

An Education Explorer for Public NCES Data

Schools, colleges and district boundaries on one map, with 25+ years of history

Pukar Bhandari

pukar.bhandari@outlook.com

2026-10-19

Table of contents

1 Try it	1
2 What is in it	1
3 How it works	2

The National Center for Education Statistics publishes a remarkable amount of open data on schools and colleges, but spread across separate surveys with different file layouts, identifiers and release schedules. Looking up “what has happened to enrollment near this school?” means stitching several of them together.

Education Explorer does that stitching ahead of time and puts the result on a map. It is a sibling of my freight explorer: same idea (a static site, a local data pipeline, DuckDB in the browser), different data.

1 Try it

Open app [GitHub](#)

2 What is in it

Layer	Source	Years
Public schools and districts	NCES Common Core of Data (CCD) + EDGE locations	1997-98 onward
Private schools	NCES Private School Universe Survey (PSS)	every survey, 1989-90 to 2023-24
Colleges	NCES IPEDS	1980 onward in the tables; map layer from 2000
School district boundaries	Census TIGER/Line	2010-11 onward


Each layer has its own shape *and* color, so the map stays readable with all three on. Beyond the map you get:

- A **year slider** with play, and markers for the biennial private-school survey years.
- **Detail panels** with enrollment trends by race/ethnicity, sex and grade, free and reduced-price lunch, staff, finance, graduates, and IPEDS enrollment, awards and admissions.
- **Compare mode**, a table view with CSV export, and shareable URLs that carry the year, selection, filters, basemap and theme.
- Light, dark and system themes, and five OpenFreeMap basemaps.

i What it does not do

It shows data and caveats only. It does not rank schools, compute quality scores or recommend anything, and a Data and methods page documents every source and release.

3 How it works



The first paint comes from small JSON “points” files; DuckDB-WASM loads afterwards and reads only the Parquet byte ranges a query needs. There is no backend, no GitHub Actions and no key-gated service, so the site costs nothing to run.

Updating after NCES publishes a new release is two commands:

```
cd pipeline
uv run python -m nces_pipeline fetch --check-only
uv run python -m nces_pipeline update
```

Line 2

List what is new without downloading anything.

Line 3

Fetch, build, geocode, build boundaries and the dictionary, then validate. It prints a changelog of row-count changes.

NCES changes file layouts often, so the build includes a **schema-drift guard**. When a release adds, drops or renames a column, the build stops and names the nearest known column (+ `NEWCOL` not in the map (nearest known: `OLDCOL`)). Validation also rejects row counts that fall more than 20% from the committed baseline and coordinates that land outside the state’s bounding box.

Nothing in the code is specific to a state; the state shown is set by one config file. The code is on GitHub.