An Accessible (WCAG 2.1 AA) LaTeX Template for Auburn University Theses and Dissertations

by

Franklin Wang

A dissertation submitted to the Graduate Faculty of
Auburn University
in partial fulfillment of the
requirements for the Degree of
Doctor of Philosophy

Auburn, Alabama
May 14, 2026

Keywords: accessibility, WCAG 2.1 AA, tagged PDF, LaTeX, Auburn University

Approved by

Dr. Ash Abebe, Chair, Professor of Mathematics and Statistics
Libby Anderson, Communications & Marketing Specialist III
Dr. Maria Witte, Professor of Adult Education
Dr. Michael Brown, Associate Professor of Mathematics and Statistics

Abstract

This manual introduces the accessible LaTeX template for Auburn University's Electronic Theses and Dissertations (ETDs). While digital accessibility is an essential requirement for Graduate School submissions, adapting complex mathematical and technical typesetting for screen readers has traditionally been a challenge. This guide bridges that gap by leveraging the latest LaTeX tagging capabilities to generate natively accessible, WCAG 2.1 AA compliant PDF documents straight from the compiler. It outlines the core principles of document accessibility, establishes the "Eight Beatitudes" of accessible authoring, and demonstrates these practices through a fully tagged mathematical showcase. By hiding the complexities of PDF tagging under the hood, this template empowers students to focus entirely on their research and meet Graduate School formatting standards with their usual workflow.

Artificial Intelligence (AI) Use Disclosure Statement

In the preparation of this thesis / dissertation, no Artificial Intelligence (AI) tools were used.

Acknowledgments

I am grateful to Dr. Ash Abebe for his strong support of this LaTeX template and for his care for students. Long before this project began, at a time when I was facing serious difficulties with my vision, he stepped in and helped me find a way forward. I will always remember that help.

I am also grateful to Libby Anderson, Dr. Maria Witte, Sadaf Jamalzadeh, Jay Neese, Soran Salimiyan, and Mikael Arvola, each of whom helped this project in important ways. Libby was always someone I could turn to when I needed help; her steady support, quick responses, and practical advice helped this project take shape. Dr. Witte read carefully, asked useful questions, and offered suggestions that made the template clearer and stronger. Jay, Sadaf and Soran shared their professional knowledge whenever I needed guidance, and their advice helped me solve problems that I could not have handled alone. Mikael Arvola of the University of Illinois Urbana-Champaign Graduate College provided detailed guidance and answers to my many questions. I am grateful for the time, care, and expertise he generously shared.

Most of all, I want to thank you, the people using this template now. By taking the time to make your documents more accessible, you are helping make academic work easier to read, easier to use, and more open to everyone.

Contents

List of Tables

List of Figures

List of Abbreviations

ETD Electronic Thesis or Dissertation

MathML Mathematical Markup Language

PDF/UA PDF/Universal Accessibility (ISO 14289)

WCAG Web Content Accessibility Guidelines

Chapter 1

Why We Do This: Writing for Everyone

> “Literature is my Utopia. Here I am not disenfranchised. No barrier of the senses shuts me out from the sweet, gracious discourse of my book friends.” — Helen Keller, *The Story of My Life* [Kel03]

Years from now, a student may open your thesis at two in the morning. You may never meet that reader. One reader may see the page. Another may listen through a screen reader. Someone else may follow the text with her fingertips on a braille display.

Whether your words can reach them is decided now. It may cost an afternoon, a good coffee, and a few choices made while you write. This chapter is about why that afternoon matters. It begins in Alabama, beside a water pump.

1.1 A word at the pump

Helen Keller was born in 1880 in Tuscumbia, Alabama. When she was nineteen months old, an illness took away her sight and her hearing. For the next five years, she later wrote, she lived “at sea in a dense fog” [Kel03].

Then came an April morning in 1887. Anne Sullivan took her to the family’s water pump. Water ran over one of Helen’s hands. Into the other, Sullivan spelled w-a-t-e-r. Again and again.

Something opened.

“Everything had a name,” Keller later remembered, “and each name gave birth to a new thought.” The child did not become intelligent on that morning. She had always been intelligent. What changed was the form in which the world reached her.

Keller went on to Radcliffe College, where in 1904 she became the first deafblind person to earn a Bachelor of Arts degree. She wrote more than a dozen books. But before the books, there was a word. Before the word, there was a hand waiting to understand.

Figure 1.1: Helen Keller (1880–1968), author and advocate, born in Tuscumbia, Alabama. (Photograph: Los Angeles Times, c. 1920.)

1.2 A mother's voice

Lev Semyonovich Pontryagin was born in Moscow in 1908. At fourteen, a stove explosion took his sight. His mother, Tatyana Andreyevna, was a seamstress. She had no mathematical training. But she sat beside him and became his reader.

For years she read mathematical papers aloud to him. When she met an unfamiliar symbol, she described its shape. When the printed notation had no natural sound, they gave it a name of their own [OR14].

Pontryagin duality, Pontryagin classes, the Pontryagin maximum principle—these now belong to mathematics. But before they entered textbooks and lecture halls, they passed through a mother’s voice.

Access, for Pontryagin, took the form of someone sitting beside him and telling him, line by line, what was on the page.

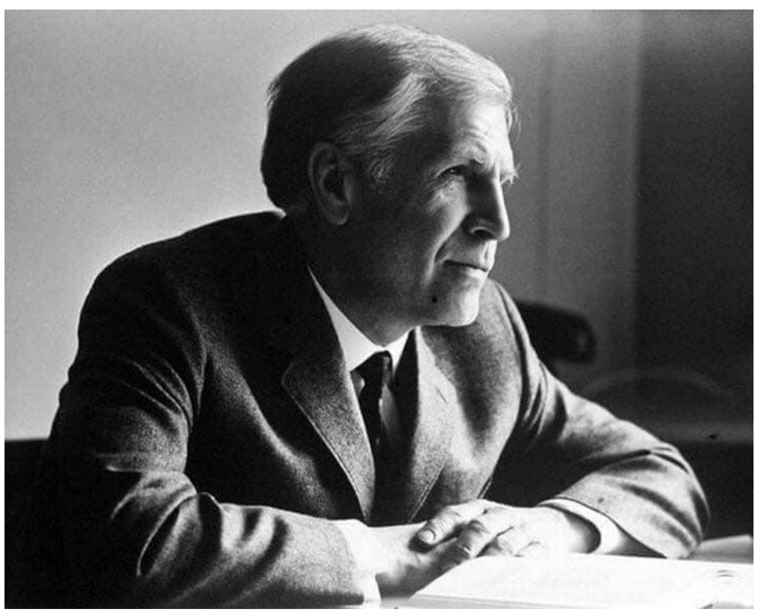

Figure 1.2: Lev Semyonovich Pontryagin (1908–1988), who lost his sight at fourteen. (Photograph: unknown author, Moscow State University history of mathematics collection; CC BY 4.0, via Wikimedia Commons.)

1.3 When the page went dark

Leonhard Euler lost the sight of his right eye around 1738. He is said to have remarked, “Now I will have fewer distractions” [Cal16]. Later, a cataract in his left eye took most of the rest. From 1766 onward he was almost totally blind.

Still, the work continued. Nearly half of his roughly eight hundred and fifty writings were produced after 1766, dictated to assistants. He calculated in memory. He composed without seeing the page. The formulas came from his mind, but they reached the world through other people’s hands.

Figure 1.3: Leonhard Euler (1707–1783), in the 1753 pastel by Jakob Emanuel Handmann. (Kunstmuseum Basel; public domain, via Wikimedia Commons.)

1.4 A way in

Put the three stories beside one another. Sullivan spelled into Keller's hand. Tatyana Pontryagin read mathematics she did not fully understand. Euler's assistants wrote down work that he could no longer see.

In each story, knowledge had to change form before it could travel. For Keller, language entered through touch. Pontryagin received the printed page through his mother's voice. Euler's thoughts reached the page through dictation.

That is easy to miss when we talk about accessibility today, because the helper is often no longer a person in the room. Today, the intermediary may be a screen reader such as VoiceOver, a braille display, or software that reads a document aloud. But the old problem remains: the reader still needs the document to arrive in a form they can use.

Software cannot guess what the author meant. It reads what the file contains. If a figure has no description, a screen reader may say only "image" or "Figure 1.2". That tells the reader that something is on the page, but not what the figure shows, why it was included, or how it supports the argument.

This is why structure matters. Alt text tells the reader what the figure shows and why it matters. A table header tells the software which entries belong together. Readable mathematics gives a formula a chance to be spoken, searched, or sent to braille.

The point is not to make the PDF look more polished. The point is to make the document usable when sight is not the only way in.

1.5 An afternoon

The template supplies much of the structure. The author still has to describe the figures, use headings correctly, and give tables the information that software needs. Over a dissertation, this may add up to an afternoon, perhaps two.

A dissertation is one of the most permanent documents most of us will write. It may remain in a university repository for decades. A future reader will not know how much time the author spent preparing the file. The reader will know only whether the document works.

Making it accessible is part of making that document complete. The Auburn Creed gives language to the responsibility behind this work: “I believe in the human touch, which cultivates sympathy with my fellow men and mutual helpfulness and brings happiness for all” [Pet43].

The next chapter turns from why this matters to how it is done.

Chapter 2

How to Do It: The Building Blocks

This chapter focuses on execution. We cover the one-time setup, how to verify your document, and the "Eight Beatitudes" of accessibility: headings, figures, tables, colour, links, lists, mathematics, and code.

2.1 Before you start

The template is already configured. You only need to follow two rules to keep the tagging system working:

1. **Do not delete** the initialization block at the very top of `main.tex`.
2. **Compile with LuaLaTeX.** Run `latexmk -lualatex main.tex`. On Overleaf, set the compiler to LuaLaTeX *and* the TeX Live version to "Rolling TeX Live" (the default release is too old).

2.2 How to check?

Auburn University Graduate School currently supports checking your finished PDF with **PAC** (PDF Accessibility Checker) or **Adobe Acrobat Pro**. PAC is used to evaluate compliance with WCAG 2.1 accessibility requirements, and Acrobat checks against WCAG 2.0. Both tools have known limitations and may produce false errors, especially in LaTeX-generated PDFs. For detailed checking instructions and guidance on interpreting the results, please consult the official Auburn University Graduate School ETD website.

If you need alternative testing software, browse the official W3C WAI Web Accessibility Evaluation Tools List to find a solution that fits your workflow.

2.3 The Eight Beatitudes

There are exactly eight core elements you must format manually. Master these eight beatitudes of accessibility, and your document will remain structurally sound.

2.3.1 Headings

Why: Screen-reader users jump from heading to heading to skim a document. Formatting regular text in bold to merely look like a heading deprives them of this structural navigation.

How: Use the native LaTeX sectioning commands, and do not skip heading levels.

```
\chapter{...}        % Heading 1
\section{...}        % Heading 2
\subsection{...}     % Heading 3
\subsubsection{...}  % Heading 4
```

Example:

- **Right:** `\section{Methods}`
- **Wrong:** `{\large\bfseries Methods}`—this only mimics the visual appearance.

2.3.2 Figures

Why: Screen readers rely on alt text to describe visual content. Without it, the software will remain silent or read the raw file name.

How: Place your image in `figures/`, then configure it as follows:

```
\begin{figure}[htbp]
  \centering
  \includegraphics[
    width=0.6\linewidth,
    alt={Curved petal-like sheets that fold
         and meet near a central hub}]
    {figures/calabi-yau.png}
  \caption{A cross-section of a Calabi-Yau manifold.}
  \label{fig:calabiyau}
\end{figure}
```

A meaningful image requires `alt={...}`: provide one descriptive sentence, and avoid redundant phrases like "image of". For purely decorative images, use the `artifact` tag (`\includegraphics[artifact]{...}`) so the screen reader skips it entirely. The caption explains *why* the figure is relevant; the alt text describes *what* it visually depicts. Cross-reference using `\ref`, not directional language like "the figure below".

Example:

- **Good alt:** "Aubie the Tiger high-fiving children outside the stadium."
- **Poor alt:** "Image 1.2"

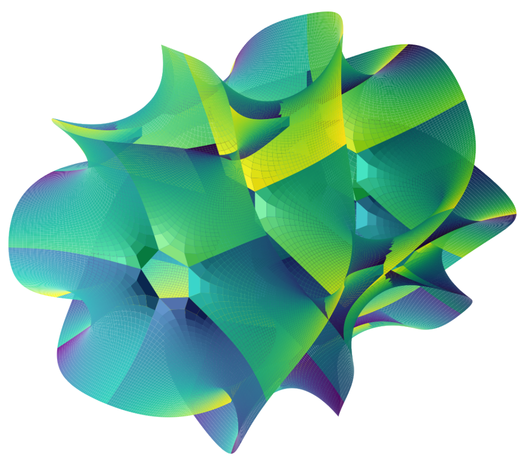

Figure 2.1: A two-dimensional cross-section of a Calabi-Yau manifold—the Fermat quintic $z_1^5 + z_2^5 = 1$—projected from four dimensions to three; colour shows the projected-out fourth coordinate.

2.3.3 Tables

Why: Without designated header rows, a table is read as a continuous stream of data. With headers marked, the screen reader announces context (e.g., "Krull dimension: 2") as it enters a new cell.

How: Add a setup line before your `tabular` environment to declare the header rows. Wrap individual header cells in `\auhead`.

```
\begin{table}[htbp]
  \centering
  \caption{Some rings and their properties.}
  \label{tab:rings}
  \tagpdfsetup{table/header-rows=1} % {1,2} for two header rows
  \begin{tabular}{lcc}
    \toprule
    \auhead{Ring} & \auhead{Krull dim} & \auhead{UFD?} \\
    \midrule
    $\mathbb{Z}$ & 1 & yes \\
    ...
    \bottomrule
  \end{tabular}
\end{table}
```

Tables are for data, not visual alignment. If you must use a `tabular` environment purely for layout, instruct the screen reader to ignore the grid by adding this line before it:

```
\tagpdfsetup{table/tagging=presentation}
```

Example: In Table 2.1, the header row is properly tagged.

Table 2.1: Some commutative rings, with their Krull dimension and whether they are unique factorization domains (UFDs).

Ring R	**Krull dimension**	**UFD?**
$\mathbb{Z}$	1	yes
$k[x, y]$	2	yes
$k[x, y]/(y^2 - x^3)$	1	no

Complex tables: For a large or multi-part table, also add a *summary*—a short, screen-reader-only sentence describing how the table is organised. WCAG recommends one for complex data tables; it is optional for simple ones and is not required by PDF/UA. Place `\tablesummary` right after `\end{tabular}`, inside the float:

```
\begin{table}[htbp]
  \centering
  \caption{...}
  \tagpdfsetup{table/header-rows=1}
  \begin{tabular}{lcc}
    ... header row and data ...
  \end{tabular}
  \tablesummary{What the table shows and how it is
    arranged, in a sentence or two.}
\end{table}
```

This writes the summary in exactly the form Adobe Acrobat reads (a `/O /Table /Summary` attribute), so it appears under the table's "Edit Table Summary" field and is announced by screen readers, with nothing changed on the printed page.

2.3.4 Color

Why: WCAG mandates strict minimum contrast ratios and prohibits using colour as the sole method of conveying information.

How: Keep primary body text black. For added colour, select dark shades that meet the requirements in Table 2.2 (pure red, green, and yellow fail on white). Ensure any colour cue is paired with a secondary indicator, such as specific wording, shapes, or font weight.

```
\definecolor{AccessibleNavy}{RGB}{16,45,98} % ~12:1
\colorlet{AccessibleRed}{red!80!black}      % ~5.9:1
\colorlet{AccessibleGreen}{green!40!black}  % ~7.2:1
```

Example:

- **Right:** Use accessible colours together with bold labels, line styles, or marker shapes so that colour is not the only visual cue.
- **Wrong:** Use colour alone to distinguish meanings, categories, or data series.

Table 2.2: Minimum contrast ratios required by WCAG 2.1.

Text	Minimum ratio
Normal, Level AA	4.5:1
Large, Level AA	3:1
Normal, Level AAA	7:1
Large, Level AAA	4.5:1

2.3.5 Links

Why: Vague phrases like "click here" provide screen-reader users with no context about a link's destination.

How: Use descriptive language that identifies the destination. The `\hrefu` command automatically underlines the link for visual clarity.

```
\hrefu{<destination URL>}
      {the author's homepage}
```

Example:

- **Right:** Visit Auburn University Graduate School website.
- **Wrong:** "To visit the homepage, click here".

2.3.6 Lists

Why: A properly formatted list announces its structure (e.g., "list, 3 items"). Manually typing dashes provides no structural metadata.

How: Use standard LaTeX environments (`itemize`, `enumerate`, or `description`). They are automatically tagged.

```
\begin{itemize}
  \item First point
  \item Second point
\end{itemize}
```

Example:

- First point.
- Second point.
- Third point.

2.3.7 Mathematics

Why: Screen readers often ignore standard equations in PDFs. This template resolves the issue natively, ensuring equations are read aloud as mathematics.

How: Just write math normally using `equation`, `align`, or inline `$...$`. No special per-equation tags are needed.

Example:

$$e^{i\pi} + 1 = 0, \qquad \int_0^1 x^2\,dx = \frac{1}{3}. \tag{2.1}$$

Theorem 2.1 (Splitting). Every short exact sequence of vector spaces splits.

2.3.8 Code

Why: The `Verbatim` environment is fully tagged. Packages like `listings` and `minted` generate untagged content and will fail accessibility checks.

How: Standardize on `Verbatim` for all code blocks.

```
def greet(name):
    return f"Hello, {name}!"
```

2.4 The rest is automatic

The template handles typography, spacing, margins, and pagination automatically. Just keep two final best practices in mind:

- Mark foreign language phrases using `\foreignlanguage{french}{...}`. This ensures screen readers apply correct pronunciation rules.

- Convey emphasis through word choice and font weight, never by relying exclusively on color or ALL CAPS.

That concludes the toolkit. Chapter 3 puts these building blocks to work.

Chapter 3

A Sample Chapter: Sheaves and Sheafification

This chapter is a demonstration. It is a short, self-contained introduction to sheaves, in the spirit of Vakil [Vak17], and at the same time it applies every building block of Chapter 2 in context: definitions, propositions, and theorems with proofs; displayed equations; three figures with alt text, one of them a commutative diagram; a table with a header row; a code listing; links; footnotes; a block quotation; all three kinds of list; and headings down to level four. Replace it with your own writing; keep the patterns.

3.1 Presheaves

A sheaf is a bookkeeping device for local information: it records the data that live over each open set of a space, and how those data restrict to smaller open sets [Vak17; Har77].

Definition 3.1 (Presheaf). Let X be a topological space. A *presheaf* $\mathcal{F}$ of abelian groups on X assigns to each open set $U \subseteq X$ an abelian group $\mathcal{F}(U)$, and to each inclusion of open sets $V \subseteq U$ a *restriction* homomorphism $\mathrm{res}_{U,V} \colon \mathcal{F}(U) \to \mathcal{F}(V)$, such that $\mathrm{res}_{U,U} = \mathrm{id}$ and, for $W \subseteq V \subseteq U$,

$$\mathrm{res}_{V,W} \circ \mathrm{res}_{U,V} = \mathrm{res}_{U,W}. \tag{3.1}$$

Elements of $\mathcal{F}(U)$ are called *sections* of $\mathcal{F}$ over U, and one writes $s|_V$ for $\mathrm{res}_{U,V}(s)$.

Example 3.2. On any space X, let $\mathcal{C}(U)$ be the continuous functions $U \to \mathbb{R}$, with the usual restriction of functions. Restricting in two steps is the same as restricting in one, which is exactly (3.1), so $\mathcal{C}$ is a presheaf.

3.2 Sheaves

A sheaf is a presheaf whose sections are determined by local data.

Definition 3.3 (Sheaf). A presheaf $\mathcal{F}$ on X is a *sheaf* if, for every open $U \subseteq X$ and every open cover $U = \bigcup_i U_i$:

1. **Identity.** If $s, t \in \mathcal{F}(U)$ satisfy $s|_{U_i} = t|_{U_i}$ for all i, then $s = t$.
2. **Gluing.** If sections $s_i \in \mathcal{F}(U_i)$ agree on overlaps—$s_i|_{U_i \cap U_j} = s_j|_{U_i \cap U_j}$ for all i, j—then some $s \in \mathcal{F}(U)$ has $s|_{U_i} = s_i$ for all i.

Together the axioms say: compatible local data patch to a unique global datum.

Example 3.4. The presheaf $\mathcal{C}$ of continuous functions is a sheaf. Continuity is a local condition, so functions that agree on overlaps glue to a unique continuous function.

Example 3.5 (A presheaf that is not a sheaf). On $X = \mathbb{R}$, let $\mathcal{B}(U)$ be the *bounded* continuous functions on U. Cover $\mathbb{R}$ by the intervals $U_n = (-n, n)$ and take $s_n = x|_{U_n}$, bounded on each U_n. The s_n agree on overlaps, but the only function they could glue to is x itself, which is not bounded on $\mathbb{R}$. Gluing fails, because boundedness is not a local condition.

Example 3.6 (The constant presheaf). Fix an abelian group A with at least two elements. The *constant presheaf* sets $\mathcal{F}(U) = A$ for every nonempty open U and $\mathcal{F}(\emptyset) = 0$, with identity restrictions between nonempty sets. Suppose $U = U_1 \sqcup U_2$ is a disjoint union of nonempty open sets, and pick $a \neq b$ in A. The sections $a \in \mathcal{F}(U_1)$ and $b \in \mathcal{F}(U_2)$ agree on the overlap $U_1 \cap U_2 = \emptyset$, where both restrict to 0—but no single element of A restricts to a on U_1 and to b on U_2. Gluing fails.

Table 3.1 collects these examples and two more. The pattern: a presheaf is a sheaf exactly when its defining condition is local. Two standard places to read further are the Stacks Project, Tag 006A [Stacks] and the sheaf chapter of *The Rising Sea* [Vak17].

3.3 Stalks and germs

What does a sheaf know about a single point? Everything that happens arbitrarily close to it.

Table 3.1: Some presheaves on $X = \mathbb{R}$, and whether they are sheaves.

Presheaf $\mathcal{F}(U)$	**Sheaf?**	**Reason**
continuous functions	yes	continuity is local
smooth functions	yes	smoothness is local
locally constant functions	yes	the condition is local
bounded functions	no	gluing fails (Ex. 3.5)
constant presheaf A	no	gluing fails (Ex. 3.6)

Definition 3.7 (Stalk). The *stalk* of a presheaf $\mathcal{F}$ at a point $p \in X$ is the direct limit over the open neighbourhoods of p,

$$\mathcal{F}_p \;=\; \operatorname*{colim}_{p \in U} \mathcal{F}(U), \tag{3.2}$$

taken along the restriction maps (Figure 3.1).

Concretely, an element of $\mathcal{F}_p$ is represented by a pair (U, s) with $p \in U$ and $s \in \mathcal{F}(U)$; two pairs (U, s) and (V, t) represent the same element when

$$s|_W = t|_W \qquad \text{for some open } W \subseteq U \cap V \text{ with } p \in W. \tag{3.3}$$

This common element is the *germ* of s at p. Two sections have the same germ exactly when they agree on some neighbourhood of p, however small (Figure 3.2).

Example 3.8. For real-analytic functions, the germ at p carries the same information as the Taylor series at p. For smooth functions the germ remembers strictly more.[1]

For a sheaf—not a mere presheaf—the germs collectively remember everything:

Proposition 3.9 (Sections are determined by their germs). Let $\mathcal{F}$ be a sheaf and U open. The map $\mathcal{F}(U) \to \prod_{p \in U} \mathcal{F}_p$ sending a section to its family of germs is injective.

1

The classic example is $f(x) = e^{-1/x^2}$, with $f(0) = 0$: it is smooth and nonzero arbitrarily close to 0, yet every term of its Taylor series at 0 vanishes—the germ sees what the series cannot.

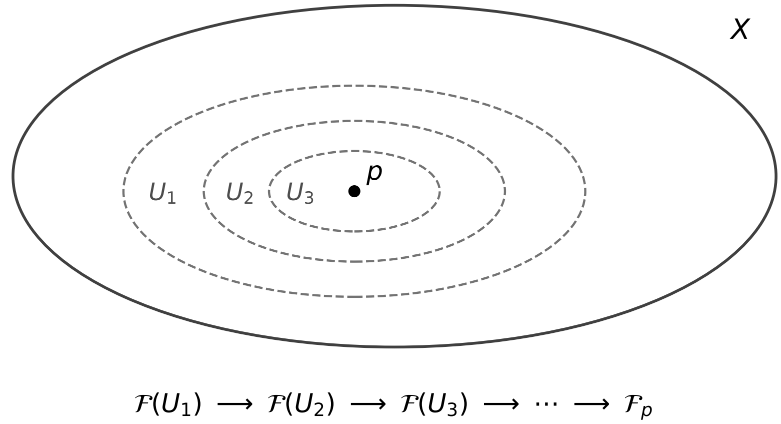

Figure 3.1: The stalk $\mathcal{F}_p$ is the limit of the groups $\mathcal{F}(U)$ over smaller and smaller neighbourhoods U of p.

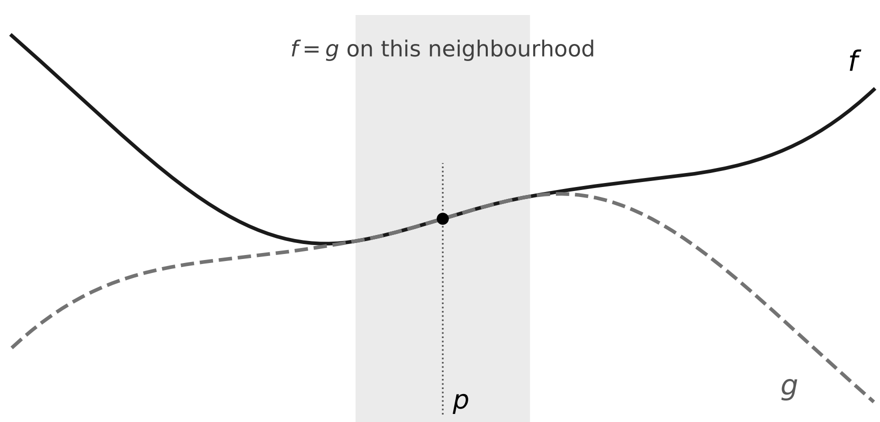

Figure 3.2: Two functions with the same germ at p: they agree on a neighbourhood of p, and may do as they please elsewhere.

Proof. If $s, t \in \mathcal{F}(U)$ have the same germ at every $p \in U$, then each point has a neighbourhood on which s and t agree. These neighbourhoods cover U, so $s = t$ by the identity axiom. □

This is the precise sense in which a sheaf is local: a section is nothing more than a compatible family of germs—which is exactly the description that sheafification, in Section 3.4, will turn into a definition.

A small vocabulary, for reference:

Section an element of $\mathcal{F}(U)$: one piece of data over the open set U.

Restriction the map $\mathcal{F}(U) \to \mathcal{F}(V)$ that forgets everything outside V.

Germ the equivalence class of a section near a point.

Stalk the group of all germs at a point.

3.4 Sheafification

Natural constructions often produce presheaves that are not sheaves—the constant presheaf is the standard example. Sheafification repairs this, changing as little as possible.

3.4.1 The universal property

Theorem 3.10 (Sheafification). Let $\mathcal{F}$ be a presheaf on X. There exist a sheaf $\mathcal{F}^{\mathrm{sh}}$ and a morphism $\mathrm{sh} : \mathcal{F} \to \mathcal{F}^{\mathrm{sh}}$ such that every morphism from $\mathcal{F}$ to a sheaf factors uniquely through sh. Moreover, sh is an isomorphism on every stalk:

$$\mathcal{F}_p \xrightarrow{\sim} (\mathcal{F}^{\mathrm{sh}})_p \qquad \text{for all } p \in X. \tag{3.4}$$

Proof sketch. Build $\mathcal{F}^{\mathrm{sh}}$ out of compatible families of germs: for U open, let

$$\mathcal{F}^{\mathrm{sh}}(U) = \Big\{ (s_p)_{p \in U} \in \textstyle\prod_{p \in U} \mathcal{F}_p \ : \ \text{the family is locally a section} \Big\}, \tag{3.5}$$

where “locally a section” means every point of U has a neighbourhood V and a section $t \in \mathcal{F}(V)$ whose germ at each $q \in V$ is s_q. The identity axiom holds because a section here is a family indexed by the points of U; the gluing axiom holds because membership in (3.5) is a local condition. The map sh sends a section to its family of germs, and the universal property and (3.4) follow by unwinding the definitions. Details are in Hartshorne [Har77, Section II.1] or Vakil [Vak17]. □

Three facts worth keeping in mind:

- Sheafification changes nothing at any stalk—only how local data assemble into sections.
- If $\mathcal{F}$ is already a sheaf, sh is an isomorphism.
- The construction is idempotent: $(\mathcal{F}^{\mathrm{sh}})^{\mathrm{sh}} \cong \mathcal{F}^{\mathrm{sh}}$, canonically.

3.4.2 Worked examples

The constant presheaf, repaired

Let A be an abelian group and $\mathcal{F}$ the constant presheaf of Example 3.6. Every stalk is A, so by (3.5) the sheafification is the sheaf of *locally constant* functions $U \to A$, written $\underline{A}$. On an open subset of $\mathbb{R}$ with n connected components,

$$\underline{A}(U) \cong A^n, \tag{3.6}$$

one constant per component.

Corollary 3.11. Let $X = \{1, 2\}$ with the discrete topology. The constant presheaf $\mathbb{Z}$ has global sections $\mathbb{Z}$, while its sheafification has global sections $\underline{\mathbb{Z}}(X) \cong \mathbb{Z} \oplus \mathbb{Z}$. Sheafification can change global sections, even though it changes no stalk.

How to recognise a sheafification

In practice one rarely uses (3.5) directly. The working tool is a standard fact: a morphism of *sheaves* that is an isomorphism on every stalk is an isomorphism. To identify $\mathcal{F}^{\mathrm{sh}}$:

1. compute the stalks of $\mathcal{F}$;
2. find a sheaf $\mathcal{G}$ and a morphism $\mathcal{F} \to \mathcal{G}$ inducing an isomorphism on every stalk;
3. conclude that $\mathcal{G} \cong \mathcal{F}^{\mathrm{sh}}$: Theorem 3.10 factors the morphism through a map $\mathcal{F}^{\mathrm{sh}} \to \mathcal{G}$ of sheaves that is still an isomorphism on every stalk—so, by the standard fact, an isomorphism.

3.5 Morphisms and the snake lemma

Sheaves on X form a category, and a remarkably well-behaved one: the sheaves of abelian groups form an *abelian* category, so kernels, cokernels, exact sequences, and diagram chases all make sense [Vak17, §2.6].

Definition 3.12 (Morphism of sheaves). A *morphism* $\varphi : \mathcal{F} \to \mathcal{G}$ of sheaves on X is a homomorphism $\varphi_U : \mathcal{F}(U) \to \mathcal{G}(U)$ for each open U, commuting with the restriction maps. Taking limits, φ induces a map $\varphi_p : \mathcal{F}_p \to \mathcal{G}_p$ on every stalk.

Proposition 3.13. For any morphism $\varphi : \mathcal{F} \to \mathcal{G}$, the presheaf $U \mapsto \ker \varphi_U$ is a sheaf.

Proof. Identity is inherited from $\mathcal{F}$. For gluing, compatible sections $s_i \in \ker \varphi_{U_i}$ glue to some $s \in \mathcal{F}(U)$; then $\varphi_U(s)|_{U_i} = \varphi_{U_i}(s|_{U_i}) = 0$ for all i, so $\varphi_U(s) = 0$ by the identity axiom in $\mathcal{G}$. □

The cokernel is a different story: the presheaf $U \mapsto \operatorname{coker} \varphi_U$ can fail to be a sheaf, and the cokernel of a morphism of sheaves is *defined* as its sheafification—the first place the construction of Section 3.4 earns its keep. Exactness is then measured where sheaves live:

Definition 3.14 (Exact sequence). A sequence of sheaves $\mathcal{F} \to \mathcal{G} \to \mathcal{H}$ is *exact* if the induced sequence of stalks $\mathcal{F}_p \to \mathcal{G}_p \to \mathcal{H}_p$ is exact for every $p \in X$.

3.5.1 The exponential sequence

Why must the cokernel be sheafified? Because a morphism can be surjective on every stalk yet miss sections. The classic witness lives on the punctured plane $X = \mathbb{C} \setminus \{0\}$. Write $\mathcal{O}_X$ for the sheaf of holomorphic functions and $\mathcal{O}_X^\times$ for the nowhere-vanishing ones, and let exp send f to e^f:

$$0 \longrightarrow \underline{2\pi i \mathbb{Z}} \longrightarrow \mathcal{O}_X \xrightarrow{\text{exp}} \mathcal{O}_X^\times \longrightarrow 0. \tag{3.7}$$

The kernel is the locally constant sheaf of Section 3.4.2: $e^f = 1$ exactly when f takes values in $2\pi i \mathbb{Z}$. The sequence is exact: on a small disc every nowhere-vanishing function has a logarithm, so exp is surjective on every stalk. But over X itself the function $g(z) = z$ has no logarithm, so exp_X is not surjective on global sections—and the image presheaf $U \mapsto \text{exp}\, \mathcal{O}_X(U)$ fails gluing in exactly the manner of Example 3.5: cover X by two slit planes, take the restrictions of z, and the local logarithms refuse to patch.

3.5.2 The snake lemma

The reward for setting up an abelian category is that the standard tools of homological algebra apply verbatim. The first and most famous is the snake lemma [Vak17, §1.7].

Theorem 3.15 (Snake lemma). In an abelian category—sheaves of abelian groups included—suppose the rows of the commutative diagram in Figure 3.3 are exact. Then there is a connecting morphism δ making the sequence

$$0 \to \ker f \to \ker g \to \ker h \xrightarrow{\delta} \operatorname{coker} f \to \operatorname{coker} g \to \operatorname{coker} h \to 0 \tag{3.8}$$

exact.

Proof idea. A diagram chase, best performed personally rather than read.[2] Chase an element of $\ker h$ through Figure 3.3: lift it to B, push it down to B', observe that it comes from a

[2] The chase even has a film credit: the snake lemma is proved on screen in the opening scene of *It's My Turn* (1980).

unique element of A', and define δ to be its class in coker f. $\square$

$$
\begin{tikzcd}
 & \ker f \arrow[r] \arrow[d] & \ker g \arrow[r] \arrow[d] & \ker h \arrow[d] \arrow[dddll, rounded corners, "\delta"'] & \\
0 \arrow[r] & A \arrow[r] \arrow[d, "f"'] & B \arrow[r] \arrow[d, "g"'] & C \arrow[r] \arrow[d, "h"'] & 0 \\
0 \arrow[r] & A' \arrow[r] \arrow[d] & B' \arrow[r] \arrow[d] & C' \arrow[r] \arrow[d] & 0 \\
 & \operatorname{coker} f \arrow[r] & \operatorname{coker} g \arrow[r] & \operatorname{coker} h &
\end{tikzcd}
$$

Figure 3.3: The snake lemma. The connecting morphism δ gives the lemma its name.

Remark 3.16. The snake lemma is the seed of sheaf cohomology: it is the engine that produces long exact sequences from short ones. Applied to (3.7), it explains where the missing logarithm went—the failure of exp on global sections is measured by the first cohomology of $\underline{2\pi i\mathbb{Z}}$, which on the punctured plane records precisely the winding number.

By machine: a first taste of cohomology

The machinery is already on your computer. In Macaulay2 [GS], the cohomology of line bundles on the projective line takes three lines:

```
R = QQ[x,y];          -- homogeneous coordinates
X = Proj R;           -- the projective line
HH^0(OO_X(2))         -- global sections: QQ^3
```

Macaulay2 answers $\mathbb{Q}^3$—the quadratics x^2, xy, y^2—and `HH^1(OO_X(-2))` likewise returns $\mathbb{Q}^1$. Where those numbers come from is the business of the next course; that they come out of the definitions in this chapter is the point.

As an aside on how to read this chapter—and any mathematics:

> The substance of mathematics lies in worked examples; a definition is only understood once it has been turned over in the hand a few times.

References

[Cal16] R. S. Calinger. *Leonhard Euler: Mathematical Genius in the Enlightenment*. Princeton, NJ: Princeton University Press, 2016.

[GS] D. R. Grayson and M. E. Stillman. *Macaulay2, a software system for research in algebraic geometry*. URL: `https://macaulay2.com/` (visited on 06/12/2026).

[Har77] R. Hartshorne. *Algebraic Geometry*. Graduate Texts in Mathematics 52. New York: Springer-Verlag, 1977.

[Kel03] H. Keller. *The Story of My Life*. New York: Doubleday, Page & Company, 1903.

[OR14] J. J. O'Connor and E. F. Robertson. *Lev Semenovich Pontryagin*. MacTutor History of Mathematics Archive, University of St Andrews. 2014. URL: `https://mathshistory.st-andrews.ac.uk/Biographies/Pontryagin/` (visited on 06/12/2026).

[Pet43] G. Petrie. *The Auburn Creed*. Auburn University. 1943. URL: `https://www.auburn.edu/main/welcome/creed.php` (visited on 06/12/2026).

[Stacks] The Stacks Project Authors. *The Stacks Project*. URL: `https://stacks.math.columbia.edu` (visited on 06/12/2026).

[Vak17] R. Vakil. *The Rising Sea: Foundations of Algebraic Geometry*. Lecture notes. 2017.

Appendix 1

A Minimal Accessible Preamble

You do not need this whole class to make an accessible document. The skeleton below is the smallest starting point: the metadata line (PDF/UA-2, so formulas enter the tag tree as readable MathML), the Lua font setup with a sans-serif body face, and the two content rules (alt text, table headers). Compile it with LuaLaTeX.

```
\DocumentMetadata{lang=en-US,
  pdfstandard=ua-2, tagging=on,
  tagging-setup={math/setup={mathml-AF,mathml-SE}}}
\documentclass{article}
\usepackage{fontspec}      % LuaLaTeX fonts
\setmainfont{Arial}        % sans-serif body font
\usepackage{unicode-math}  % OpenType math
\usepackage{graphicx}
\usepackage{booktabs}
\begin{document}

\section{A heading}
Some text, and math $e^{i\pi}+1=0$.

\begin{figure}\centering
  \includegraphics[width=.5\linewidth,
    alt={Describe the image}]{pic}
  \caption{Why this figure is here.}
\end{figure}

\tagpdfsetup{table/header-rows={1}}
\begin{tabular}{lr}
  \toprule
  Name & Value \\
  \midrule
  A & 1 \\
  \bottomrule
\end{tabular}

\end{document}
```

A1.1 Where to read more

The accessibility machinery used here comes from the LaTeX Tagging Project; for the authoritative, continually updated instructions, see the Tagging Project documentation. For another university template built on the same machinery – PDF/UA-2, tagged tables, and screen-readable mathematics – see the University of Illinois Graduate College's `uofithesis` class. And for a broader collection of LaTeX tutorials, reference sheets, and accessibility-minded thesis, homework, and poster templates, see Paul Hurtado's LaTeX resource page.