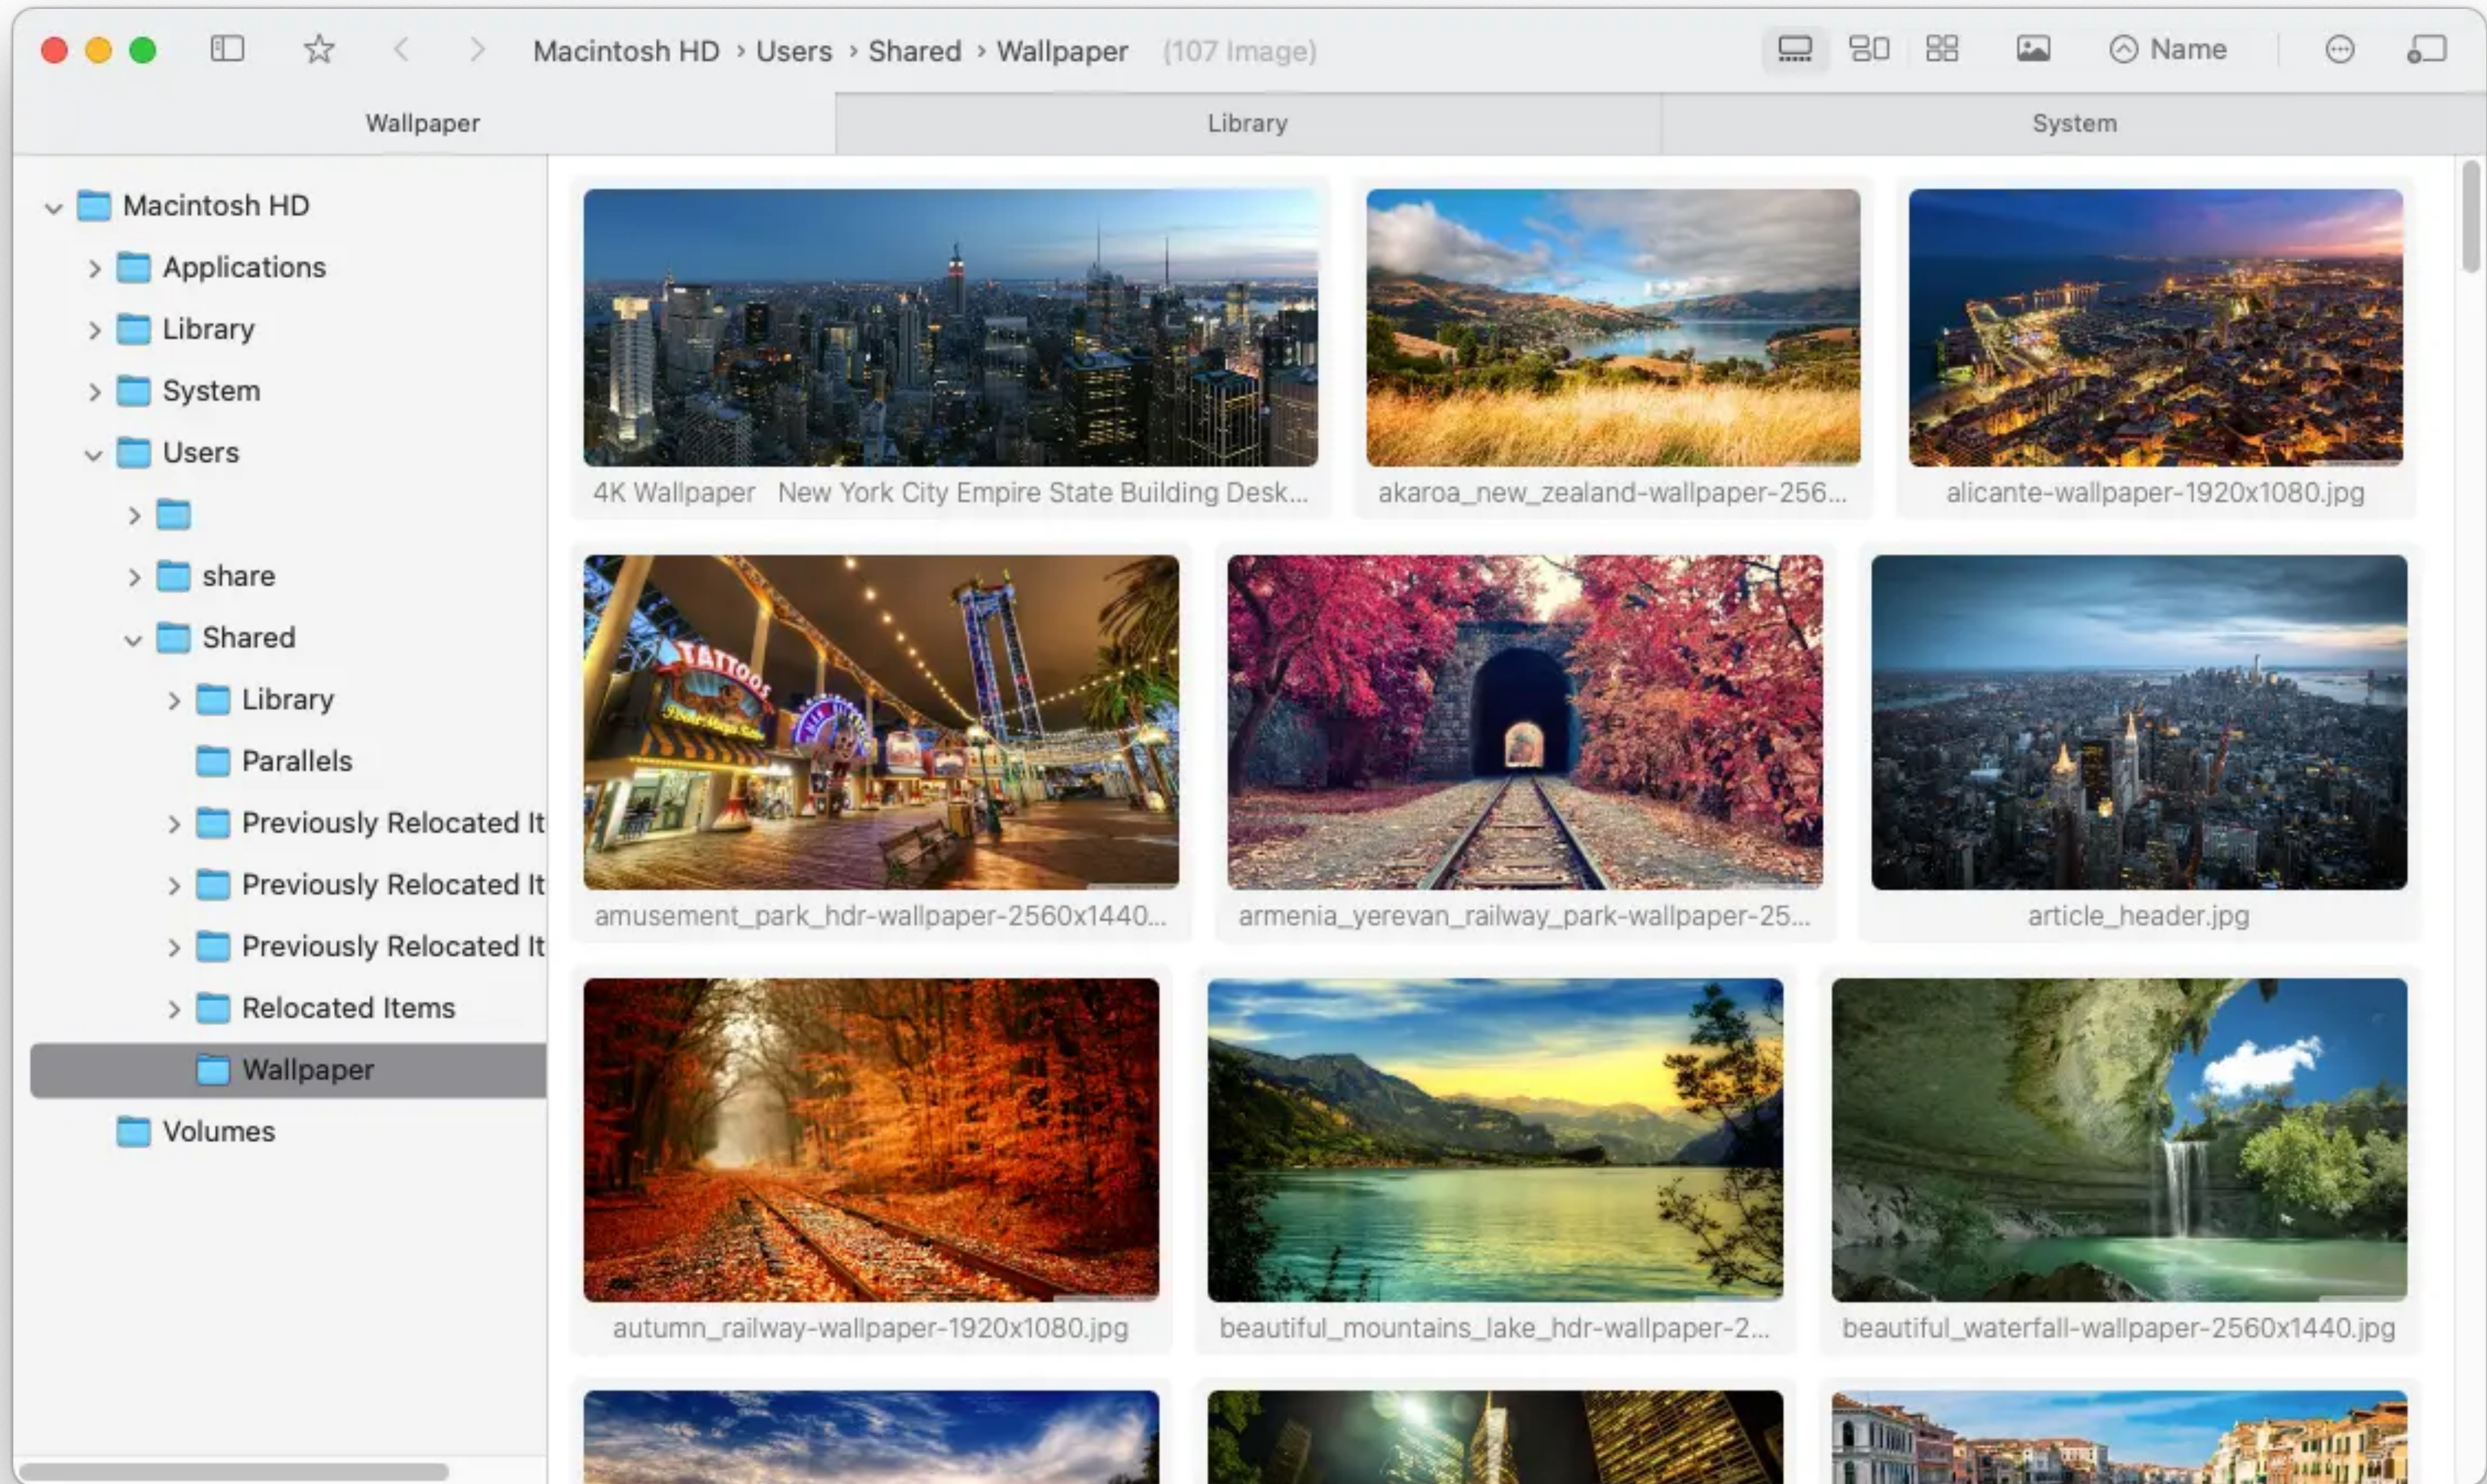
Competition

- Mixed
- Database and AI-wise
 - digiKam: an incomprehensible mess for nerds
 - Apple's Photos.app
 - Windows Photos
 - Eagle: commercial, specialised
- Raw camera photos have Lightroom, Darktable, RawTherapee





Windows Photos



<https://flowvision.app>

Future

The second system

- Merge the projects
- Rewrite it all, probably in something as low-level as SDL3
- Server-client architecture, utilise all CPU cores, utilise the GPU
- I want a tagged database of everything
- Content server?
- Will I find the time again? _(ツ)_/

Thank you

Next talks

- Colour management
- File formats

[While riding on a train looking out the window, I was happy...]

To not be studying source files in an old CVS directory tree for their use of Pango API calls, only to discover that the code was completely broken and that I had just perpetrated these lies into my own source tree for someone else to do the same, an unintentional virus that would infect three or four of us for several days.

—why the lucky stiff, [CLOSURE.pdf](#)