Exercises from Vakil

Ian Coley

February 6, 2015

2 Sheaves

2.1 Motivating example: The sheaf of differentiable functions

Exercise 2.1.A.
Show that the only maximal ideal of the germ of differentiable functions $\mathcal{O}_p$ is $\mathfrak{m}_p$, the ideal of functions vanishing at p.

Solution.
We take for granted that $\mathfrak{m}_p$ is maximal. We will show that $\mathcal{O}_p^\times = \mathcal{O}_p \setminus \mathfrak{m}_p$. Suppose that $f \in \mathcal{O}_p$ does not vanish at p, and say $f(p) = \alpha$. Then since f is differentiable, it is continuous, so it is nonzero in some neighbourhood U of p. Let $K \subset U$ be a compact neighbourhood of p, which exists since we are working in locally Euclidean space. Then we may define $g : U \to \mathbb{R}$ by

$$g(x) = \begin{cases} 1/f(x) & x \in K \\ 0 & x \in U \setminus K \end{cases}.$$

We see that $g(x)$ is certainly continuous in K, but it may not be differentiable if $f'(x) = 0$ somewhere in K. However in this case, we must shrink U and try again until we get it right. (CHECK) Having established that, we may define g on an open subset of $p \in K$ to obtain the appropriate germ. □

Exercise 2.1.B.
Prove that $\mathfrak{m}_p/\mathfrak{m}_p^2$ is naturally isomorphic to the cotangent space at p as an $\mathbb{R}$-module.

Solution.
We would like to establish the exact sequence

$$0 \to \mathfrak{m}_p^2 \to \mathfrak{m}_p \to T_p^* X \to 0.$$

The first map is clearly given by inclusion, but the second map is not as obvious. We will take a function f and send it to its derivation df, which is a member of the cotangent space. This map d is surjective by facts from differential geometry. It is clear that $\mathfrak{m}_p^2 \subset \ker d$ by Leibniz' rule: $d(f^2) = 2f\,df$, so $d(f^2)_p = 0$ since $f(p) = 0$. (TO BE FINISHED) □

2.2 Definition of sheaf and presheaf

Exercise 2.2.A.
Prove: a presheaf is the same thing as a contravariant functor.

Solution.
This is reasonably clear. Suppose we have a contravariant functor $\mathscr{F} : \mathcal{X} \to \mathcal{C}$, where $\mathcal{C}$ is some (concrete) category and $\mathcal{X}$ the category associated to a topological space X. Then we know that given an inclusion $i : V \subset U$, we have a morphism $\mathscr{F}(i) : \mathscr{F}(U) \to \mathscr{F}(V)$. Further, if we have the identity inclusion $\mathrm{id}_U : U \subset U$, then $\mathscr{F}(\mathrm{id}_U) : \mathscr{F}(U) \to \mathscr{F}(U)$ must also be the identity. Further, we know that $\mathscr{F}$ respects (and reverses) compositions. This is all the data that we needed. □

Exercise 2.2.B.
Show that the following are presheaves on $\mathbb{C}$ but not sheaves: (a) bounded functions and (b) holomorphic functions admitting a holomorphic square root.

Solution.
It is clear that both are presheaves by the same reasoning that differentiable functions were a presheaf; functions can be restricted. For (a), we see this is not a presheaf since it violates the gluability axiom. Suppose we have any open cover of $\mathbb{C}$ by bounded open sets. Define $f_i : U_i \to \mathbb{C}$ be given by $f_i(x) = |x|$. Then on any U_i, we have a bounded function. However, we cannot glue all these functions together, since eventually $|x| \to \infty$.

For (b), since he's said this exercise is unimportant and I'm not sure what a holomorphic square root is, I will skip it. □

Exercise 2.2.C.
The identity and gluability axioms may be rephrased to say that $\mathscr{F}(\bigcup_i U_i)$ is what limit?

Solution.
We need $\mathscr{F}(\bigcup_i U_i)$ to be the limit of the diagram J with objects $\mathscr{F}(U_i)$ and maps the restriction maps res_{U_i,U_j}. Let $U = \bigcup_i U$. If this is the case, then the universal property of limits tells us that, given two $f, g \in \mathscr{F}(U)$ which restricts onto each of the U_i in the appropriate way, we have a unique morphism $\mathscr{F}(U) \to \mathscr{F}(U)$ sending f to g.

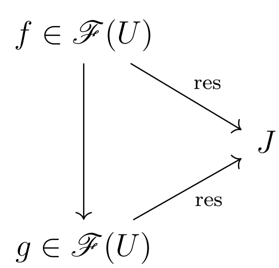

Since the identity map works for this unique morphism, this must be our choice, so $f = g$. This proves identity. Gluability follows from the existence of the limit at all, i.e. there is some element in $\mathscr{F}(U)$ restricting to each of the $\mathscr{F}(U_i)$ in the appropriate ways. □

Exercise 2.2.D.

(a) Verify that differentiable functions, continuous functions, smooth functions, and ordinary functions are shaves on a manifold $X \subset \mathbb{R}^n$.

(b) Show that real-valued continuous functions on (open sets of) a topological space X form a sheaf.

Solution.
We assume that everything is a presheaf already. The method for (b) will also work for (a), so we will just do that.

In a topological space, we do not have coordinate charts, but everything works fine still. The identity axiom is clear. The gluability axiom also works since the function obtained by gluing is still continuous. If we were dealing with smooth functions on a manifold, then the result would still be smooth, as the functions are glued together on open sets and not on closed sets (where we might have problems). □

Exercise 2.2.E.
Let $\mathscr{F}(U)$ be maps from U to a set S that are locally constant. Show that this is a sheaf.

Solution.
We take the description that $\mathscr{F}(U)$ is the set of continuous maps $U \to S$, where S is endowed with the discrete topology. This is clearly a presheaf, so we need only check that it is a sheaf. But this is clear, as the identity axiom is trivial and the gluability axiom follows by the following: glue together $f_i : U_i \to S$ into one function $f : U \to S$. Then for each $p \in U$, choose the constant open neighbourhood around p to be the one guaranteed to exist in $U_i \subset U$. □

Exercise 2.2.F.
Suppose Y is a topological space. Show that "continuous maps to Y" form a sheaf of sets on X.

Solution.
That it is a presheaf and the identity axiom are obvious. To test gluability, suppose that we have an appropriate cover $\{U_i\}_{i \in I}$ and functions $f_i : U_i \to Y$. Define $f : U \to Y$ in the obvious way. Then let $V \subset Y$ be an open set. We know that $f_i^{-1}(V)$ is open in U_i for all $i \in I$. Thus

$$f^{-1}(V) = \bigcup_{i \in I} f_i^{-1}(V)$$

is still open, since an open subset of an open subset is still open in U. Hence the function is continuous. We can apply this proof to the above cases. □

Exercise 2.2.G.

(a) Suppose we are given a continuous map $\mu : Y \to X$. how that "sections of μ" form a sheaf.

(b) Suppose that Y is a topological group. Show that continuous maps to Y form a sheaf of *groups*.

Solution.

(a) The identity axiom is clear as it ever was, since sections are a type of continuous functions to a topological space. Suppose now that we have our set $\{s_i : U_i \to Y\}$. Glue together the functions in the necessary way to create $s : U \to Y$. Then we need s to be injective to be left invertible, but this is easily the case. Suppose $s(p) = s(q)$ with $p \neq q$. Then p and q cannot lie in the same U_i, since the s_i are all injective. Therefore let $p \in U_i$ and $q \in U_j$. We could not have $(\mu \circ s)(p) = (\mu \circ s)(q)$ since p and q lie in different open sets, which implies that $s(p) \neq s(q)$, a contradiction.

(b) We can put a group structure on the set of continuous maps by pointwise multiplication. The presheaf axioms are still clearly satisfied, and the identity element of any $\mathscr{F}(U)$ is the constant map sending every element to e_Y, so we have a contravariant functor to the category of groups. There is no reason any of the other axioms should fail because we have added the rest of this structure, and indeed they do not.

□

Exercise 2.2.H.
Suppose $\pi : X \to Y$ is a continuous map, and $\mathscr{F}$ is a presheaf on X. Then define $\pi_*\mathscr{F}$ by $\pi_*\mathscr{F}(V) = \mathscr{F}(\pi^{-1}(V))$, where $V \subset Y$ is open. Show $\pi_*\mathscr{F}$ is a presheaf on Y, and is a sheaf of $\mathscr{F}$ is.

Solution.
If we have such a presheaf, then we can define the restriction maps in the following way: let $U \subset V \subset Y$ be open subsets. Then since $\pi^{-1}(U) = U'$ and $\pi^{-1}(V) = V'$ are open subsets in X satisfying $U' \subset V'$, there is a restriction map $r' : \mathscr{F}(V') \to \mathscr{F}(U')$. As such, we can construct the restriction map $r : \pi_*\mathscr{F}(V) \to \pi_*\mathscr{F}(U)$ with this in the only sensible way. Thus we have a presheaf structure on $\pi_*\mathscr{F}$. If $\mathscr{F}$ is a sheaf, then taking 2.2.C, everything still works out fine. □

Exercise 2.2.I.
Suppose $\pi : X \to Y$ is a continuous map, and $\mathscr{F}$ is a sheaf of sets on X. If $\pi(p) = q$, describe the natural morphism of stalks $(\pi_*\mathscr{F})_q \to \mathscr{F}_p$.

Solution.
We will take the definition of the stalk as a colimit. We have that

$$(\pi_*\mathscr{F})_q = \varinjlim \pi_*\mathscr{F}(U) = \varinjlim \mathscr{F}(\pi^{-1}(U))$$

Since each U is an open set containing p, we have the natural map we wanted by the map induced by $\pi^{-1}(U) \to U$, i.e. π itself. □

Exercise 2.2.J.
If $(X, \mathcal{O}_X)$ is a ringed space, and $\mathscr{F}$ is an $\mathcal{O}_X$-module, describe for each $p \in X$ how $\mathscr{F}_p$ is an $\mathcal{O}_{X,p}$-module.

Solution.
We have that

$$\mathscr{O}_{X,p} = \varinjlim \mathscr{O}_X(U) \qquad \text{and} \qquad \mathscr{F}_p = \varinjlim \mathscr{F}(U).$$

Each $\mathscr{F}(U)$ has a well-defined structure as a $\mathscr{O}_X(U)$-module. Further, if $V \subset U$, we can turn $\mathscr{F}(U)$ into a $\mathscr{O}_X(V)$-module by moving forward in the limit diagram, since we know everything must commute. □

2.3 Morphisms of presheaves and sheaves

Exercise 2.3.A.
If $\phi : \mathscr{F} \to \mathscr{G}$ is a morphism of sheaves on X and $p \in X$, described an induced morphism of stalks $\phi_p : \mathscr{F}_P \to \mathscr{G}_p$.

Solution.
Since we have, for $p \in U$ open,

$$\mathscr{F}_p = \varinjlim \mathscr{F}(U),$$

we can define ϕ_p by elements by appealing to $\phi(U) : \mathscr{F}(U) \to \mathscr{G}(U)$. Taking the germ definition, can take a representative for an equivalence class (f, U) and map it according to $\phi(U)$, same as above. □

Exercise 2.3.B.
Suppose $\pi : X \to Y$ is a continuous map of topological spaces. Show that the pushforward gives a functor $\pi_* : \mathbf{Sets}_X \to \mathbf{Sets}_Y$, where **Sets** is an arbitrary choice.

Solution.
We have seen that the pushforward gives us a map of objects, but we need to check that it respects morphisms. Let $\mathscr{F}, \mathscr{G} \in \mathbf{Sets}_X$, and let $\phi : \mathscr{F} \to \mathscr{G}$ be a sheaf morphism. Then we see that $\pi_*\phi$ must fit in the diagram below (for every open set (U)):

$$\begin{array}{ccc} \mathscr{F} & \xrightarrow{\phi} & \mathscr{G} \\ {\scriptstyle \pi_*}\big\downarrow & & {\scriptstyle \pi_*}\big\downarrow \\ \pi_*\mathscr{F} & \xrightarrow{\pi_*\phi} & \pi_*\mathscr{G} \end{array}$$

Dissecting this, suppose we have an element $\alpha \in \pi_*\mathscr{F}(U) = \mathscr{F}(\pi^{-1}(U))$. Then we must have $\pi_*\phi(U)(\alpha) = \phi(\pi^{-1}(U))(\alpha)$, which is an element of $\mathscr{G}(\pi^{-1}(U))$. Hence under π_* it is send to an element of $\pi_*\mathscr{G}(U)$, as required. It is clear from this description that if ϕ is the identify on $\mathscr{F}$ that $\pi_*\phi$ is the identity on $\pi_*\mathscr{F}$ so we are done. □

Exercise 2.3.C.
Suppose $\mathscr{F}$ and $\mathscr{G}$ are two sheaves of sets on X, though it suffices they be presheaves. Let $\mathscr{H}om(\mathscr{F}, \mathscr{G})$ be the collection of data

$$\mathscr{H}om(\mathscr{F}, \mathscr{G})(U) := \mathrm{Mor}(\mathscr{F}|_U, \mathscr{G}|_U),$$

where $\mathscr{F}|_U$ is the restriction of $\mathscr{F}$ to the open set U. Show that this is a sheaf of sets on X.

Solution.
Let us look at $\mathscr{H}om(\mathscr{F},\mathscr{G})$ as a functor $X \to \mathbf{Sets}$. Given an inclusion $V \subset U$, we can define the following map on the morphism sets: if $f : \mathscr{F}|_U \to \mathscr{G}|_U$, then we have a natural restriction of $\mathscr{F}|_U$ to $\mathscr{F}|_V$, which will in turn map onto $\mathscr{G}|_V$. This is how we map $\mathscr{H}om(\mathscr{F},\mathscr{G})(U) \to \mathscr{H}om(\mathscr{F},\mathscr{G})(V)$. This gives us a presheaf structure on $\mathscr{H}om(\mathscr{F},\mathscr{G})$. The sheaf axioms follow from the limit definition of sheaf axioms on $\mathscr{F}$ and $\mathscr{G}$. The $\mathscr{H}om$ functor that we could develop now, but Vakil hints will be developed soon after, should commute properly with limits. □

Exercise 2.3.D.

(a) If $\mathscr{F}$ is a sheaf of sets on X, then show that $\mathscr{H}om(\underline{\{p\}},\mathscr{F}) \cong \mathscr{F}$, where $\underline{\{p\}}$ is the constant sheaf associated with the one element set $\{p\}$.

(b) If $\mathscr{F}$ is a sheaf of abelian groups on X, then show that $\mathscr{H}om_{\mathbf{Ab}_X}(\underline{\mathbb{Z}},\mathscr{F}) \cong \mathscr{F}$ as sheaves of abelian groups.

(c) If $\mathscr{F}$ is an $\mathcal{O}_X$-module, then show that $\mathscr{H}om_{\mathcal{O}_X}(\mathcal{O}_X,\mathscr{F}) \cong \mathscr{F}$ as sheaves of $\mathcal{O}_X$-modules.

Solution.
For (a), we need to check what is happening on any open set $U \subset X$. Consider $\mathscr{H}om(\underline{\{p\}},\mathscr{F})(U)$. We have

$$\mathscr{H}om(\underline{\{p\}},\mathscr{F})(U) = \mathrm{Mor}(\underline{\{p\}}|_U,\mathscr{F}|_U) = \{\phi : \{p\} \to \mathscr{F}(U)\} \cong \mathscr{F}(U).$$

The Mor set is just the inclusion of one point into $\mathscr{F}(U)$, so these are isomorphic as sets. As described above, given $V \subset U$, the restriction maps would then just map $\phi : \{p\} \to \mathscr{F}(V)$, and the identity and gluability should follow easily too.

We are now going to generalise the above, since all we are considering is the sheaf based around the object in a category representing the identity functor. We can address (b) and (c) simultaneous by generally considering sheaves of R-modules. In general, $\mathrm{Hom}_{\mathbf{R\text{-}Mod}}(R,M) \cong M$ as R-modules. Therefore we have for every open U,

$$\mathscr{H}om_{\mathbf{R\text{-}Mod}_X}(\underline{R},\mathscr{F})(U) = \{\phi : R \to \mathscr{F}(U)\} \cong \mathscr{F}(U).$$

Applying this to the case of $R = \mathbb{Z}, \mathcal{O}_X$ is what we want. □

Exercise 2.3.E.
Check that, for a morphism of presheaves of abelian groups $\phi : \mathscr{F} \to \mathscr{G}$, $\ker_{\mathrm{pre}} \phi$ is a presheaf.

Solution.
To write down what this object actually is, we have $(\ker_{\mathrm{pre}})\phi(U) := \ker(\phi(U) : \mathscr{F}(U) \to \mathscr{G}(U))$. We will take Vakil's hint and use the diagram, for $V \subset U$,

$$\begin{array}{ccccccc}
0 & \longrightarrow & \ker_{\mathrm{pre}}\phi(U) & \longrightarrow & \mathscr{F}(U) & \xrightarrow{\phi(U)} & \mathscr{G}(U) \\
 & & \Big\downarrow{\scriptstyle ?} & & \Big\downarrow{\scriptstyle \mathrm{res}_{U,V}} & & \Big\downarrow{\scriptstyle \mathrm{res}_{U,V}} \\
0 & \longrightarrow & \ker_{\mathrm{pre}}\phi(V) & \longrightarrow & \mathscr{F}(V) & \xrightarrow{\phi(V)} & \mathscr{G}(V)
\end{array}$$

The map ? should be the restriction map for $\ker_{\text{pre}} \phi$ if we manage to endow it with a presheaf structure. Indeed, we can take $\text{res}_{U,V}$ and restrict its domain to $\ker_{\text{pre}} \phi(U)$, but it is not clear that it lands in $\ker_{\text{pre}} \phi(V)$. However, we have $\alpha \in \ker_{\text{pre}} \phi(U)$, we have $\phi(U)(\alpha) = 0$, hence $\text{res}_{U,V} \circ \phi(U)(\alpha) = 0$ as well. By commutativity, $\phi(V) \circ \text{res}_{U,V}(\alpha) = 0$, whence $\text{res}_{U,V}(\alpha) \in \ker \phi(V)$ as we wanted. $\square$

Exercise 2.3.F.
Show that the presheaf cokernel $\operatorname{coker}_{\text{pre}} \phi$ satisfies the universal property of cokernels.

Solution.
This is clear when we look at what is happening on open sets. In that case, we have the picture

$$\begin{tikzcd} \mathscr{F}(U) \arrow[r, "\phi"] \arrow[rr, bend right, "0"'] & \mathscr{G}(U) \arrow[r, "p(U)"] & \operatorname{coker}_{\text{pre}} \phi(U) \arrow[r] & 0. \end{tikzcd}$$

Suppose that we had some sheaf $\mathscr{T}$ and $\mathscr{F} \xrightarrow{\phi} \mathscr{G} \xrightarrow{q} \mathscr{T}$ with $q \circ \phi = 0$. Then on every open set, we have the picture

$$\begin{tikzcd} \mathscr{F}(U) \arrow[r, "\phi(U)"] \arrow[rr, bend left, "0"] \arrow[rrd, "0"'] & \mathscr{G}(U) \arrow[r, "p(U)"] \arrow[rd, "q(U)"] & \operatorname{coker}_{\text{pre}} \phi(U) \arrow[d, dashed, "\exists!"] \\ & & \mathscr{T}(U) \end{tikzcd}$$

Since $\operatorname{coker}_{\text{pre}} \phi(U)$ is the cokernel of $\phi(U)$, we have a unique map shown above. This allows us to define a unique map $\operatorname{coker}_{\text{pre}} \phi \to \mathscr{T}$ globally. $\square$

Exercise 2.3.G.
Show (or observe) that for a topological space X with open set U, $\mathscr{F} \mapsto \mathscr{F}(U)$ gives a functor from presheaves of abelian groups on X, $\mathbf{Ab}_X^{\text{pre}}$, to $\mathbf{Ab}$. Show that this functor is exact.

Solution.
We have now taken for granted that $\mathbf{Ab}_X^{\text{pre}}$ is an abelian category, otherwise this question would make no sense. Also, I have observed the fact above, so we move on to exactness.

Suppose that

$$0 \to \mathscr{F} \xrightarrow{a} \mathscr{G} \xrightarrow{b} \mathscr{H} \to 0$$

is an exact sequence of abelian group presheaves on X. Since a and b are defined on open sets, this means that

$$0 \to \mathscr{F}(U) \xrightarrow{a(U)} \mathscr{G}(U) \xrightarrow{b(U)} \mathscr{H}(U) \to 0$$

must still be exact. That is all that we wanted to show. $\square$

Exercise 2.3.H.
Show that a sequence of presheaves $0 \to \mathscr{F}_1 \to \mathscr{F}_2 \to \cdots \to \mathscr{F}_n \to 0$ is exact if and only if $0 \to \mathscr{F}_1(U) \to \mathscr{F}_2(U) \to \cdots \to \mathscr{F}_n(U) \to 0$ is exact for all U.

Solution.
I think based on Vakil's statement that *Homological algebra (exact sequences and so forth) works, and also "works open set by open set"* makes this trivial. The forward direction is obvious (and we have been using it this entire time), and we can build up the appropriate kernels and cokernels to show exactness open set by open set. □

Exercise 2.3.I.
Suppose $\phi : \mathscr{F} \to \mathscr{G}$ is a morphism of sheaves. Show that the presheaf kernel $\ker_{\text{pre}} \phi$ is in fact a sheaf. Show that it satisfies the universal property of kernels.

Solution.
As the problem suggests, assuming that $\ker_{\text{pre}} \phi$ is a sheaf, then it must satisfy the universal property as $\mathbf{Ab}_X$ is a full subcategory of $\mathbf{Ab}_X^{\text{pre}}$.

To show it is a sheaf, we just need to figure out identity and gluability. Since $\ker_{\text{pre}} \phi$ is a subobject of $\mathscr{F}$, we know exactly how to do this. □

Exercise 2.3.J.
Let X be $\mathbb{C}$ with the classical topology, and let $\underline{\mathbb{Z}}$ be the constant sheaf on X associated to $\mathbb{Z}$, $\mathscr{O}_X$ the sheaf of holomorphic functions, and $\mathscr{F}$ the presheaf of functions admitting a holomorphic logarithm. Describe an exact sequence of presheaves on X

$$0 \to \underline{\mathbb{Z}} \to \mathscr{O}_X \to \mathscr{F} \to 0$$

where $\underline{\mathbb{Z}} \to \mathscr{O}_X$ is the natural inclusion and $\mathscr{O}_X \to \mathscr{F}$ is given by $f \mapsto \exp(2\pi i f)$. Show that $\mathscr{F}$ is *not* a sheaf.

Solution.
The natural inclusion is clearly injective. Suppose a function $g \in \mathscr{F}(U)$ admits a holomorphic logarithm $\log g$. Then we may let $f = \frac{1}{2\pi i} \log g \in \mathscr{O}_X(U)$. The map is surjective on open sets, so is overall surjective.

Suppose $\exp(2\pi i f) = 0$ identically on U. Then we must have $f(z) \in \mathbb{Z}$ for all $z \in U$, hence f comes from the inclusion $\underline{\mathbb{Z}}(U) \to \mathscr{O}_X(U)$. This proves exactness.

We are told that $\mathscr{F}$ does not satisfy gluability, so we check this. If we let D_r be the open unit disc centred at the origin of radius r, let $U = D_2 \setminus \overline{D_1}$, an open annulus. Take an open cover of U by

$$A = \{(x, y) \in U : x \in (-1.5, 2)\}, \qquad B = \{(x, y) \in U : x \in (-2, 1.5)\}.$$

Let f_A, f_B be the constant function. Since on both these sets we may make a branch cut either along $\arg z = 0$ or $\arg z = \pi$ to define a holomorphic logarithm, but it cannot be done on the entire annulus, so gluability fails. □

2.4 Properties determined at the level of stalks, and sheafification

Exercise 2.4.A.
Prove that a section of a sheaf of sets is determined by its germs, i.e., the natural map

$$\mathscr{F}(U) \to \prod_{p \in U} \mathscr{F}_p$$

is injective.

Solution.
According to the hints, we will need only the identity axiom to prove this. First, the natural map f above maps an element $\alpha \in \mathscr{F}(U)$ to its germ (α, U) over p. Indeed, this map is induced by the restriction map from $\mathscr{F}(U)$ to some $\mathscr{F}(V)$ for some neighbourhood $p \in V_p \subset U$. Since the identity axioms tells us that two sections in $\mathscr{F}(U)$ restricting identically to an open cover of U (which we have here) are in fact equal, this proves injectivity. $\square$

Exercise 2.4.B.
Show that $\operatorname{supp} s := \{p \in X : s_p \neq 0 \text{ in } \mathscr{F}_p\}$ is a closed subset of X.

Solution.
We will show its complement is open. Suppose that we have $p \in X$ so that $s_p = 0$ in $\mathscr{F}_p$. Then there must be some neighbourhood V of p so that $s|_V = 0$ constantly. Therefore s is locally constantly 0, so $\{p \in X : s_p = 0\}$ is open. $\square$

Exercise 2.4.C.
Prove that any choice of compatible germs for a sheaf of sets $\mathscr{F}$ over U is the image of a section of $\mathscr{F}$ over U.

Solution.
We are pushed towards using gluability. Suppose that we have such a bunch of compatible germs. Then from the definition, we obtain an open cover $\{U_p\}$ of p and sections s'_p so that we have sections (s'_p, U_p). Then we may glue these together to obtain a section on U, i.e. some (s, U) so that $\operatorname{res}_{U,U_p} s = s'_p$. This implies that the value on each of the stalks is als correct so we are done. $\square$

Exercise 2.4.D.
If ϕ_1 and ϕ_2 are morphisms from a presheaf of sets $\mathscr{F}$ to a sheaf of sets $\mathscr{G}$ that induce the same maps on each stalk, show that $\phi_1 = \phi_2$.

Solution.
We use the diagram

$$\begin{array}{ccc}
\mathscr{F}(U) & \xrightarrow[\phi_2]{\phi_1} & \mathscr{G}(U) \\
\downarrow & & \downarrow \\
\prod_{p \in U} \mathscr{F}_p & \xrightarrow{\prod \phi_1 = \prod \phi_2} & \prod_{p \in U} \mathscr{G}_p
\end{array}$$

We may do even better: since the right vertical map is injective, it is invertible from its image. Therefore we need only check that taking ϕ_1 on a stalk yields a compatible germ. But this is obvious. Therefore since we have equality on the bottom maps, we have equality of the top maps. □

Exercise 2.4.E.
Show that a morphism of sheaves of sets is an isomorphism if and only if it induces an isomorphism of all stalks.

Solution.
Using the above picture again, let $\phi : \mathscr{F}(U) \to \mathscr{G}(U)$ be a morphism of sheaves of sets, and let $\iota_{\mathscr{F}}$ and $\iota_{\mathscr{G}}$ be the natural inclusions of the set of sections into the product of stalks.

$$\begin{array}{ccc} \mathscr{F}(U) & \xrightarrow{\phi} & \mathscr{G}(U) \\ \big\downarrow \iota_{\mathscr{F}} & & \big\downarrow \iota_{\mathscr{G}} \\ \prod_{p\in U} \mathscr{F}_p & \xrightarrow{\prod \phi_p} & \prod_{p\in U} \mathscr{G}_p \end{array}$$

We refine this to

$$\begin{array}{ccc} \mathscr{F}(U) & \xrightarrow{\phi} & \mathscr{G}(U) \\ \big\downarrow \sim & & \big\downarrow \sim \\ \operatorname{im} \iota_{\mathscr{F}} & \xrightarrow{\prod \phi_p|_{\operatorname{im} \iota_{\mathscr{F}}}} & \operatorname{im} \iota_{\mathscr{G}} \end{array}$$

If either the top or the bottom of this diagram is an isomorphism, then the other one is an isomorphism as well. As stated above, it is clear that the bottom map takes compatible germs to compatible germs, so the bottom map is well-defined, so we are done. □

Exercise 2.4.F.

(a) Show that Exercise 2.4.A is false for general presheaves.

(b) Show that Exercise 2.4.D is false for general presheaves.

(c) Show that Exercise 2.4.E is false for general presheaves.

Solution.

(a) The hint overall is to use the two-point space $X = \{p, q\}$ with the discrete topology. Define $\mathscr{F}$ as follows: we let $\mathscr{F}(\{p\} = \mathscr{F}(\{q\}) = \{\alpha\}$, some one point set, and let $\mathscr{F}(X) = \{a, b\}$, some two point set. This is a presheaf, since we let the restriction maps $\operatorname{res}_{X,\{p\}}$ and $\operatorname{res}_{X,\{q\}}$ be the constant maps with value α (as they must since one point sets are a final object in **Sets**).

Now, the map $\mathscr{F}(X) \to \mathscr{F}_p \times \mathscr{F}_q = \{(\alpha, \alpha)\}$ is a set map from a two point set to a one point set, which cannot be injective.

(b) We will continue using the above sheaf. Let $\mathscr{G}$ be very similar to $\mathscr{F}$: we will let $\mathscr{G}(X) = \{c, d\}$ and $\mathscr{G}(\{p\}) = \mathscr{G}(\{q\}) = \{\beta\}$. Let $\phi_1, \phi_2 : \mathscr{F} \to \mathscr{G}$ be defined as follows: we must have $\phi_1(\alpha) = \beta$ on $\mathscr{F}_p$ and $\mathscr{F}_q$. On $\mathscr{F}(X)$, let $\phi_1(a) = \phi_1(b) = c$ and let $\phi_2(a) = \phi_2(b) = d$. Then we have two different sheaf morphisms which agree on stalks, but are not equal.

(c) Take ϕ_1 above. By construction, the bottom map is an isomorphism on stalks, but we do not have an isomorphism overall as $\phi_1(X)$ is not surjective.

$\square$

Exercise 2.4.G.
Show that sheafification is unique up to unique isomorphism, assuming it exists. Show that if $\mathscr{F}$ is a sheaf, then the sheafification is $\mathrm{id} : \mathscr{F} \to \mathscr{F}$.

Solution.
Suppose that $\mathscr{F}_1^{\mathrm{sh}}$ and $\mathscr{F}_2^{\mathrm{sh}}$ are two sheafifications of $\mathscr{F}$. Then these are both sheaves, so we have a diagram

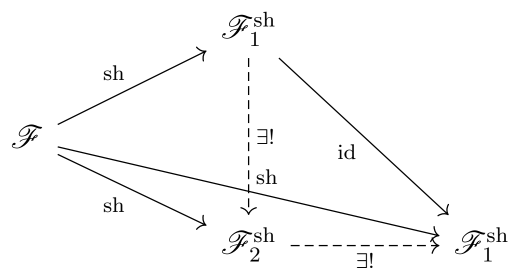

The two dashed morphisms must compose to be the identity, since that is the unique map making the top triangle commute, so we have a unique isomorphism from $\mathscr{F}_1^{\mathrm{sh}}$ to $\mathscr{F}_2^{\mathrm{sh}}$. Further, if $\mathscr{F}$ is itself a sheaf, $\mathscr{F}$ satisfies the universal property of $\mathscr{F}^{\mathrm{sh}}$ by letting $\mathrm{sh} : \mathscr{F} \to \mathscr{F}$ be given by the identity, so it must be its own sheafification. $\square$

Exercise 2.4.H.
Assume for now that sheafification exists. Use the universal property to show that for any morphism of presheaves $\phi : \mathscr{F} \to \mathscr{G}$, we get a natural induced morphism of sheaves $\phi^{\mathrm{sh}} : \mathscr{F}^{\mathrm{sh}} \to \mathscr{G}^{\mathrm{sh}}$. Show that sheafification is a functor from presheaves on X to sheaves on X.

Solution.
We use the diagram below:

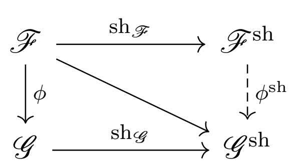

We know there is a composite map $\mathscr{F} \to \mathscr{G}^{\mathrm{sh}}$ by $\phi \circ \mathrm{sh}_{\mathscr{G}}$. By the universal property of sheafification, there is a unique map $\mathscr{F}^{\mathrm{sh}} \to \mathscr{G}^{\mathrm{sh}}$ making the diagram commute, which we denote by the dashed morphism ϕ^{sh}.

Since we have assumed how sh acts on objects of $\mathbf{Ab}_X^{\mathrm{pre}}$ and now have shown how it works on morhpisms, we have a good candidate for a functor. That composition works is clear, and that sh takes the identity to the identity follows from 2.4.G. There fore we have a functor $\mathrm{sh} : \mathbf{Ab}_X^{\mathrm{pre}} \to \mathbf{Ab}_X$. $\square$

Exercise 2.4.I.
Show that $\mathscr{F}^{\mathrm{sh}}$ forms a sheaf.

Solution.
We already know what the restriction maps should be, so we have a presheaf. Suppose that $\{U_i\}$ is a open cover of U, and suppose that $f, g \in \mathscr{F}^{\text{sh}}(U)$ are two sections that agree on each restriction onto $\mathscr{F}^{\text{sh}}(U_i)$, which we denote f_i, g_i respectively. We know that we have

$$f_i = (f_{i,x} \in \mathscr{F}_x : x \in U_i,\ \exists V_i \subset U_i \text{ and } s_i \in \mathscr{F}(V_i) \text{ s.t. } s_{i,y} = f_{i,y}\, \forall y \in V_i)$$

and similar for g_i. Because of this constructions, we see that all the $f_{i,x}$ and $g_{i,x}$ must be the same, which would make $f = g$ overall.

Now suppose that we have $f_i \in \mathscr{F}^{\text{sh}}(U_i)$ with the intersections agreeing appropriately. Then suppose we take the section $f \in \mathscr{F}^{\text{sh}}(U)$ obtained by taking each member of the stalks $f_{i,x}$, since we know that if $x \in U_i \cap U_j$, then by assumption $f_{i,x} = f_{j,x}$. We know that this collection is a set of compatible germs because, at each $x \in U$, there is some U_i with $x \in U_i$. As such there is an open $V_i \subset U_i$ such that the set of stalks is the image of a section. Since $V_i \subset U$ is still open, we may take this as our neighbourhood. Thus we have glued our sections together as required, so we have a sheaf. □

Exercise 2.4.J.
Describe a natural map of presheaves $\text{sh} : \mathscr{F} \to \mathscr{F}^{\text{sh}}$.

Solution.
We know that we have a natural maps

$$\mathscr{F}(U) \to \prod_{p \in U} \mathscr{F}_p \to \mathscr{F}^{\text{sh}}(U).$$

This should be our natural map of presheaves. Since $\mathscr{F}^{\text{sh}}(U)$ is a subobject of $\prod_{p \in U} \mathscr{F}_p$, the second map will make sense in abelian categories (or **Sets**), so we are fine. For it to be a map of presheaves, it would have to play well with restriction maps, i.e. for $V \subset U$

$$\begin{array}{ccccc} \mathscr{F}(U) & \longrightarrow & \prod_{p \in U} \mathscr{F}_p & \longrightarrow & \mathscr{F}^{\text{sh}}(U) \\ \downarrow{\scriptstyle \text{res}_{U,V}} & & \downarrow & & \downarrow{\scriptstyle \text{res}_{U,V}} \\ \mathscr{F}(V) & \longrightarrow & \prod_{q \in V} \mathscr{F}_q & \longrightarrow & \mathscr{F}^{\text{sh}}(V) \end{array}$$

It is clear that the left and right squares commute individually, so we have the required commutative diagram above. □

Exercise 2.4.K.
Show that the map sh satisfies the universal property of sheafification.

Solution.
Suppose that $\mathscr{G}$ is a sheaf and $\phi : \mathscr{F} \to \mathscr{G}$ a map of presheaves. We will first construct the map $\psi : \mathscr{F}^{\text{sh}} \to \mathscr{G}$ then show it is unique. Let $f \in \mathscr{F}^{\text{sh}}(U)$. Construct an open cover U_i of U so that on each U_i we have a section $f_i \in \mathscr{F}(U_i)$ and $f_x = f_{i,x}$ for all $x \in U_i$.

Now, consider $g_i = \phi(f_i) \in \mathscr{G}(U_i)$. Since $\mathscr{G}$ is a sheaf and it is clear that the g_i agree on intersections by the above work, we may glue the g_i together into a unique $g \in \mathscr{G}(U)$ which agrees on each of the restrictions. We want to define $\psi(f) := g$ for this g. It is also clear that $\psi \circ \mathrm{sh} = \phi$ from this, and it is unique because we were forced by commutativity to define ψ this way. □

Exercise 2.4.L.
Show that the sheafification functor is left adjoint to the forgetful functor from sheaves on X to presheaves on X.

Solution.
Let $u : \mathcal{C}_X \to \mathcal{C}_X^{\mathrm{pre}}$ be the forgetful functor. Then we need to show that, for a presheaf $\mathscr{F}$ and sheaf $\mathscr{G}$,

$$\mathscr{H}om(\mathrm{sh}(\mathscr{F}), \mathscr{G}) \cong \mathscr{H}om(\mathscr{F}, u(\mathscr{G})).$$

We know that given any map from a presheaf to a sheaf, we have a unique map from its sheafification to that sheaf. This gives an inclusion of the right set in the left. Now, given a map from $\mathscr{F}^{\mathrm{sh}} \to \mathscr{G}$, we can precompose with the natural map (abusing the hell out of notation) $\mathrm{sh} : \mathscr{F} \to \mathscr{F}^{\mathrm{sh}}$ to obtain a map $\mathscr{F} \to \mathscr{G}$, which means the opposite inclusion also holds. Therefore we have a bijection between these two sets. Depending on the category $\mathcal{C}$ we will want a stronger isomorphism, but this follows easily. □

Exercise 2.4.M.
Show $\mathscr{F} \to \mathscr{F}^{\mathrm{sh}}$ induces an isomorphism of stalks.

Solution.
We will take the 'set of germs' definition of stalks. Recall that $\mathscr{F}_p$ is the set of equivalence classes (f, U) where $p \in U$ is open, $f \in \mathscr{F}(U)$, and $(f, U) \sim (g, V)$ if there is an open $W \subset U \cap V$ with $\mathrm{res}_{U,W}(f) = \mathrm{res}_{V,W}(g)$.

We know that, locally, each $f \in \mathscr{F}^{\mathrm{sh}}(U)$ is some tuple of stalk elements (f_x) which (locally enough) are the image a section $\mathscr{F}(V)$ with $V \subset U$. Therefore associated to a stalk element f_x, $x \in U$, we have a section $s \in \mathscr{F}(V)$ so that $s_x = f_x$. That is, f_x is representable by (s, V). Therefore we have an injection $\mathscr{F}_x^{\mathrm{sh}} \to \mathscr{F}_x$. Further, given some germ (s, V) in $\mathscr{F}_x$, we may construct the element $(s_y, y \in V) \in \prod_{p \in U} \mathscr{F}_p$. Then this is a compatible germ, hence is an element of $\mathscr{F}^{\mathrm{sh}}(U)$. Therefore s_x is an element of the stalk $\mathscr{F}_x^{\mathrm{sh}}$. This is the bijection which implies an isomorphism. □

Exercise 2.4.N.
Suppose $\phi : \mathscr{F} \to \mathscr{G}$ is a morphism of sheaves of sets on a topological space X. Show that the following are equivalent:

(a) ϕ is a monomorphism in the category of sheaves.

(b) ϕ is injective on the level of stalks.

(c) ϕ is injective on the level of open sets.

Solution.
We will take Vakil's recommendations: since we have shown that morphisms are determined by the action on stalks for sheaves, (b) implies (a) and (b) implies (c), since we can view stalks either globally or locally. (c) also implies (a). We therefore need only show that (a) implies (c).

Vakil suggests for this to use the 'indicator sheaf' for U, which we will denote $\mathscr{T}_U$ for lack of any other notation. ϕ being a monomorphism means that, for any two maps $\psi_1, \psi_2 : \mathscr{T} \to \mathscr{F}$, we have

$$\phi \circ \psi_1 = \phi \circ \psi_2 \implies \psi_1 = \psi_2.$$

Therefore let us try this with the indicator sheaf $\mathscr{T}_U$. Suppose that $\phi(U) : \mathscr{F}(U) \to \mathscr{G}(U)$ is not injective, so that $\phi(U)(s) = \phi(U)(t)$ for some $s \neq t \in \mathscr{F}(U)$. Then let $\psi_1, \psi_2 : \mathscr{T}_U \to \mathscr{F}$ be defined as follows: if $\{a\}$ is the set associated to each $\mathscr{T}_U(V)$ for $V \subset U$ open, then we would like

$$\psi_i(V)(a) = \begin{cases} \mathrm{res}_{U,V}\, s & i = 1 \\ \mathrm{res}_{U,V}\, t & i = 2 \end{cases},$$

where we include the possibility $V = U$ so that $\mathrm{res}_{U,U} = \mathrm{id}$. This does define two sheaf morphisms such that $\phi \circ \psi_1 = \phi \circ \psi_2$. However, we have $\psi_1 \neq \psi_2$ since $s \neq t$ by assumption, so we have a contradiction. Therefore $\phi(U)$ must be injective on each U. □

Exercise 2.4.O.
Continuing the notation of the previous exercise, show that the following are equivalent:

(a) ϕ is an epimorphism in the category of sheaves.

(b) ϕ is surjective on the level of stalks.

Solution.
(b) implies (a) is clear, as it was above. For (a) implies (b), we take Vakil's hint and use the skyscraper sheaf with value $\{a, b, c\}$ at p, which we denote $\mathscr{T}_p$. If ϕ is an epimorphism then for any $\psi_1, \psi_2 : \mathscr{G} \to \mathscr{T}$ such that $\psi_1 \circ \phi = \psi_2 \circ \phi$, then we have $\psi_1 = \psi_2$.

Suppose $\phi_p : \mathscr{F}_p \to \mathscr{G}_p$ is not surjective, and let $g \in \mathscr{G}_p$ not be in the image of ϕ_p. Then let us define maps $\psi_1, \psi_2 : \mathscr{G} \to \mathscr{T}_p$ by defining what happens on stalks: since $\mathscr{T}_{p,q}$, that is the stalk of $\mathscr{T}_p$ at q, is equal to $\{a, b, c\}$ for $q = p$ and the final object in **Sets** elsewhere, let us define for $s \in \mathscr{G}_p$,

$$\psi_{i,p}(s) = \begin{cases} a & s = g, i = 1 \\ b & s = g, i = 2 \\ c & s \neq g \end{cases}.$$

There are sheaf morphisms ψ_1 and ψ_2 corresponding to these sheaf maps, and these satisfy $\psi_1 \circ \phi = \psi_2 \circ \phi$. However by construction $\psi_1 \neq \psi_2$, so ϕ cannot be an epimorphism. This proves the contrapositive of the converse, so we are done. □

Exercise 2.4.P.
Let $X = \mathbb{C}$ with the classical topology, let $\mathcal{O}_X$ be the sheaf of holomorphic functions, and let $\mathcal{O}_X^\times$ be the sheaf of invertible (nowhere zero) holomorphic functions. Show that $\mathcal{O}_X \xrightarrow{\exp} \mathcal{O}_X^\times$ describes $\mathcal{O}_X^\times$ as a quotient sheaf of $\mathcal{O}_X$. Find an open set on which this map is not surjective.

Solution.
We know that the exp sheaf map is an epimorphism if and only if $\exp_p$ is surjective for all $p \in X$. Therefore we will show this. From Exercise 2.3.J, we know that we have an exact sequence of (pre)sheaves if we replace $\mathscr{O}_X^\times$ by holomorphic functions admitting a holomorphic logarithm. From complex analysis, we can define the log of any function whose codomain is a subset of $\mathbb{C} \setminus \{0\}$ by

$$(\log f)(z) = \int_\gamma \frac{f'(z)}{f(z)}\, dz,$$

where γ is a path between 0 and z (which one does not matter since $\mathbb{C}$ is simply connected). Therefore we can write $f = \exp(\log f)$. Since the stalks in this case are just locally-defined functions supported at p, this shows that every $\exp_p$ is surjective.

We will run into the same problem if we try to define functions on an annulus. Let U be the open annulus centred at 0 constructed in Exercise 2.3.J. Then the constant function has no logarithm, so the map $\exp(U) : \mathscr{O}_X(U) \to \mathscr{O}^\times(U)$ cannot be surjective. □

2.5 Sheaves of abelian groups, and $\mathscr{O}_X$-modules, form abelian categories

Exercise 2.5.A.
Show that the stalk of the kernel is the kernel of the stalks.

Solution.
We want a natural isomorphism

$$(\ker(\mathscr{F} \to \mathscr{G}))_x \cong \ker(\mathscr{F}_x \to \mathscr{G}_x).$$

Let $\phi : \mathscr{F} \to \mathscr{G}$ for ease of notation. Let $\ker\phi$ denote the object of the kernel of ϕ in the category of sheaves. Suppose we have an element $f \in (\ker\phi)_x$ (we will let $\ker\phi_x$ be the right hand side of the above). Then since we have a natural monomorphism $\ker\phi \to \mathscr{F}$, we have an injection $(\ker\phi)_x \to \mathscr{F}_x$ (assuming we are working over a nice category, otherwise we just have a monomorphism). As such we can map $(\ker\phi)_x \to \mathscr{G}_x$ via ϕ_x. We summarise this in the picture

$$\begin{array}{ccccccc} 0 & \longrightarrow & \ker\phi & \longrightarrow & \mathscr{F} & \xrightarrow{\phi} & \mathscr{G} \\ & & \downarrow & & \downarrow & & \downarrow \\ 0 & \longrightarrow & (\ker\phi)_x & \longrightarrow & \mathscr{F}_x & \xrightarrow{\phi_x} & \mathscr{G}_x \end{array}$$

The top row of this diagram is exact. By diagram chasing, we can see the bottom is exact too. Therefore $(\ker\phi)_x$ satisfies the requirements of $\ker\phi_x$, so the two must be isomorphic. □

Exercise 2.5.B.
Show that the stalk of the cokernel is naturally isomorphic to the cokernel of the stalk.

Solution.
Take the above proof and look at cokernels instead. Using additionally the fact that the presheaf cokernel and the sheaf cokernel have isomorphic stalks (Poblem 2.4.M), we are done. □

Exercise 2.5.C.
Suppose that $\phi : \mathscr{F} \to \mathscr{G}$ is a morphism of sheaves of abelian groups. Show that the image sheaf $\operatorname{im}\phi$ is the sheafification of the image presheaf. Show that the stalk of the image is the image of the stalk.

Solution.
We have now shown that we are working in an abelian category. We know that $\operatorname{im}\phi = \ker\operatorname{coker}\phi$ in this case. Taking kernels will always give us a sheaf, but we do not know that $\operatorname{coker}\phi$ is a sheaf a priori. Writing $_{\text{pre}}$ when necessary, we see

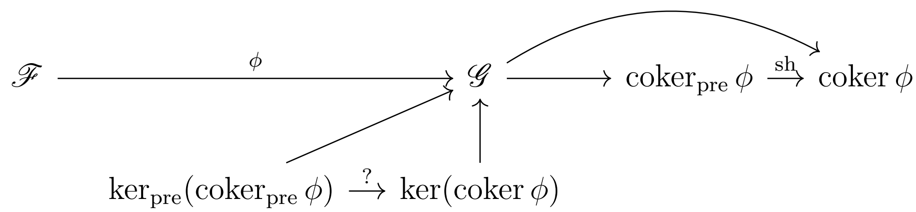

We want ? to be the sheafification map. As such, we will show that $\ker(\operatorname{coker}\phi)$ satisfies the universal property of the sheafification of $\ker_{\text{pre}}(\operatorname{coker}_{\text{pre}}\phi)$. Write $\mathscr{H}^{\text{pre}}$ for $\ker_{\text{pre}}(\operatorname{coker}_{\text{pre}}\phi)$ and write $\mathscr{H}$ for its sheafification. By the above, the composite map $\mathscr{H}^{\text{pre}} \to \operatorname{coker}\phi$ is zero, which is a map from a presheaf into a sheaf. Therefore it factors uniquely through to $\mathscr{H} \to \operatorname{coker}\phi$. However, since $\ker_{\text{pre}}(\operatorname{coker}\phi) = \ker(\operatorname{coker}\phi)$, we know that $\mathscr{H}^{\text{pre}}$ must also factor uniquely through $\ker(\operatorname{coker}(\phi))$. Since this is another map between $\mathscr{H}^{\text{pre}}$ and a sheaf, we have another factorisation through $\mathscr{H}$. This gives

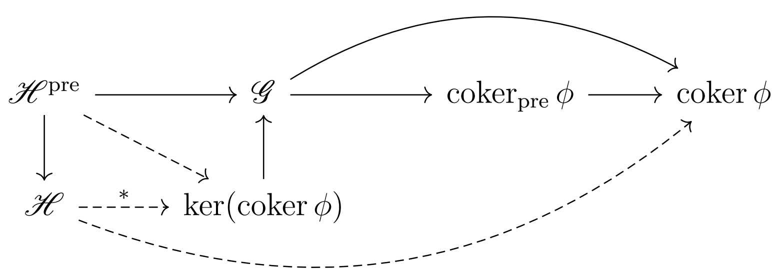

Since this diagram commutes, if the starred dashed arrow is not the identity, then we have two different ways to move from $\mathscr{H}^{\text{pre}}$ through $\mathscr{H}$ to $\operatorname{coker}\phi$. Thus it must be the identity, so $\ker(\operatorname{coker}\phi)$ is the sheafification of the presheaf image.

That the stalk of the image is the image of the stalk follows from the previous two exercises. □

Exercise 2.5.D.
Show that talking the stalk of a sheaf of abelian groups is an exact functor.

Solution.
Consider the sheaf category $\mathbf{Ab}_X$. Let $p \in X$ be a point. Suppose that

$$0 \to \mathscr{F} \xrightarrow{\phi} \mathscr{G} \xrightarrow{\psi} \mathscr{H} \to 0$$

is an exact sequence of sheaves. We want to show that

$$0 \to \mathscr{F}_p \xrightarrow{\phi_p} \mathscr{G}_p \xrightarrow{\psi_p} \mathscr{H}_p \to 0$$

is also exact. A criterion for exactness in the first sequence is that $\operatorname{im}\phi = \ker\psi$. We know that

$$(\operatorname{im}\phi)_p = \operatorname{im}\phi_p, \qquad (\ker\psi)_p = \ker\psi_p$$

by the above problems. Therefore since $(\operatorname{im}\phi)_p = (\ker\psi)_p$, the sequence in **Ab** is exact since its maps satisfy the above criterion. □

Exercise 2.5.E.
To check the exponential sequence

$$0 \to \underline{\mathbb{Z}} \xrightarrow{\cdot 2\pi i} \mathscr{O}_x \xrightarrow{\exp} \mathscr{O}_X^* \to 1$$

is exact, we know that we may check it on stalks. Let $p \in \mathbb{C}$ be any point and let $\iota = \cdot 2\pi i$. We have shown before that the first map is a monomorphism and the second map is an epimorphism, i.e. ι_p is injective on stalks and $\exp_p$ is surjective on stalks. We need only show that $\operatorname{im}\iota_p = \ker\exp_p$.

We know that $\operatorname{im}\iota_p \subset \ker\exp_p$ since $\exp(2\pi i m) = 0$ for all $m \in \mathbb{Z}$, so the function described is the zero function. Further, let $f \in \mathscr{O}_X$ such that $\exp(f) = 0$. Then we know that $f(x) \in 2\pi i\mathbb{Z}$ for all $x \in U$, where U is some small open set containing p. Since f takes values on a discrete set, we know that there is some smaller neighbourhood of p so that f is constant, with value $f(p) = 2\pi i n$ (for some n). This neighbourhood can be described as the connected component containing p of $f^{-1}(2\pi i n)$. Because we are working on stalks, f and the constant function $2\pi i n$ represent the same germ, so we know that f comes from $n \in \underline{\mathbb{Z}}_p$. This completes the proof.

Exercise 2.5.F.
Suppose $U \subset X$ is an open set, and $0 \to \mathscr{F} \to \mathscr{G} \to \mathscr{H}$ is an exact sequence of sheaves of abelian groups. Show that

$$0 \to \mathscr{F}(U) \to \mathscr{G}(U) \to \mathscr{H}(U)$$

is exact. Show that the section functor need not be exact.

Solution.
For the second part of this problem, we showed in 2.4.P that $\mathscr{O}_X(U) \to \mathscr{O}_X^*(U)$ is not always surjective, e.g. for U the open annulus. Therefore we would not have right exactness for that section functor.

We will do this the long way, not using the fact that all right adjoints are left exact. We started this process in Exercise 2.4.N: if $\mathscr{F} \to \mathscr{G}$ is a monomorphism, then $\mathscr{F}(U) \to \mathscr{G}(U)$ is injective. Therefore this sequence is exact at $\mathscr{F}(U)$. Therefore we need only show that the image of $\mathscr{F}(U)$ in $\mathscr{G}(U)$ is the kernel of $\mathscr{G}(U) \to \mathscr{H}(U)$. We know that any exact sequence can be factored into a short exact sequence. Exactness at $\mathscr{G}$ follows is equivalent to

$$\text{for } \mathscr{F} \xrightarrow{\phi} \mathscr{G} \xrightarrow{\psi} \mathscr{H}, \quad 0 \to \operatorname{im}\phi \to \mathscr{G} \to \operatorname{im}\psi \to 0 \text{ is short exact.}$$

Since we have shown that $\operatorname{im}\phi = \ker\psi$,we need to show that $\ker\psi(U) \to \mathscr{G}(U) \to \operatorname{im}\psi(U)$ is short exact. But this is obvious. □

Exercise 2.5.G.
Suppose $0 \to \mathscr{F} \xrightarrow{\phi} \mathscr{G} \xrightarrow{\psi} \mathscr{H}$ is an exact sequence in $\mathbf{Ab}_X$. If $\pi : X \to Y$ is a continuous map, show that

$$0 \to \pi_*\mathscr{F} \xrightarrow{\pi_*\phi} \pi_*\mathscr{G} \xrightarrow{\pi_*\psi} \pi_*\mathscr{H}$$

is exact.

Solution.
We need to show that $\operatorname{im} \pi_*\phi = \ker \pi_*\psi$. It would be nice to check this on stalks, but there is no clear way of writing down what the stalks in this case are, so we will check it on open sets instead. We can check this on open stalks because we do not need right exactness of π_*; that would cause problems only at exactness at $\pi_*\mathscr{H}$. Now, we know that

$$\pi_*\phi(U) = \phi(\pi^{-1}(U)), \quad \pi_*\psi(U) = \psi(\pi^{-1}(U)).$$

The condition we have to check boils down thus:

$$\operatorname{im} \pi_*\phi(U) = \ker \pi_*\psi(U) \iff \operatorname{im} \phi(\pi^{-1}(U)) = \ker \psi(\pi^{-1}(U)).$$

But this is obviously true by assumption. □

Exercise 2.5.H.
Suppose $\mathscr{F}$ is a sheaf of abelian groups on a topological space X. Show that $\mathscr{H}om(\mathscr{F}, -)$ is a left exact covariant functor $\mathbf{Ab}_X \to \mathbf{Ab}_X$. Show that $\mathscr{H}om(-, \mathscr{F})$ is a left exact contravariant functor.

Solution.
Let $\mathscr{G}$ and $\mathscr{H}$ be two other sheaves, and let $\phi : \mathscr{G} \to \mathscr{H}$ be a sheaf morphism. Since it is already obvious that $\mathscr{H}om(\mathscr{F}, -)$ takes one sheaf to the homomorphism sheaf, we need to demonstrate its behaviour on morphisms. We see that $\mathscr{H}om(\mathscr{F}, -)$ induces a map $\bar{\phi} : \mathscr{H}om(\mathscr{F}, \mathscr{G}) \to \mathscr{H}om(\mathscr{F}, \mathscr{H})$ given by $\psi \mapsto \phi \circ \psi$. We see that $\mathscr{H}om(-, \mathscr{F})$ induces a map $\bar{\phi} : \mathscr{H}om(\mathscr{H}, \mathscr{F}) \to \mathscr{H}om(\mathscr{G}, \mathscr{F})$ given by $\psi \mapsto \psi \circ \phi$.

We need only show left-exactness. We cannot use stalks again, because $\mathscr{H}om(\mathscr{F}, \mathscr{G})_p \neq \operatorname{Hom}(\mathscr{F}_p, \mathscr{G}_p)$ in general. We can use the same open set argument as we did above, though, since we do not need to worry about the last place. The proof is virtually identical. □

Exercise 2.5.I.
Show that if $(X, \mathcal{O}_X)$ is a ringed space, then $\mathcal{O}_X$-modules form an abelian category.

Solution.
This is almost immediate since we know that sheaves of abelian groups form an abelian category. We only need to notice that subobjects, quotient objects, coproducts, and products of $\mathcal{O}_X$-modules have a natural $\mathcal{O}_X$-module structure, so the sheaves of abelian groups satisfying various universal properties actually lie in the category of $\mathcal{O}_X$-modules. □

Exercise 2.5.J.

(a) Suppose $\mathcal{O}_X$ is a sheaf of rings on X. Define (categorically) what should mean by the *tensor product of two $\mathcal{O}_X$-modules.* Give an explicit construction, and show that it satisfies your categorical definition.

(b) Show that the tensor product of stalks is the stalk of the tensor product.

Solution.

(a) What it should be is the following: let $\mathscr{F}$ and $\mathscr{G}$ be two $\mathcal{O}_X$-modules. Then suppose we have a map $\phi : \mathscr{F} \oplus \mathscr{G} \to \mathscr{T}$ of $\mathcal{O}_X$-modules so that on each $U \subset X$, ϕ is $\mathcal{O}_X(U)$-linear in both arguments. There may be a better way to define $\mathcal{O}_X$-linearity in this case, but I do not see what it is.

We want to define the tensor product as the morphism $\mathscr{F} \oplus \mathscr{G} \to \mathscr{F} \otimes_{\mathcal{O}_X} \mathscr{G}$ such that for any such ϕ as above, there is a unique map ψ such that the diagram below commutes:

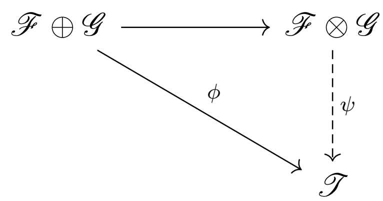

Our guess for this explicit construction is, for $p \in X$, define the stalks of the tensor sheaf to be

$$(\mathscr{F} \otimes \mathscr{G})_p = \mathscr{F}_p \otimes \mathscr{G}_p.$$

There is absolutely no guarantee what we obtain is a sheaf, so we sheafify for good measure. We will keep the same notation. We need to show that this definition works. Since morphisms are defined on stalks, we know that there is a unique morphism of sheaves that corresponds to the stalk morphisms we propose. Therefore it satisfies the universal property we described.

$\square$

2.6 The inverse image sheaf

Exercise 2.6.A.

Let $\pi : X \to Y$ be a continuous map and let $\mathscr{G}$ be a sheaf on Y. Then let

$$\pi^{-1}\mathscr{G}^{\text{pre}}(U) = \varinjlim_{V \supset \pi(U)} \mathscr{G}(V).$$

Note that $\pi(U)$ is not open in general. Show that this is a presheaf, but not necessarily a sheaf, on X.

Solution.

We will show that we have the structure of a contravariant functor. Fortunately, the restriction maps are easy: suppose that $V \subset U \subset X$. Then clearly $\pi(V) \subset \pi(U)$. Therefore if we have $f \in \pi^{-1}\mathscr{G}^{\text{pre}}(U)$, we can represent it by an element (x, W) for some open $W \supset \pi(U)$.

Then $W \supset \pi(V)$ is still an open set, so in fact (x, W) represents an element of $\pi^{-1}\mathscr{G}^{\text{pre}}(V)$. We need only figure out what element of the colimit it corresponds to, and this is the value of $\text{res}_{U,V}(f)$.

To show it is not a sheaf in general, consider $Y = \{p\}$ and $X = \{a, b\}$, where π is the constant map. Give X the discrete topology. If $\mathscr{G}$ is the constant sheaf S on Y, where S is some set, then since the only (nontrivial) open set in Y is Y itself, we must have

$$\pi^{-1}\mathscr{G}^{\text{pre}}(U) = \mathscr{G}(Y) = S$$

for any open U in X. This makes $\pi^{-1}\mathscr{G}^{\text{pre}}$ the constant sheaf on X. We know from previous work that this is not a sheaf if we choose S to be a set with more than one element. □

Exercise 2.6.B.
If $\pi : X \to Y$ is a continuous map, and $\mathscr{F}$ is a sheaf on X and $\mathscr{G}$ is a sheaf on Y, describe a bijection

$$\text{Mor}_X(\pi^{-1}\mathscr{G}, \mathscr{F}) \leftrightarrow \text{Mor}_Y(\mathscr{G}, \pi_*\mathscr{F}).$$

Observe that this bijection is functorial in both $\mathscr{F}$ and $\mathscr{G}$ (i.e. 'natural'). Thus π^{-1} satisfies the universal property of the left adjoint of π_*.

Solution.
We will do this using Vakil's hint. Consider $U \subset X$ and $V \subset Y$ open, and let $\phi_{VU} : \mathscr{G}(V) \to \mathscr{F}(U)$ be any map. Then we call a collection $\phi = \{\phi_{UV} : U \subset X, V \subset Y \text{ open}\}$ *compatible* if the following holds: for all open $U' \subset U$ and $V' \subset V$ with $\pi(U') \subset V'$, we require the diagram

$$\begin{array}{ccc} \mathscr{G}(V) & \xrightarrow{\phi_{VU}} & \mathscr{F}(U) \\ \downarrow{\scriptstyle \text{res}_{V,V'}} & & \downarrow{\scriptstyle \text{res}_{U,U'}} \\ \mathscr{G}(V') & \xrightarrow{\phi_{V'U'}} & \mathscr{F}(U') \end{array}$$

to commute. We let $\text{Mor}_{YX}(\mathscr{G}, \mathscr{F})$ be the set of all compatible collections of maps. Vakil suggest we show that both of the above sets are equal to this one we have just created.

Now, suppose we have a map $\phi : \pi^{-1}\mathscr{G} \to \mathscr{F}$. We want to show that we can turn this map into a compatible collection. Let $U \subset X$. Since $\pi^{-1}\mathscr{G}(U)$ is a direct limit of $\mathscr{G}(V)$ for $V \supset \pi(U)$ open, we can construct a map ϕ_{VU} for all of these U: each $x \in \mathscr{G}(V)$ corresponds to some element $y \in \pi^{-1}\mathscr{G}(U)$, so let $\phi_{VU}(x) = \phi(U)(y)$. This commutes with restriction maps because we know that $\phi(U)$ does. This process clearly works backwards as well.

Suppose we have a map $\psi : \mathscr{G} \to \pi_*\mathscr{F}$. Let $V \subset Y$ be open. Since $\pi_*\mathscr{F}(V) = \mathscr{F}(\pi^{-1}(V))$, we can construct a map $\psi_{V\pi^{-1}(V)}$ in the obvious way. For $x \in \mathscr{G}(V)$, let $\psi_{V\pi^{-1}(V)}(x) = \psi(V)(x)$. This process is clearly reversible as well. Therefore we have shown these two sets are the same. Showing functoriality should be trivial, because we require the nice commutativity of these diagrams (which would be the only snag). □

Exercise 2.6.C.
Show that the stalks of $\pi^{-1}\mathscr{G}$ are the same as the stalks of $\mathscr{G}$. More precisely, if $\pi(p) = q$, describe a natural isomorphism $\mathscr{G}_q \cong (\pi^{-1}\mathscr{G})_p$.

Solution.
Using adjointness should be the easiest way. We know that left adjoints commute with colimits, and that stalks are defined as colimits. Therefore since

$$\pi^{-1}\mathscr{G}^{\text{pre}}(U) = \varinjlim_{V \supset \pi(U)} \mathscr{G}(V),$$

we have

$$(\pi^{-1}\mathscr{G}^{\text{pre}})_p = \varinjlim_{p \in U} \pi^{-1}\mathscr{G}^{\text{pre}}(U) = \pi^{-1} \varinjlim_{p \in U} \mathscr{G}^{\text{pre}}(U) = \pi^{-1}(\mathscr{G}^{\text{pre}})_p = (\mathscr{G}^{\text{pre}})_q.$$

Since presheaves and sheaves have the same stalks, this completes the proof. $\square$

Exercise 2.6.D.
If U is an open subset of Y, $i : U \to Y$ is the inclusion, and $\mathscr{G}$ is a sheaf on Y, show that $i^{-1}\mathscr{G}$ is naturally isomorphic to $\mathscr{G}|_U$.

Solution.
Let us look at the definition. Suppose that W is an open set in U. Then by definition,

$$i^{-1}\mathscr{G}^{\text{pre}}(W) = \varinjlim_{V \supset i(W)} \mathscr{G}(V).$$

In particular, if we look at the diagram implied by the colimit, it has a terminal object, namely W itself. Since the whole diagram must commute through the arrow $W \to i(W)$, we have that the colimit itself must be $\mathscr{G}(W)$. Therefore the data of $i^{-1}\mathscr{G}^{\text{pre}}$ is the sheaf data of every open set contained in U, which is just $\mathscr{G}|_U$. $\square$

Exercise 2.6.E.
Show that π^{-1} is an exact functor from sheaves of abelian groups on Y to sheaves of abelian groups on X.

Solution.
Taking the hint, we know that we can check exactness on stalks, and by 2.6.C, we have an isomorphism of stalks under π^{-1}. Therefore suppose we have an exact sequence

$$0 \to \mathscr{F} \to \mathscr{G} \to \mathscr{H} \to 0$$

in $\mathbf{Ab}_Y$. Then at every point $p \in X$ with $\pi(p) = q$, we have

$$\begin{array}{ccccccccc} 0 & \longrightarrow & \mathscr{F}_q & \longrightarrow & \mathscr{G}_q & \longrightarrow & \mathscr{H}_q & \longrightarrow & 0 \\ & & \sim\downarrow & & \sim\downarrow & & \sim\downarrow & & \\ 0 & \longrightarrow & (\pi^{-1}\mathscr{F})_p & \longrightarrow & (\pi^{-1}\mathscr{G})_p & \longrightarrow & (\pi^{-1}\mathscr{H})_p & \longrightarrow & 0 \end{array}$$

Since the top row is exact, the bottom row is also exact. Therefore π^{-1} is an exact functor. $\square$

Exercise 2.6.F.

(a) Suppose $Z \subset Y$ is a closed subset, and $i : Z \to Y$ is the inclusion. If $\mathscr{F}$ is a sheaf of sets on Z, then show that the stalk $(i_*\mathscr{F})_q$ is a one element set if $q \notin Z$, and $\mathscr{F}_q$ if $q \in Z$.

(b) Define the *support* of a sheaf $\mathscr{G}$ of sets, denoted $\operatorname{supp}\mathscr{G}$, as the locus where the stalks are not the one-element set:

$$\operatorname{supp}\mathscr{G} := \{p \in X : |\mathscr{G}_p| \neq 1\}.$$

Suppose $\operatorname{supp}\mathscr{G} \subset Z$ where Z is closed. Show that the natural map $\mathscr{G} \to i_*i^{-1}\mathscr{G}$ is an isomorphsm. Thus a sheaf supposed on a closed subset can be considered a sheaf on that closed subset.

Solution.

(a) We have, for V open in Y,

$$(i_*\mathscr{F})_q = \varinjlim_{q \in V} i_*\mathscr{F}(V) = \varinjlim_{q \in V} \mathscr{F}(i^{-1}(V)).$$

Since $i^{-1}(V) = Z \cap V$, which is open in Z, then we have two cases. If $q \in Z$, then every open set (in Z) containing q can be described as $Z \cap V$ for some V open in Y. Therefore we lose no data here, and we are taking the colimit over (essentially) the same diagram, whence $(i_*\mathscr{F})_q \cong \mathscr{F}_q$. If $q \notin Z$, then $V \cap Z$ will be empty for sufficiently small V (because Z is a closed set). In this case, the colimit is being taken over empty sets at the end of it, whence $(i_*\mathscr{F})_q$ must be the final object of **Sets**, i.e. a one-element set.

(b) This is asking if the unit is a natural isomorphism. This occurs when the left adjoint is fully faithful, which is provable in an abstract categorical way. To see this: let L and R be an adjoint pair between $\mathcal{C}$ and $\mathcal{D}$. If L is fully faithful, then we have for any $A, B \in \mathcal{C}$,

$$\operatorname{Mor}_{\mathcal{C}}(A, B) \cong \operatorname{Mor}_{\mathcal{D}}(L(A), L(B)) \cong \operatorname{Mor}_{\mathcal{C}}(A, R \circ L(B)),$$

where the second isomorphism follows from adjunction. Since this holds for all $A, B \in \mathcal{C}$, we must have $R \circ L$ naturally isomorphic to $1_{\mathcal{C}}$.

In our case, we just need to show that i^{-1} is fully faithful, i.e. that the first isomorphism as above actually holds. But in this case, since $\operatorname{supp}\mathscr{G} \subset Z$, taking i^{-1} does not affect the data of any 'important' open sets; any sheaf map ϕ can only be defined in one way outside of $\operatorname{supp}\mathscr{G}$ since the one-point set is final. Therefore i^{-1} is fully faithful in this case, so we have the required natural isomorphism.

□

Exercise 2.6.G.

Suppose $i : U \to Y$ is the inclusion of an open set into Y. Define the *extension of i by zero* $i_! : \mathbf{Mod}_{\mathcal{O}_U} \to \mathbf{Mod}_{\mathcal{O}_Y}$ as follows. Suppose $\mathscr{F}$ is an $\mathcal{O}_U$-module. For open $W \subset Y$, define $(i_!^{\text{pre}}\mathscr{F})(W) = \mathscr{F}(W)$ if $W \subset U$ and 0 otherwise. This is clearly a presheaf $\mathcal{O}_Y$-module. Define $i_!$ as $(i_!^{\text{pre}})^{\text{sh}}$. Note that $i_!\mathscr{F}$ is an $\mathcal{O}_Y$-module, and that this defines a functor.

(a) Show that $i_!^{\text{pre}}\mathscr{F}$ need not be a sheaf.

(b) For $q \in Y$, show that $(i_!\mathscr{F})_q = \mathscr{F}_q$ if $q \in U$ and 0 otherwise.

(c) Show that $i_!$ is an exact functor.

(d) If $\mathscr{G}$ is an $\mathscr{O}_Y$-module, describe an inclusion $i_! i^{-1}\mathscr{G} \to \mathscr{G}$.

(e) Show that $(i_!, i^{-1})$ is an adjoint pair, so there is a natural bijection

$$\operatorname{Hom}_{\mathscr{O}_Y}(i_!\mathscr{F}, \mathscr{G}) \leftrightarrow \operatorname{Hom}_{\mathscr{O}_U}(\mathscr{F}, \mathscr{G}|_U)$$

for any $\mathscr{O}_U$-module $\mathscr{F}$ and $\mathscr{O}_Y$-module $\mathscr{G}$.

Solution.
Even though Vakil says we shouldn't do this question right now, it's worth a shot.

(a) The recommendation here is to try the sheaf of continuous functions on $\mathbb{R}$ to find a counterexample. Let $\mathscr{F}$ be this sheaf. Let U be the disjoint union $(0,1) \cup (2,3)$. Then $i_!^{\text{pre}}\mathscr{F}(1,2) = 0$. However, we can find continuous functions on $(0,1)$ and $(2,3)$ that cannot be glued together, since the restriction of the gluing on $(1,2)$ must be 0, which is not necessarily true. Concretely, the constant function 1 on $(0,1)$ and $(2,3)$ cannot be glued (continuously) to be 0 on $(1,2)$.

(b) If we don't need this point to prove (e), then it follows because left adjoints commute with colimits. As it stands, however, I think we should prove this directly. Recalling that the stalks of a presheaf and its sheafification are isomorphic,

$$(i_!^{\text{pre}}\mathscr{F})_q = \varinjlim_{q \in V} i_!^{\text{pre}}\mathscr{F}(V).$$

If $q \in U$, then eventually all open V with $q \in V$ are also contained in U, so that $i_!^{\text{pre}}\mathscr{F}(V) = \mathscr{F}(V)$. Then it is clear that $(i_!^{\text{pre}}\mathscr{F})_q \cong \mathscr{F}_q$. If $q \notin U$, then sufficiently many V with $q \in V$ yields $i_!^{\text{pre}}\mathscr{F}(V) = 0$ to force $\mathscr{F}_q = 0$.

(c) (b) shows that $i_!$ is exact since it is easily seen to be exact on the level of stalks. If $q \in U$, then it follows from the reasoning of 2.6.E. If $q \notin U$, then the 'sequence' on stalks is identically zero, which is trivially exact.

(d) We can try to define this on the level of open sets. Let $W \subset Y$ be open. We have

$$i_! i^{-1}\mathscr{G}(W) = i_!\mathscr{G}|_U(W) = \begin{cases} \mathscr{G}|_U(W) & W \subset U \\ 0 & \text{else} \end{cases}$$

The inclusion $\mathscr{G}|_U(W) \subset \mathscr{G}(W)$ or $0 \subset \mathscr{G}(W)$ is perfectly well defined.

(e) We will establish this bijection as best we can. Suppose that $\phi : i_!\mathscr{F} \to \mathscr{G}$ is an $\mathscr{O}_Y$-module homomorphism. Suppose that we have $W \subset U$ open. Then we have a map

$$\phi(W) : i_!\mathscr{F}(W) = \mathscr{F}(W) \to \mathscr{G}(W).$$

Since $W \subset U$, we have $\mathscr{G}|_U(W) = \mathscr{G}(W)$, so $\phi(W)$ is a perfectly well defined map on $\mathscr{F}(W) \to \mathscr{G}|_U(W)$ as well. If W is not contained in U, then $i_!\mathscr{F}(W) = 0$. Therefore $\phi(W)$ must be the (initial) zero map. But fortunately, this makes $\mathscr{G}|_U(W) = 0$ too, so $\phi(W) : \mathscr{F}(W) \to \mathscr{G}|_U(W)$ is the (final) zero map. Therefore we have an inclusion of $\mathrm{Hom}_{\mathscr{O}_Y}(i_!\mathscr{F}, \mathscr{G}) \subset \mathrm{Hom}_{\mathscr{O}_U}(\mathscr{F}, \mathscr{G}|_U)$. The reverse inclusion follows identically. That we have a bijection follows easily as well: the pairs $i_!\mathscr{F}$ and $\mathscr{F}$ and $\mathscr{G}$ and $\mathscr{G}|_U$ have the same stalks at the applicable open sets, so we can check to see if we get the same map back by checking stalks.

□

2.7 Recovering sheaves from a "sheaf on a base"

Exercise 2.7.A.
How can you recover a sheaf $\mathscr{F}$ from the partial information on some base $\{B_i\}$ of the topology on X?

Solution.
We know how to recover the stalks now. The stalk $\mathscr{F}_p$ is obtained by taking the colimit over all open sets containing p. But we know for every arbitrary open U containing p, there is a base element B so that $p \in B \subset U$. Put another way, a germ over p can be represented by some (f, U), but it can also be represented as (g, B) where B is a base element. This is enough information to give us the data of the stalks. We can reconstruct compatible germs by gluing the stalks back together for each open $U \subset X$ and, if $U = \bigcup i \in IB_i$, using the known data of $\mathscr{F}(B_i)$. This is probably precise enough. □

Exercise 2.7.B.
Verify that $F(B) \to \mathscr{F}(B)$ is an isomorphism, likely by showing that it is injective and surjective.

Solution.
We have defined F to be a sheaf (of sets) on the base $\{B_i\}$. To write back down the formal definition of $\mathscr{F}$:

$$\mathscr{F}(U) := \{(f_p \in F_p)_{p\in U} : \forall p \in U, \exists B \text{ with } p \in B \subset U, s \in F(B), \text{ with } s_q = f_q \,\forall q \in B\}.$$

It is the set of all compatible germs where we make sure that the open subset of U we pick is in the base. That sentence alone pretty much proves what we'd like.

Let us look at the natural map $F(B) \to \mathscr{F}(B)$. Let us first show it is surjective. Given some set (f_p) of the above form, we may examine the elements (s_q) guaranteed by the definition. This is an element of $F(B')$ for some $B' \subset B$. As such, there is an element of $F(B)$ which restricts down to (s_q), and this would then be a preimage of (f_p).

Injectivity is also clear: if we have (g_p) and (h_p) mapping to the same $(f_p) \in \mathscr{F}(B)$, then we know that (g_p) and (h_p) must agree on all subbase elements of $F(B)$, which (by identity) would make them the same element. This gives us the required isomorphism. $\square$

Exercise 2.7.C.
Suppose $\{B_i\}$ is a base for the topology of X. A morphism $F \to G$ of sheaves on the base is a collection of maps $F(B_k) \to G(B_k)$ such that the diagram

$$\begin{array}{ccc} F(B_i) & \longrightarrow & G(B_i) \\ {\scriptstyle \mathrm{res}_{B_i,B_j}}\big\downarrow & & \big\downarrow{\scriptstyle \mathrm{res}_{B_i,B_j}} \\ F(B_j) & \longrightarrow & G(B_j) \end{array}$$

commutes for all $B_j \subset B_i$.

(a) Verify that a morphism of sheaves is determined by the induced morphism of sheaves on the base.

(b) Show that a morphism of sheaves on the base gives a morphism of the induced sheaves.

Solution.
We may now assume Theorem 2.7.1, which is very good. I think we can be very clever about this. We know that $\mathscr{H}om(\mathscr{F}, \mathscr{G})$ is itself a sheaf on X, and so it is uniquely determined by its data on a base, i.e. the restriction maps and the data of

$$\mathscr{H}om(\mathscr{F}, \mathscr{G})(B) = \mathrm{Hom}(\mathscr{F}(B), \mathscr{G}(B)) \cong \mathrm{Hom}(F(B), G(B)),$$

where the last isomorphism comes from 2.7.B. Therefore it seems that this proves (a) and (b) simultaneously. This might be a little too clever, and in particular incorrect. $\square$

Exercise 2.7.D.
Suppose $X = \bigcup U_i$ is an open cover of X, and we have sheaves $\mathscr{F}_i$ on U_i along with isomorphisms $\phi_{ij} : \mathbb{F}_i|_{U_i \cap U_j} \to \mathscr{F}_j|_{U_i \cap U_j}$ that agree on triple overlaps, i.e. $\phi_{jk} \circ \phi_{ij} = \phi_{ik}$ on $U_i \cap U_j \cap U_k$. Show that these sheaves can be glued together into a sheaf $\mathscr{F}$ on X (unique up to unique isomorphism) such that $\mathscr{F}_i \cong \mathscr{F}|_{U_i}$, and the isomorphisms over $U_i \cap U_j$ are the obvious ones.

Solution.
We know that each $\mathscr{F}_i$ (for $i \in I$) is determined by its information on a base $\{B_j^i\}$ for $j \in J_i$ some index set. Then we claim that

$$B = \bigcup_{i \in I} \bigcup_{j \in J_i} B_j^i$$

is a base for X. Indeed, let $V \in X$ be open. Then we can write $V = \bigcup V_i$, where $V_i = U_i \cap V$ is open. In turn, each $V_i = \bigcup B_k^i$ since each of those sets are a base for U_i. Therefore we have

$$V = \bigcup_{i \in I} \bigcup_{k \in J_i} B_k^i,$$

so we indeed have a base for X. The conditions above guarantee that we can come up with a well defined sheaf F on the base B. To be more explicit, we gluability and identity hold because they hold on each U_i and they agree on double and triple intersections. Therefore we have a sheaf $\mathscr{F}$ on X from F on B. Everything else, including unique up to unique isomorphism, follows from the theorem. □

Exercise 2.7.E.
Suppose a morphism of sheaves $F \to G$ on a base B_i is surjective for all B_i. Show that the corresponding morphism of sheaves is surjective (or more precisely, an epimorphism). The converse is not true, unlike the case for injectivity.

Solution.
Suppose we had some $\phi(U) : \mathscr{F}(U) \to \mathscr{G}(U)$ which was not surjective. Let $q \in \mathscr{G}(U)$ be an element not in the image of $\phi(U)$. Then $q \in \mathscr{G}(B)$ for some basis element B. As such q is in the image of $\phi(B)$. This is a flagrant contradiction. □

3 Toward affine schemes: the underlying set, and topological space

3.1 Toward schemes

Exercise 3.1.A.
Suppose that $\pi : X \to Y$ is a continuous map of differentiable manifolds (as topological spaces). Show that π is differentiable if differentiable functions pull back to differentiable functions, i.e. if pullback by π gives a map $\mathcal{O}_Y \to \pi_*\mathcal{O}_X$.

Solution.
Since we are working with differentiable manifolds, let V be an open subset of Y diffeomorphic to $\mathbb{R}^n$ (for whatever appropriate n). Then we shall construct the map $\pi^\sharp : \mathcal{O}_Y \to \pi_*\mathcal{O}_X$ on V, using the terminology from 3.1.1.

First, what is $\pi_*\mathcal{O}_X(V)$? It is precisely $\mathcal{O}_X(\pi^{-1}(V))$, $\pi^{-1}(V)$ being (for a properly chosen V) a chart on X. Therefore what is the map $\pi^\sharp$? Given a function $f \in \mathcal{O}_Y$, we should have $\pi^\sharp(f) = f \circ \pi$. If this map is differentiable, then we know by the chain rule that π itself must be differentiable, as

$$(f \circ \pi)'(x) = f'(\pi(x)) \cdot \pi'(x).$$

In particular, we can extract $\pi'(x)$ using only the data from the derivatives of f and $f \circ \pi$ that we are guaranteed exist. Since we have shown what we must on a base of the topology for Y, we are done because now we need only glue everything together. □

Exercise 3.1.B.
Show that a morphism of differentiable manifolds $\pi : X \to Y$ with $\pi(p) = q$ induces a morphism of stalks $\pi^\sharp : \mathcal{O}_{Y,q} \to \mathcal{O}_{X,p}$. Show that $\pi^\sharp(\mathfrak{m}_{Y,q}) \subset \mathfrak{m}_{X,p}$.

Solution.
Consider a germ $(f, V) \in \mathcal{O}_{Y,q}$. Then since π is a morphism of differentiable manifolds, we are to assume that it is a differentiable function, so we have the map $\pi^\sharp$ described above. We such, we can take $(f, V) \mapsto (f \circ \pi, V)$. Then we know that this is a perfectly fine differentiable map in this neighbourhood V from 3.1.A, so this is how we get a map on stalks. If we had $f(p) = 0$, then we would have $f(\pi(q)) = f(p) = 0$, so $\pi^\sharp(f, V)$ would also vanish. This shows the inclusion of the maximal ideals. □

3.2 The underlying set of affine schemes

Exercise 3.2.A.

(a) Describe the set $\operatorname{Spec} k[\varepsilon]/(\varepsilon^2)$. The ring $k[\varepsilon]/(\varepsilon^2)$ is called the ring of *dual numbers*. You should think of ε as a very small number, so small that its square is 0 (although it itself is not 0).

(b) Describe the set $\operatorname{Spec} k[x]_{(x)}$.

Solution.

(a) The ring consists of all $a + b\varepsilon$, where $a, b \in k$. We know that all choices where $b = 0$ force the ideal $(a + b\varepsilon) = k$. If $b \neq 0$, then we can represent $(a + b\varepsilon)$ by $(a/b + \varepsilon)$, so we can just write $(a + \varepsilon)$. We see that

$$(a + \varepsilon)(b + \varepsilon) = ab + (a + b)\varepsilon,$$

since the ε^2 term vanishes. In particular, if $a \neq 0$, then we can choose $b = -a$ so that $(a + \varepsilon)(b + \varepsilon) \in k$.

All this is to say that the only prime ideal seems to be (ε). Therefore $\operatorname{Spec} k[\varepsilon]/(\varepsilon^2) = \{(\varepsilon)\}$.

(b) Recall what $k[x]_{(x)}$ means: we invert all elements outside of (x), since this is a multiplicatively closed set. Since the only prime ideals in $k[x]$ were of the form $(x - a)$ for $a \in k$, this means the only prime ideals we have left are (x) and (0). This spectrum therefore also has only two points.

□

Exercise 3.2.B.
Show that for prime ideals in $\mathbb{R}[x]$ of the form $(x^2 + ax + b)$ with that polynomial irreducible, the quotient $\mathbb{R}[x]/(x^2 + ax + b) \cong \mathbb{C}$ always.

Solution.
Up to a linear change of variables, (in particular, $x \mapsto y - a/2$), we can rewrite

$$x^2 + ax + b = y^2 + c$$

for some other constant c. This c will necessarily be positive, because a linear change of variables does not change the discriminant of the polynomial. Since $\Delta = b - 4ac$ and we require $\Delta < 0$ for irreducibility, we have $c > 0$. We know that $\mathbb{R}[y]/(y^2 + c) = \mathbb{R}[\sqrt{-c}] \cong \mathbb{C}$. Since the linear change of variables is also an isomorphism on the quotient spaces, we are done. □

Exercise 3.2.C.
Describe the set $\mathbb{A}^1_{\mathbb{Q}}$.

Solution.
What are the prime ideals here? We are still in a PID, so we need only find the irreducible polynomials. We know we need only to consider monic polynomials in $\mathbb{Q}[x]$. We know that monic irreducible polynomials are completely determined by their roots. Further, we know that these roots must come in Galois-conjugate (to use Vakil's term) tuples; e.g., we could not have $\sqrt{2}$ a root without $-\sqrt{2}$ also a root.

Therefore it seems that every irreducible polynomial f corresponds to some set $\{a_1, \ldots, a_n\}$ of Galois-conjugate tuples of complex numbers. Further, given any particular a_i, we can recover the rest of the set of conjugates because they must be the roots of the minimal polynomial of a_i in $\mathbb{Q}[x]$, which is the f we started with.

This means we have a surjection $\bar{\mathbb{Q}} \to \mathbb{A}^1_{\mathbb{Q}}$, and the kernel of this map is the action of the absolute Galois group $\operatorname{Gal}(\bar{\mathbb{Q}}, \mathbb{Q})$. Trying to make this analogous to $\mathbb{A}^1_{\mathbb{R}}$, in that case we had to glue together all complex conjugates. In this case, we have to glue together all Galois-conjugates. Indeed, we had

$$\mathbb{A}^1_{\mathbb{R}} = \bar{\mathbb{R}} / \operatorname{Gal}(\bar{\mathbb{R}}, \mathbb{R}) = \mathbb{C}/(x = \bar{x}).$$

This situation here is analogous. □

Exercise 3.2.D.
If k is a field, show that $\operatorname{Spec} k[x]$ has infinitely many primes.

Solution.
We know that prime ideals correspond to irreducible polynomials. If k is a field of characteristic zero, it is plain to see that $(x - p)$ for each $p \in \mathbb{N}$ prime are an infinite set of prime ideals. If k has characteristic p, then the existence of finite fields of order p^n for all n guarantees irreducible polynomials of every degree in $k[x]$. These too are an infinite set of prime ideals. □

Exercise 3.2.E.
Show that the only prime ideals of $\mathbb{C}[x, y]$ are of the form $(x - a, y - b)$ or $(f(x, y))$ for an irreducible polynomial $f \in \mathbb{C}[x, y]$.

Solution.
We might not have a PID, but $\mathbb{C}[x, y]$ is a UFD. We know that every $(f(x, y))$ is a prime ideal, so assume that $\mathfrak{p}$ is an ideal that is not principal. Suppose that no $f(x, y)$ and $g(x, y)$ in $\mathfrak{p}$ were relatively prime. Then there would be some highest degree polynomial $h(x, y) \in \mathfrak{p}$ dividing every element of $\mathfrak{p}$, which implies that $\mathfrak{p} \subset ((h(x, y))$. Since $h(x, y) \in \mathfrak{p}$, this shows that $\mathfrak{p}$ is principal, a contradiction.

Now, let $f(x, y)$ and $g(x, y)$ be these relatively prime polynomials. Then $(f(x, y), g(x, y)) \subset \mathfrak{p}$. As implied by the hint, since $\mathbb{C}[x, y] = \mathbb{C}[x][y]$, we can use obtain the greatest common divisor of $f(x, y)$ and $g(x, y)$ in $\mathbb{C}[x]$, which is some $h(x) \neq 0$. Since $h(x) \in \mathfrak{p}$ and $\mathfrak{p}$ is a prime ideal, and because $h(x)$ splits over $\mathbb{C}$, some prime divisor (i.e. irreducible factor) of

$h(x)$ is contained in $\mathfrak{p}$, say $(x-a)$. The same process for $\mathbb{C}[x,y] = \mathbb{C}[y][x]$ shows that there is a $(y-b) \in \mathfrak{p}$. But $(x-a, y-b)$ is a maximal ideal, so if $\mathfrak{p}$ is a proper ideal of $\mathbb{C}[x,y]$ and $(x-a, y-b) \subset \mathfrak{p}$, this must actually be an equality. □

Exercise 3.2.F.
Show that the Nullstellensatz implies the Weak Nullstellensatz

Solution.
Writing down what they want, we assume that if k is a field, then every maximal ideal of $k[x_1, \ldots, x_n]$ has a residue field a finite extension of k.

Since $\mathbb{C}$ is an algebraically closed field, it has no finite extensions. Therefore every maximal ideal of $\mathbb{C}[x_1, \ldots, x_n]$ must have residue field isomorphic to $\mathbb{C}$ itself. This only occurs in the stated case, where $\mathfrak{m} = (x_1 - a_1, \ldots, x_n - a_n)$. It suffices to check what happens if we take one generator of a maximal ideal to be a non-linear irreducible polynomial. In this case, the residue field clearly contains a nontrivial residue class for some x_i, so it cannot be isomorphic to $\mathbb{C}$. For example, if we had $f(x,y) = y^2 - x^3$, then $x \in \mathbb{C}[x,y]/(f(x,y))$ is nontrivial, so $\mathbb{C}[x,y]/(f(x,y))$ is not isomorphic to $\mathbb{C}$. □

Exercise 3.2.G.
Prove that any integral domain A which is a finite k-algebra must be a field (not requiring the Nullstellensatz).

Solution.
We're almost all of our way to a field anyway. We need only show that we are working in a division algebra, i.e. every element is invertible. This is implied by showing, taking the hint that $\ell_x : A \to A$, $y \mapsto x \cdot y$ is an isomorphism for all nonzero $x \in A$.

We certainly know that ℓ_x is injective; since A is an integral domain, we have cancellation. If $\ell_x(a) = \ell_x(b)$, then

$$\ell_x(a) - \ell_x(b) = \ell_x(a-b) = x(a-b) = 0.$$

Since we took $x \neq 0$, then we must have $a - b = 0$, i.e. $a = b$. Since ℓ_x is a ring homomorphism (not hard to see), it is a fortiori a vector space homomorphism. An injective map between vector spaces of the same dimension must be a vector space isomorphism. In particular, the map is surjective. Therefore it is also surjective as a ring homomorphism, and therefore a ring isomorphism. This shows that there exists $c \in A$ so that $\ell_x(c) = 1_A$ for every $x \in A \backslash \{0\}$. Therefore A is a commutative division domain, so it is a field. □

Exercise 3.2.H.
Describe the maximal ideal of $\mathbb{Q}[x,y]$ corresponding to $(\sqrt{2}, \sqrt{2})$ and $(-\sqrt{2}, -\sqrt{2})$. Describe the maximal ideal of $\mathbb{Q}[x,y]$ corresponding to $(\sqrt{2}, -\sqrt{2})$ and $(-\sqrt{2}, \sqrt{2})$. What are the residue fields in both cases?

Solution.
From what was discussed in 3.2.C, we know that for $\mathbb{A}^1_{\mathbb{Q}}$, $\sqrt{2}$ and $-\sqrt{2}$ correspond to the ideal generated by $x^2 - 2$ in $\mathbb{Q}[x]$. Similarly, the ideal corresponding to $(\sqrt{2}, \sqrt{2})$ should be

just $(xy-2, x^2-2)$, which is definitely maximal because we see its quotient is a field. When we mod out by this ideal, we get $K = \mathbb{Q}(\sqrt{2})$ from the x^2-2 contribution, and y is forced to be $\sqrt{2}$ when we let $x = \sqrt{2}$ and $-\sqrt{2}$ when $x = -\sqrt{2}$. Therefore the whole quotient is just K.

In the other case, we want to pick the maximal ideal $(xy+2, x^2-2)$. We get the same field extension but y is forced to be the opposite sign as x for the two choices of x. $\square$

Exercise 3.2.I.
Consider the map of sets $\phi : \mathbb{C}^2 \to \mathbb{A}^2_{\mathbb{Q}}$ defined as follows: (z_1, z_2) is sent to the prime ideal of $\mathbb{Q}[x, y]$ consisting of polynomials vanishing at (z_1, z_2).

(a) What is the image of (π, π^2)?

(b) Show that ϕ is surjective.

Solution.

(a) The image should be the zero ideal. It is known that π is not an algebraic number. We proved above that prime ideals of $\mathbb{Q}[x, y]$ correspond to $\bar{\mathbb{Q}}^2$ under the gluing of Galois-conjugates. Since $\pi \notin \bar{\mathbb{Q}}$, we must map (π, π^2) trivially to (0).

(b) Vakil says we need to use other machinery to solve this, but it seems doable regardless. We know that prime ideals of $\mathbb{Q}[x, y]$ are going to correspond to irreducible polynomials in x, in y, or in both. We know that every irreducible polynomial in $\mathbb{Q}[x]$ has a root in $\mathbb{C}$. If we have a prime ideal of the form $(f(x), g(y))$, then choosing a root z_1 of $f(x)$ and a root z_2 of $g(y)$ will make a preimage (z_1, z_2) for this ideal. This is because any polynomial (in x) vanishing at z_1 must contain the factor

$$\prod_{\sigma \in \mathrm{Gal}(\bar{\mathbb{Q}}, \mathbb{Q})} (x - \sigma(z_1))$$

for all Galois conjugates of z_1 (without superfluous overrepresentation). This is the minimal polynomial of z_1, hence irreducible, so it must be the $f(x)$ we started with. The same goes for $g(y)$.

For ideals of the form $(f(x, y))$ for an irreducible f with nontrivial degree in both x and y, we need to be slightly more creative. Suppose that $f(x, y)$ did not have a root in $\mathbb{C}^2$. Then for any fixed $b \in \mathbb{C}$, the (degree at least one) polynomial $f(x, b) \in \mathbb{C}[x]$ would have no root. But this is a contradiction. Therefore $f(x, y)$ has a root (a, b) in $\mathbb{C}^2$, and this root is a preimage. This is probably enough work to satisfy the solution at this point in time.

$\square$

Exercise 3.2.J.
Suppose A is a ring, and I an ideal of A. Let $\phi : A \to A/I$. Show that ϕ^{-1} gives an inclusion-preserving bijection between primes of A/I and primes of A containing I.

Solution.
This is a standard commutative algebra problem. Let $\mathfrak{p} \in A/I$ be a prime ideal. Then $\phi^{-1}(\mathfrak{p})$ is prime for the following reason: suppose that $ab \in \phi^{-1}(\mathfrak{p})$. Then $\phi(ab) = \phi(a)\phi(b) \in \mathfrak{p}$, hence (without loss of generality) $\phi(a) \in \mathfrak{p}$. Therefore $a \in \phi^{-1}(\mathfrak{p})$. Clearly $\phi^{-1}(\mathfrak{p})$ contains $\phi^{-1}((0)) = I$, so this proves one direction.

Conversely, if $\mathfrak{q}$ is a prime ideal of A, then $\phi(\mathfrak{q})$ is prime in A/I. This generally holds for any surjective map. Suppose that $ab \in \phi(\mathfrak{q})$. Then $\phi^{-1}(a)$ and $\phi^{-1}(b)$ are both nonempty, since ϕ is surjective, so let $a' \in \phi^{-1}(a)$ and $b' \in \phi^{-1}(b)$. Then $a'b' \in \mathfrak{q}$, so we have (without loss of generality) $a' \in \mathfrak{q}$, whence $a \in \phi(\mathfrak{q})$. If $\mathfrak{q}$ does not contain I, then taking $\phi^{-1}(\phi(\mathfrak{q}))$ gives a prime ideal containing I, so we get back something different from what we began with. Therefore if we demand that $\mathfrak{q} \supset I$, we have the required bijection. That inclusion is preserved is completely obvious. □

Exercise 3.2.K.
Suppose S is a multiplicative subset of A. Describe an order-preserving bijection of the primes of $S^{-1}A$ with the primes of A that do not meet S.

Solution.
We have a natural inclusion (though strictly speaking it may not be injective) $\phi : A \to S^{-1}A$ such that $a \mapsto a/1$. Suppose that $\mathfrak{p}$ is a prime ideal of $S^{-1}(A)$. Then we know that $\mathfrak{p}$ does not contain any invertible elements. In particular, $\mathfrak{p}$ does not contain any elements of the form $s/1$ since these are invertible. Since every noninvertible element of $S^{-1}A$ is in the image of ϕ, we can take $\phi^{-1}(\mathfrak{p})$, which (as discussed above) is a prime ideal that necessarily avoids S.

Now suppose that $\mathfrak{q}$ is a prime ideal of A. If $\mathfrak{q} \cap S \neq \varnothing$, then $\phi(\mathfrak{q})$ contains an invertible element, whence the ideal generated by this set cannot be a prime ideal. If $\mathfrak{q} \cap S = \varnothing$, then we claim the ideal generated by $\phi(\mathfrak{q})$ is prime ideal. Let $\phi(\mathfrak{q}) = \mathfrak{p}$. Suppose that

$$\frac{a}{s} \cdot \frac{b}{t} \in \mathfrak{p}.$$

Then we have $st \cdot \frac{a}{s} \cdot \frac{b}{t} = ab \in \mathfrak{p}$. Then since ab is in the image of $\mathfrak{q}$, we know that (without loss of generality) $a \in \mathfrak{q}$. Therefore $a/s \in \mathfrak{p}$. This is the desired bijection. □

Exercise 3.2.L.
Show that $(\mathbb{C}[x, y]/(xy))_x \cong \mathbb{C}[x]_x$ (in some natural way).

Solution.
$\mathbb{C}[x]_x$ should have a pretty easy spectrum, since from previous problems we know $\operatorname{Spec} \mathbb{C}[x]_x$ is just those prime ideals in $\mathbb{C}[x]$ avoiding $\{1, x, x^2, \ldots\}$. Since (nonzero) prime ideals of $\mathbb{C}[x]$ are of the form $(x - a)$ for any $a \in \mathbb{C}$, we just need to avoid the choice $a = 0$. Any other choice of a gives an ideal not containing x, for if $x \in (x - a)$ then $-a \in (x - a)$ which implies $(x - a) = \mathbb{C}[x]$, so it is not a maximal ideal. Therefore we can realise $\operatorname{Spec} \mathbb{C}[x]_x = \mathbb{C} \setminus \{0\}$.

We know from the last few problems that $\operatorname{Spec} \mathbb{C}[x, y]/(xy)$ can be viewed as a subset of $\operatorname{Spec} \mathbb{C}[x, y]$, which we have discussed. Specifically, $\operatorname{Spec} \mathbb{C}[x, y]/(xy)$ 'contains' those prime ideals of $\operatorname{Spec} \mathbb{C}[x, y]$ which contain (xy). Since prime ideals of $\mathbb{C}[x, y]$ are (by 3.2.E) just of the form $(x-a, y-b)$ or $(f(x,y))$, we know what to ignore. Since xy is irreducible, it never shows up in anything of the form $(f(x,y))$ unless $f(x,y) = xy$, and this corresponds to the zero element in $\operatorname{Spec} \mathbb{C}[x,y]/(xy)$. Therefore the nonzero prime ideals in $\operatorname{Spec} \mathbb{C}[x,y]/(xy)$ are limited to those of the form $(x-a, y-b)$. But we know that the element

$$(x-a)(y-b) = xy - bx - ay + ab = ab - bx - ay \in (x-a, y-b).$$

We do some more manipulation to show

$$ab - bx - ay + b(x-a) + a(y-b) = -ab \in (x-a, y-b).$$

This shows in particular that we must have $a = 0$ or $b = 0$. Therefore we want ideals of the form $(x, y-b)$ or $(x-a, y)$ for $a, b \in \mathbb{C}$. Now, to deal with the localisation, we need to throw out all choices of $a = 0$ and everything of the form $(x, y-b)$. Therefore the primes we have left are of the form $(x-a, y)$ for $a \in \mathbb{C} \setminus \{0\}$. That is precisely what we concluded above, and we can see how the correspondence works now:

$$\operatorname{Spec} \mathbb{C}[x]_x \ni (x-a) \leftrightarrow (x-a, y) \in \operatorname{Spec}(\mathbb{C}[x,y]/(xy))_x$$

□

Exercise 3.2.M.
If $\phi : B \to A$ is a map of rings, and $\mathfrak{p}$ is a prime ideal of A, show that $\phi^{-1}(\mathfrak{p})$ is a prime ideal of B.

Solution.
I'm not sure how we could've gotten away without proving this by this point. See Exercise 3.2.J. □

Exercise 3.2.N.
Let B be a ring.

(a) Suppose $I \subset B$ is an ideal. Show that the map $\operatorname{Spec} B/I \to \operatorname{Spec} B$ is the inclusion map implied by 3.2.J.

(b) Suppose $S \subset B$ is a multiplicative set. Show that the map $\operatorname{Spec} S^{-1}B \to \operatorname{Spec} B$ is the inclusion implied by 3.2.K.

Solution.
I'm reasonably sure I proved this well enough above to satisfy both myself and Vakil. □

Exercise 3.2.O.
Consider the map of complex manifolds sending $\mathbb{C} \to \mathbb{C}$ via $x \mapsto y = x^2$. We interpret the domain as the 'x-line' and the codomain as the 'y-line'. Interpret the corresponding map of rings given by $\mathbb{C}[y] \to \mathbb{C}[x]$ given by $y \mapsto x^2$. Verify that the fibre above the point $a \in \mathbb{C}$ is the point(s) $\pm\sqrt{a} \in \mathbb{C}$, using the definition given above.

Solution.
I think I know what they're asking for here, but I'm unclear since this is mostly a thought exercise. The spectra of $\mathbb{C}[x]$ and $\mathbb{C}[y]$ can be identified with $\mathbb{C}$ (plus the mysterious (0)). Therefore looking at the map $\phi : \mathbb{C}[y] \to \mathbb{C}[x]$ where $y \mapsto x^2$ gives us the map on spectra which is the first map described above. I'm not sure what else to write down. $\square$

Exercise 3.2.P.
Suppose k is a field, and $f_1, \ldots, f_n \in k[x_1, \ldots, x_m]$ are given. Let $\phi : k[y_1, \ldots, y_n] \to k[x_1, \ldots, x_m]$ be the ring morphism defined by $y_i \mapsto f_i$.

(a) Show that ϕ induces a map $\operatorname{Spec} k[x_1, \ldots, x_m]/I \to \operatorname{Spec} k[y_1, \ldots, y_n]/J$ for any ideals $I \subset k[x_1, \ldots, x_m]$ and $J \subset k[y_1, \ldots, y_n]$ such that $\phi(J) \subset I$.

(b) Show that the map of part (a) sends the point $(a_1, \ldots, a_m) \in k^m$ (i.e. $(x_1 - a_1, \ldots, x_n - a_n) \in \operatorname{Spec} k[x_1, \ldots, x_m]$) to

$$(f_1(a_1, \ldots, a_m), \ldots, f_n(a_1, \ldots, a_m)) \in k^n.$$

Solution.

(a) This is clearly true if $I = 0$ because we can use the canonical map. Now, we know that prime ideals of $\operatorname{Spec} k[x_1, \ldots, x_m]/I$ are in bijective correspondence with prime ideals of $\operatorname{Spec} k[x_1, \ldots, x_m]$ containing I. A prime ideal $\mathfrak{p}$ containing I will be mapped to a prime ideal containing $\phi^{-1}(I)$. For $\phi^{-1}(\mathfrak{p})$ to translate back to an ideal in $k[y_1, \ldots, y_n]/J$ we need $J \subset \phi^{-1}(\mathfrak{p})$. This occurs if $\phi(J) \subset \mathfrak{p}$, which is true because $\phi(J) \subset I \subset \mathfrak{p}$. Therefore this map is well defined.

(b) If we think about maximal ideals as the vanishing set at a given point (which is the correspondence between the Spec and points of affine space), then this shouldn't be too bad. If we consider an element $g \in k[y_1, \ldots, y_n]$, then $\phi(g(y_1, \ldots, y_n)) = g(f_1(x_1, \ldots, x_m), \ldots, g_n(x_1, \ldots, x_m))$. Therefore if $\mathfrak{p} \subset k[x_1, \ldots, x_m]$ is the vanishing set of $a_1, \ldots, a_m$, we would like $\phi^{-1}(\mathfrak{p})$ to be the vanishing set of the point $(b_1, \ldots, b_n)$ so that
$$\phi(g(b_1, \ldots, b_n)) = g(f_1(a_1, \ldots, a_m), \ldots, f_n(a_1, \ldots, a_m)).$$
Therefore these b_i should correspond to $f_i(a_1, \ldots, a_m)$.

$\square$

Exercise 3.2.Q.
Consider the map of sets $\pi : \mathbb{A}^n_{\mathbb{Z}} \to \operatorname{Spec} \mathbb{Z}$ given by the ring map $\mathbb{Z} \to \mathbb{Z}[x_1, \ldots, x_n]$. If p is prime, describe a bijection between the fibre $\pi^{-1}([(p)])$ and $\mathbb{A}^n_{\mathbb{F}_p}$. Can you interpret the fibre over $[(0)]$ as $\mathbb{A}^n_k$ for some field k?

Solution.
Let ϕ be the ring map above. Then we have

$$\pi^{-1}([(p)]) = \{\mathfrak{p} \subset \mathbb{Z}[x_1, \ldots, x_n] : \phi^{-1}(\mathfrak{p}) = (p)\}.$$

What does this mean? When we take the preimage of $\mathfrak{p}$ under ϕ, all that remains are the constant polynomials in that prime ideal. There may not be any such; in this case $\pi(\mathfrak{p}) = (0)$. Otherwise, we have some minimal constant polynomial n in $\mathfrak{p}$. As such, the ideal $n\mathbb{Z}[x_1, \ldots, x_n] \subset \mathfrak{p}$. However, it is clear that n must be a prime integer: otherwise we would have $\ell \cdot m = n \in \mathfrak{p}$ for some other constant polynomials ℓ and m, nullifying primeness.

Now, we know that $\mathbb{A}^n_{\mathbb{F}_p} = \operatorname{Spec} \mathbb{F}_p[x_1, \ldots, x_n] = \operatorname{Spec} \mathbb{Z}/p\mathbb{Z}[x_1, \ldots, x_n]$. Since prime ideals of $\mathbb{Z}/p\mathbb{Z}[x_1, \ldots, x_n]$ are in bijection with ideal of $\mathbb{Z}[x_1, \ldots, x_n]$ containing $p\mathbb{Z}[x_1, \ldots, x_n]$, we see the bijection starting to emerge. An ideal $\mathfrak{p} \in \operatorname{Spec}[x_1, \ldots, x_n]$ mapping to (p) contains the constant polynomial p. Therefore $p\mathbb{Z}[x_1, \ldots, x_n] \subset \mathfrak{p}$, so $\mathfrak{p}$ correspond to an ideal in $\mathbb{Z}/p\mathbb{Z}[x_1, \ldots, x_n] = \mathbb{F}_p[x_1, \ldots, x_n]$. This gives us a point in finite affine space $\mathbb{A}^n_{\mathbb{F}_p}$.

If $\mathfrak{p}$ is an ideal without any such constant polynomial p, as we said above it goes to the zero ideal. We can probably picture this as a point in $\mathbb{A}^n_{\mathbb{Q}}$, and we might even have a bijection there. We certainly know that for every monic irreducible polynomial in $\mathbb{Q}[x_1, \ldots, x_n]$, we can associate to it an irreducible polynomial in $\mathbb{Z}[x_1, \ldots, x_n]$ by 'clearing denominators'. Indeed, these polynomials have the same roots, and the association is unique. We can use this to come up with a bijection of closed points between $\mathbb{A}^n_{\mathbb{Z}}$ lying over $[(0)]$ and $\mathbb{A}^n_{\mathbb{Q}}$.

□

Exercise 3.2.R.

(a) Show that if I is an ideal of nilpotents, then the inclusion $\operatorname{Spec} B/I \to \operatorname{Spec} B$ is a bijection.

(b) Show that the nilpotents of a ring B form an ideal, the *nilradical* which we denote $\mathfrak{N}$.

Solution.

These are pretty standard commutative algebra exercises. It suffices to show that the nilradical is the intersection of all prime ideals for both parts. Let $x \in \mathfrak{N}$, which we only assume is a subset for now, and that $x^n = 0$. Then we have for any prime ideal $\mathfrak{p}$ that $x^n = x^{n-1} \cdot x \in \mathfrak{p}$, whence either x^{n-1} or $x \in \mathfrak{p}$. If we then have $x^{n-1} \in \mathfrak{p}$, then we can write this $x^{n-2} \cdot x \in \mathfrak{p}$, whence x^{n-2} or $x \in \mathfrak{p}$. By induction this shows $x \in \mathfrak{p}$. Therefore $\mathfrak{N} \subset \mathfrak{p}$ for every prime ideal $\mathfrak{p}$, so

$$\mathfrak{N} \subset \bigcap_{\mathfrak{p} \subset B} \mathfrak{p}.$$

Now we need to show that this is everything. Suppose $x \notin \mathfrak{N}$. Then we will show that $x \notin \mathfrak{p}$ for some prime $\mathfrak{p}$. Let S be the set of ideals not containing x^n for any $n \in \mathbb{N}$. Since $(0) \in S$ this is a nonempty set which has a partial ordering on it, hence by Zorn's lemma S has a maximal element $\mathfrak{m}$.

We claim that $\mathfrak{m}$ is a prime ideal. Suppose that we have $a, b \in B$ such that $a, b \notin \mathfrak{m}$ but $ab \in \mathfrak{m}$. Then consider the ideal

$$\mathfrak{a} = \{z \in B : a \cdot z \in B\}.$$

Then $\mathfrak{m} \subsetneq \mathfrak{a}$, hence $x^n \in \mathfrak{a}$ for some n. By the same reasoning,

$$\mathfrak{b} = \{z \in B : x^n \cdot z \in B\}$$

is an ideal strictly bigger than $\mathfrak{m}$, so this ideal contains some x^m. But then this implies that $x^n \cdot x^m = x^{n+m} \in \mathfrak{m}$, which is a contradiction. Therefore $\mathfrak{m}$ is prime and hence x does not belong to some prime ideal.

Having shown this, it is clear that the inclusion above is a bijection: since primes in B/I correspond to primes in B containing I, this is all primes. Further, the intersection of ideals is again an ideal, so $\mathfrak{N}$ is an ideal. $\square$

Exercise 3.2.S.
Prove that the nilradical is the intersection of all prime ideals.

Solution.
Oops, I didn't read ahead. See above. $\square$

Exercise 3.2.T.
Suppose we have a polynomial $f(x) \in k[x]$. Instead, we work in $k[x, \varepsilon]/(\varepsilon^2)$. What then is $f(x + \varepsilon)$?

Solution.
Let $f(x) = \sum_{i=0}^{n} a_i x^i$. Then we have

$$f(x+\varepsilon) = \sum_{i=0}^{n} a_i (x+\varepsilon)^i = \sum_{i=0}^{n} a_i \left(\sum_{j=0}^{i} \binom{i}{j} x^{i-j} \varepsilon^j \right)$$

Most of those terms, however, vanish since $\varepsilon^2 = 0$. Hence

$$= \sum_{i=0}^{n} a_i (x^i + i x^{i-1} \varepsilon) = f(x) + \varepsilon \cdot f'(x).$$

This looks like it will be very convenient. $\square$

3.4 The underlying topological space of an affine scheme

Exercise 3.4.A.
Check that the x-axis is contained in $V(xy, yz)$.

Solution.
Recall that

$$V(s) := \{[\mathfrak{p}] \in \operatorname{Spec} A : S \subset \mathfrak{p}\}.$$

Since the x-axis is given by $y = z = 0$, we know that it correspond to the prime ideal (y, z). Certainly xy and yz are contained in this ideal. $\square$

Exercise 3.4.B.
Show that if (S) is the ideal generated by S, then $V(S) = V((S))$.

Solution.
This is trivial. If $S \subset \mathfrak{p}$, then $(S) \subset \mathfrak{p}$. The converse is even more trivial. $\square$

Exercise 3.4.C.

(a) Show that $\varnothing$ and $\operatorname{Spec} A$ are both open.

(b) If I_i is a collection of ideals (as i runs over some index set), show that

$$\bigcap_i V(I_i) = V(\sum_i I_i).$$

Hence the union of any collection of open sets is open.

(c) Show that $V(I_1) \cup V(I_2) = V(I_1 \cdot I_2)$. Hence the intersection of any finite number of open sets is open.

Solution.

(a) We know that $V(\varnothing) = \operatorname{Spec} A$, hence $\varnothing$ is open. Further, $V(A) = \varnothing$, since no prime ideal contains anything, therefore $\operatorname{Spec} A$ is open.

(b) We will show double inclusion. Suppose that $\mathfrak{p} \in \bigcap_i V(I_i)$. Then $I_i \subset \mathfrak{p}$ for all I_i, therefore the ideal generated by all of these I_i must also be a subset of $\mathfrak{p}$. This is precisely $\sum_i I_i$. Conversely, if $\mathfrak{p} \in V(\sum_i I_i)$, then we have $p \in V(I_i)$ for each i as well, since $I_i \subset \sum_i I_i$. Therefore $\mathfrak{p} \in \bigcap_i V(I_i)$. This completes the equality.

(c) Suppose that $\mathfrak{p} \in V(I_1) \cup V(I_2)$. Then without loss of generality, $I_1 \subset \mathfrak{p}$. As such, any linear combination $\sum_{j=1}^n a_i x_i \in \mathfrak{p}$, where $x_i \in I_1$ and $a_i \in A$. These combinations contain all A-linear combinations of elements of I_1 and I_2, so $I_1 \cdot I_2 \subset \mathfrak{p}$, hence $\mathfrak{p} \in V(I_1 \cdot I_2)$. Conversely, suppose that $\mathfrak{p} \notin V(I_1) \cup V(I_2)$. This means that $I_1 \not\subset \mathfrak{p}$ and $I_2 \not\subset \mathfrak{p}$. Let $a \in I_1 \setminus \mathfrak{p}$ and $b \in I_2 \setminus \mathfrak{p}$. Then if we had $\mathfrak{p} \in V(I_1 \cdot I_2)$, then in particular we must have $ab \in \mathfrak{p}$. But $a, b \notin \mathfrak{p}$, which contradicts that $\mathfrak{p}$ was a prime ideal. Therefore we have $V(I_1) \cup V(I_2) \subset V(I_1 \cdot I_2)$ and $(V(I_1) \cup V(I_2))^c \subset (V(I_1 \cdot I_2))^c$, which shows double inclusion.

$\square$

Exercise 3.4.D.

If $I \subset A$ is an ideal, then define its radical by

$$\sqrt{I} := \{r \in A : r^n \in I \text{ for some } n \in \mathbb{N}\}.$$

Show that $\sqrt{I}$ is an ideal. Show that $V(\sqrt{I}) = V(I)$. We say that an ideal is *radical* if $\sqrt{I} = I$. Show that $\sqrt{\sqrt{I}} = \sqrt{I}$ and that prime ideals are radical.

Solution.

Certainly since $I \subset \sqrt{I}$, we have $V(\sqrt{I}) \subset V(I)$. We need only show that if $\mathfrak{p} \in V(I)$, then $\mathfrak{p} \in V(\sqrt{I})$. Consider some $r \in \sqrt{I}$ with $r^n \in I$. Then since $r^n \in I \subset \mathfrak{p}$, we have either r^{n-1} or $r \in \mathfrak{p}$. Going with the same procedure as dealing with the nilradical above, by induction this shows that $r \in \mathfrak{p}$, whence $\sqrt{I} \subset \mathfrak{p}$. This proves the first part of this problem.

We need only to show that $\sqrt{\sqrt{I}} \subset \sqrt{I}$, since the other inclusion is obvious. Suppose $r \in \sqrt{\sqrt{I}}$. Then there exists $m \in \mathbb{N}$ so that $r^m \in \sqrt{I}$. In turn, there is an $n \in \mathbb{N}$ such that $(r^m)^n \in I$, i.e. $r^{mn} \in I$, Therefore $r \in \sqrt{I}$. This proves the second thing.

Again, we need only to show that $\sqrt{\mathfrak{p}} \subset \mathfrak{p}$. Suppose that $r \in \sqrt{\mathfrak{p}}$ and $r^m \in \mathfrak{p}$. Then going by the same exact argument by induction, we have $r \in \mathfrak{p}$. Therefore prime ideals are radical. □

Exercise 3.4.E.
If $I_1, \ldots, I_n$ are ideals of a ring A, show that

$$\sqrt{\bigcap_{i=1}^{n} I_i} = \bigcap_{i=1}^{n} \sqrt{I_i}.$$

Solution.
Suppose that r is an element of the lefthand side. Then for some $m \in \mathbb{N}$, we have $r^m \in \bigcap_{i=1}^{n} I_i$, i.e. $r^m \in I_i$ for all i. Therefore $r \in \sqrt{I_i}$ for all i, hence r is in the righthand side.

Conversely, if r is in the righthand side, then for each i there is an m_i so that $r^{m_i} \in I_i$. Let $m = \max m_i$ (which exists because we only dealing with a finite intersection). Then $r^m \in I_i$ for all i, hence r is in the lefthand side. Note that without finiteness this might not have worked. □

Exercise 3.4.F.
Show that $\sqrt{I}$ is the intersection of all prime ideals containing I.

Solution.
The hint is to use that the nilradical of A is the intersection of all prime ideals of A. Consider the ring A/I. Then the nilradical of A/I is all $r \in A/I$ such that $r^n = 0$ for some n, i.e. $r^n \in I$ for some n. The representatives for $\mathfrak{N}(A/I)$ is precisely the radical of I. Since we know that prime ideals in A/I correspond to prime ideals in A containing I, we have

$$\sqrt{I} = \mathfrak{N}(A/I) = \bigcap_{\mathfrak{p} \subset A/I} \mathfrak{p} = \bigcap_{I \subset \mathfrak{q} \subset A} \mathfrak{q}.$$

This completes the solution. Note that while we are swapping around somewhat wildly between the two rings, there is no actual ambiguity because of the previous work we've done on the spectra of these two spaces □

Exercise 3.4.G.
Describe the topological space $\mathbb{A}^1_k$.

Solution.
We know that prime ideals in $k[x]$ correspond to irreducible polynomials, since this space is a PID. We showed earlier that this space has infinitely many points no matter which field k we consider. The open sets here are complements of (finite unions of) vanishing sets of irreducible polynomials $p(x) \in k[x]$. It is pretty much the same as $\mathbb{A}^1_{\mathbb{C}}$ in terms of its description. □

Exercise 3.4.H.
By showing that closed sets pull back to closed sets, show that $\pi : \operatorname{Spec} A \to \operatorname{Spec} B$ induced from $\phi : B \to A$ is a continuous map. Interpret Spec as a contravariant functor **Rings**→**Top**.

Solution.
It suffices to check this on a basis for the closed sets of $\operatorname{Spec} B$. Let $V(I) \in \operatorname{Spec} B$ be a such a closed set. Then we have

$$\pi^{-1}(V(I)) = \{\mathfrak{p} \in \operatorname{Spec} A : \pi(\mathfrak{p}) \in V(I)\} = \{\mathfrak{p} : I \subset \phi^{-1}(\mathfrak{p})\} = V(\phi(I)).$$

Therefore closed sets pull back to closed sets, so this map is indeed continuous. The interpretation of the contravariant functor follows readily: a ring morphism $\phi : B \to A$ induces a continuous map of topological spaces $\pi : \operatorname{Spec} A \to \operatorname{Spec} B$, with the necessary axioms about the identity morphism easily seen to be satisfied. □

Exercise 3.4.I.
Suppose that $I, S \subset B$ are an ideal and multiplicative set, respectively.

(a) Show that $\operatorname{Spec} B/I$ is naturally a *closed* subset of $\operatorname{Spec} B$. If $S = \{1, f, f^2, \ldots\}$ for some $f \in B$, show that $\operatorname{Spec} S^{-1}B$ is naturally an *open* subset of $\operatorname{Spec} B$. Show that for arbitrary S, $\operatorname{Spec} S^{-1}B$ need not be open or closed.

(b) Show that the Zariski topology on these subsets is the subspace topology induced by the inclusion.

Solution.
(a) We know that primes in B/I correspond to primes in B containing I. Therefore $\operatorname{Spec} B/I = V(I) \subset \operatorname{Spec} B$, so this is a naturally closed subset. We know that primes in $S^{-1}B$ correspond to primes in B avoiding S. If we take S as in the problem statement, then $\operatorname{Spec} S^{-1}B = \operatorname{Spec} B \setminus V((S)) = \operatorname{Spec} B \setminus V((f))$. As such, $\operatorname{Spec} S^{-1}B$ is open. This should also hold whenever S or (S) satisfies some finitely-generated criterion, so that $V((S))$ is necessarily closed as it is a union of finitely many closed sets.

To use the recommended counterexample, if we take the quotient field localisation $\mathbb{Z} \to \mathbb{Q}$, then S is not finitely generated as a multiplicative set; it is generated by the set of (positive) primes in $\mathbb{Z}$. We know that $\operatorname{Spec} \mathbb{Q} = \{(0)\}$, and this point is not closed in $\operatorname{Spec} \mathbb{Z}$. If it were, then we would have $(0) = V(I)$ for some $I \subset \mathbb{Z}$. But the only choice for I is (0) itself, but $V(I) = \operatorname{Spec} \mathbb{Z}$ in this case. The issue arose here because $\operatorname{Spec} Z$ was a countable union of points, and we took for $V((S))$ a countable union of closed sets, which is not necessarily closed.

(b) Let us compared closed subsets, as was suggested. If we look at a closed set $V(\mathfrak{a}) \subset \operatorname{Spec} B/I$ (in the inherent Zariski topology), then since we can pull back $\mathfrak{a}$ to an ideal of B, we have $V(\mathfrak{a}) = V(\mathfrak{a} + I)$ in the subspace topology. Since $V(\mathfrak{a} + I)$ is closed Exercise 3.4.C, we see that the basis of the Zariski topology on $\operatorname{Spec} B/I$ comes from the basis elements of the Zariski topology on $\operatorname{Spec} B$ contained in $\operatorname{Spec} B/I$. Therefore the topologies agree.

Similarly, if we have $V(\mathfrak{a}) \subset \operatorname{Spec} S^{-1}(B)$, then we know that $\mathfrak{a}$ is still an ideal in $\operatorname{Spec} B$, so the correspondence is even clearer in this case. All we require is that $V(\mathfrak{a})$ in the localised spectrum ignores the prime ideals intersecting S, i.e. we intersect with $\operatorname{Spec} S^{-1}B$.

□

Exercise 3.4.J.
Suppose $I \subset B$ is an ideal. Show that f vanishes on $V(I)$ if and only if $f \in \sqrt{I}$.

Solution.
We know that $V(I) = V(\sqrt{I})$, and that $\sqrt{I} = \bigcap_{I \subset \mathfrak{p} \subset A} \mathfrak{p}$. Therefore if f vanishes on $V(I)$, then we have $f \in \mathfrak{p}$ for all primes $\mathfrak{p} \supset I$, so that $f \in \sqrt{I}$. Conversely, if $f \in \sqrt{I}$, then $f \in \mathfrak{p}$ for all $I \subset \mathfrak{p}$, so f vanishes on $V(I)$. □

Exercise 3.4.K.
Describe the topological space $\operatorname{Spec} k[x]_{(x)}$.

Solution.
We know that the set of points we have is just (x) and (0). We know that the point (x) is closed, so the point (0) is open. The point (0) is not closed, however, so we know that (x) is not open. Therefore the only open sets are $\operatorname{Spec} k[x]_{(x)}$, (0), and $\varnothing$. This is a 2-point space with neither the indiscrete nor the discrete topology. □

3.5 A base of the Zariski topology on $\operatorname{Spec} A$: Distinguished open sets

Exercise 3.5.A.
Show that the distinguished open sets form a base for the (Zariski) topology.

Solution.
We define

$$D(f) = \{[\mathfrak{p}] \in \operatorname{Spec} A : f \notin \mathfrak{p}\}.$$

This is equivalent (as a subset of $\operatorname{Spec} A$) to $\operatorname{Spec} A_f$, as we saw in Exercise 3.4.I. Take some $S \subset A$ a subset. Suppose that $x \in V(S)$. Then for each $f \in S$, we must have $x \notin D(f)$. Rephrasing this in terms of open sets, since we take the sets $\operatorname{Spec} A \setminus V(S) = V(S)^c$ to be our topology, we have $x \in V(S)^c$ if and only if $x \in D(f)$ for some (perhaps many) $f \in S$. That is,

$$V(S)^c = \bigcup_{f \in S} D(f).$$

Therefore the $D(f)$ are a good base for our topology. □

Exercise 3.5.B.
Suppose $f_i \in A$ for $i \in J$ some index set. Show that $\bigcup_{i \in J} D(f_i) = \operatorname{Spec} A$ if and only if $(f_i) = A$, or equivalently and very usefully, there are a_i for $i \in J$, all but finitely many 0, such that $\sum_{i \in J} a_i f_i = 1$.

Solution.
Let S be the set of all the f_i. Then we know

$$\bigcup_{i\in J} D(f_i) = V(S)^c = \operatorname{Spec} A \setminus V(S).$$

Then this union is $\operatorname{Spec} A$ if and only if $V(S) = \varnothing$. Since $V(S) = V(\sqrt{(S)})$ (to be as general as possible), this means that S must generate all of $\operatorname{Spec} A$. We know that S generates all of $\operatorname{Spec} A$ if and only if $1 \in (S)$, which is our equivalent criterion above since the span of S contains only finite linear combinations. □

Exercise 3.5.C.
Show that if $\operatorname{Spec} A$ is an infinite union of distinguished open sets $\bigcup_{j\in J} D(f_j)$, then in fact it is a union of a finite number of these, i.e. there is a finite subset $J' \subset J$ such that $\operatorname{Spec} A = \bigcup_{j\in J'} D(f_j)$.

Solution.
Because we know that $1 \in A$ is expressible as a finite linear combination of some $f_{j_1}, \ldots, f_{j_n}$ (with nonzero coefficients), then in fact the ideal generated by $J' = \{f_{j_i} : i = 1, \ldots, n\}$ is all of $\operatorname{Spec} A$. □

Exercise 3.5.D.
Show that $D(f) \cap D(g) = D(fg)$.

Solution.
Suppose $\mathfrak{p} \in D(f) \cap D(g)$. Then we have $f \notin \mathfrak{p}$ and $g \notin \mathfrak{p}$. If $fg \in \mathfrak{p}$, then this contradicts the primality of $\mathfrak{p}$, so we have $\mathfrak{p} \in D(fg)$. Conversely, suppose $fg \notin \mathfrak{p}$. Then if $f \in \mathfrak{p}$, we would need $fg \in \mathfrak{p}$ since $\mathfrak{p} \supset \mathfrak{p}A$. The same problem holds for $g \in \mathfrak{p}$. This shows a double inclusion. □

Exercise 3.5.E.
Show that $D(f) \subset D(g)$ if and only if $f^n \in (g)$ for some $n \geq 1$, if and only if g is an invertible element of A_f.

Solution.
We can show these in order. We know that we can identify $D(f)$ with $\operatorname{Spec} A \setminus V(\{f\})$. We know that $V(\{f\}) = V(I_f)$, where $I_f = \sqrt{(f)}$. We define I_g similarly. Hence suppose that $D(f) \subset D(g)$. Then this means $V(I_f) \supset V(I_g)$, which occurs when $I_f \subset I_g$. In particular, $f \in I_g$ so there exists $n \in \mathbb{N}$ so that $f^n \in (g)$.

Now supposing that $f^n \in (g)$, i.e. $f \in I_g$, consider the element $g \in A_f$. Let $f^n = ag$. Then since we are inverting $\{1, f, f^2, \ldots\}$, we have that

$$1 = \frac{f^n}{f^n} = \frac{a}{f^n} \cdot g,$$

so $g \in A_f^\times$.

Finally, if g is invertible in A_f, then let $g^{-1} = a/f^n$ for some n. Then since $g \cdot (a/f^n) = 1$, we have $g \cdot a = f^n$. Then suppose that $r \in I_f$. Then there is $m \in \mathbb{N}$ such that $r^m = bf \in (f)$ for some $b \in A$. Since $(bf)^n = b^n f^n \in (g)$, we have

$$r^{mn} = (bf)^n = ab^n g \in (g),$$

then $r \in I_g$. Hence $I_f \subset I_g$, which completes the proof. □

Exercise 3.5.F.
Show that $D(f) = \varnothing$ if and only if $f \in \mathfrak{N}$.

Solution.
If $f \in \mathfrak{N}$, then $f \in \mathfrak{p}$ for all primes $\mathfrak{p}$ by the oft-quoted theorem about the nilradical, so $D(f) = \varnothing$. Conversely, if $D(f) = \varnothing$, then $f \in \mathfrak{p}$ for all primes $\mathfrak{p}$, whence $f \in \mathfrak{N}$. □

3.6 Topological (and Noetherian) properties

Exercise 3.6.A.
If $A = A_1 \times A_2 \times \cdots \times A_n$, describe a homeomorphism $\operatorname{Spec} A_1 \sqcup \operatorname{Spec} A_2 \sqcup \cdots \sqcup \operatorname{Spec} A_n \to \operatorname{Spec} A$ for which each $\operatorname{Spec} A_i$ is mapped onto a distinguished open subset $D(f_i)$ of $\operatorname{Spec} A$. Thus

$$\operatorname{Spec} \prod_{i=1}^{n} A_i = \bigsqcup_{i=1}^{n} \operatorname{Spec} A_i.$$

Solution.
Solving this seems to imply that Spec commutes with coproducts, which is pretty nice.

We certainly could reduce to the case where $n = 2$, since the general case follows by induction on this general principal. We need to understand what the images of $\operatorname{Spec} A_1$ and $\operatorname{Spec} A_2$ are in $\operatorname{Spec} A$. Specifically, we map primes $\mathfrak{p} \in A_1$ to $\mathfrak{p} \times A_2$ and primes $\mathfrak{q} \in A_2$ to $A_1 \times \mathfrak{q}$; this is the map induced by the projection of A onto its components. Further, this map as defined is injective, so we need only show it is surjective.

Since we know that $A \to A_i$ is surjective, a prime in A maps onto a prime in A_i (for each i). Therefore the restriction to each component is a prime ideal, so it must be of the form $\mathfrak{p} \times \mathfrak{q}$ for a prime in A_1 and a prime in A_2. But the ring $A/(\mathfrak{p} \times \mathfrak{q})$ is not a domain for two proper prime ideals $\mathfrak{p}$ and $\mathfrak{q}$: we have $(1,0) \cdot (0,1) = (0,0)$. The only way to eliminate such a possibility is if $\mathfrak{p} = A_1$ or $\mathfrak{q} = A_2$, which are not prime ideals per se in A_1 and A_2 but their product is a prime in A. Hence the 'inclusion' above is actually surjective as well, so we have an isomorphism.

The hint given in the question is that $\operatorname{Spec} A_1 = D(1,0)$ and $\operatorname{Spec} A_2 = D(0,1)$. This is certainly true: for the first case, we are allowed any prime ideal in the first coordinate and allowed only A_2 in the second coordinate. This completes the solution. □

Exercise 3.6.B.

(a) Show that in an irreducible topological space, any nonempty open set is dense.

(b) If X is a topological space and Z (with the subspace topology) is an irreducible subset, then the closure $\bar{Z}$ in X is irreducible as well.

Solution.

(a) Let $U \subset X$ be a nonempty open set. Then we know that $X \setminus U$ is closed, and further we can write
$$X = \bar{U} \cup (X \setminus U).$$
Since X is irreducible, we must have either $X \setminus U = X$ or $\bar{U} = X$. Since $U \neq \varnothing$ by assumption, we must have $\bar{U} = X$, i.e. U is dense.

(b) Suppose that $\bar{Z}$ is reducible in X. Then let $\bar{Z} = Y_1 \cup Y_2$ for Y_i closed. Since these are closed subsets of a closed set in X, they are themselves closed in X. Therefore $Y_1 \cap Z$ and $Y_2 \cap Z$ are closed in the subspace topology, showing that Z is reducible, a contradiction.

□

Exercise 3.6.C.
If A is an integral domain, show that $\operatorname{Spec} A$ is irreducible.

Solution.
Suppose that Y and Z are closed sets which span $\operatorname{Spec} A$. Then we have $[(0)] \in Y$ without loss of generality. Further, we know that we can write $Y = V(I)$. Then since $[(0)] \in Y$, we must have $I \subset (0)$, i.e. $I = (0)$. Therefore $Y = V((0)) = \operatorname{Spec} A$, so we do not have a proper union of $\operatorname{Spec} A$ as closed subsets. □

Exercise 3.6.D.
Show that an irreducible topological space is connected.

Solution.
Suppose that $\operatorname{Spec} A$ was disconnected. Then there exist $Y, Z \subset \operatorname{Spec} A$ which are open and satisfy $Y \cup Z = \operatorname{Spec} A$, $Y \cap Z = \varnothing$. Since we have $Y = \operatorname{Spec} A \setminus Z$, this means that Y is closed, and similarly Z is closed. Hence $\operatorname{Spec} A = Y \cup Z$ is a union of proper closed sets, so it is reducible. This proves the contrapositive. □

Exercise 3.6.E.
Give an example of a ring A where $\operatorname{Spec} A$ is connected but reducible.

Solution.
Consider $A = \mathbb{C}[x, y]/(xy)$. Its spectrum is a subspace of $\mathbb{A}^2_{\mathbb{C}}$ whose points are the two coordinate axes (points satisfying $xy = 0$). It is certainly reducible, since it is the union of $\operatorname{Spec} \mathbb{C}[x, y]/(x = 0)$ and $\operatorname{Spec} \mathbb{C}[x, y]/(y = 0)$.

Now we will show it is connected. Because the Zariski topology is strictly coarser than the analytic topology when it comes to the closed points of $\mathbb{A}^2_{\mathbb{C}} (\cong \mathbb{C}^2)$, a disconnection here would give rise to a disconnection in $\mathbb{C}^2$. But the coordinate axes are not disconnected in $\mathbb{C}^2$, so we must be connected in $\mathbb{A}^2_{\mathbb{C}}$ as well. □

Exercise 3.6.F.

(a) Suppose $I = (wz - xy, wy - x^2, xz - y^2) \subset k[w, x, y, z]$. show that $\operatorname{Spec} k[w, x, y, z]/I$ is irreducible by showing that $k[w, x, y, z]/I$ is an integral domain.

(b) Note that the generators of the ideal of part (a) may be written as the equations ensuring that

$$\operatorname{rank}\begin{pmatrix} w & x & y \\ x & y & z \end{pmatrix} \leq 1,$$

i.e. as the determinants of the 2×2 submatrices. Generalise part (a) to the idea of rank one $2 \times n$ matrices.

Solution.

(a) Vakil recommends we show that $k[w, x, y, z]/I$ is isomorphic to a subring of $k[a, b]$ generated by monomials of degree 3. Let us attempt that. Call this subring S, and we have $S = k[a^3, a^2b, ab^2, b^3]$. Let $\phi : k[w, x, y, z] \to S$ be a map to be defined shortly, and we will show that $\ker \phi = I$. Let $\phi(w) = a^3$, $\phi(x) = a^2b$, $\phi(y) = ab^2$, and $\phi(z) = b^3$. Then we see $\phi(wz - xy) = a^3b^3 - a^2b \cdot ab^2 = 0$, $\phi(wy - x^2) = a^3 \cdot ab^2 - a^4b^2 = 0$, and $\phi(xz - y^2) = a^2b \cdot b^3 - a^2b^4 = 0$. Therefore $I \subset \ker \phi$. To show the converse, we claim that the three relations making up I are the only (linearly independent) relations between the generators of S. For instance, $wxz^3 - y^5 = 0$, but since we may substitute $y^2 = xz$, we have

$$wxz^3 - y^5 = wxz^3 - x^2yz^2 = xz^2(wz - xy),$$

which we have already counted. So supposing that this is sufficient, we have $\ker \phi = I$. Since S is a domain, we have $k[w, x, y, z]/I$ is a domain as well. This makes its spectrum irreducible.

(b) I believe it wants us to check the following situation:

$$\operatorname{rank}\begin{pmatrix} x_0 & x_1 & \cdots & x_{n-1} \\ x_1 & x_2 & \cdots & x_n \end{pmatrix} \leq 1.$$

Perhaps this should be $= 1$. Let I be the ideal generated by the determinants of the 2×2 submatrices of the above matrix. Then $k[x_0, \ldots, x_n]/I \cong S_n \subset k[a, b]$, where S_n is generated by degree n monomials. Is this enough? Maybe?

□

Exercise 3.6.G.

(a) Show that $\operatorname{Spec} A$ is quasicompact.

(b) Show that in general $\operatorname{Spec} A$ can have nonquasicompact open sets.

Solution.

(a) Note that Vakil's definition of compact is that X is quasicompact if every open cover of X has a finite subcover. Exercise 3.5.C tells us that if $\operatorname{Spec} A$ is a infinite union of distinguished open sets $D(f_i)$ for some index set $i \in I$, then there is a finite subset $S \subset I$ such that $\operatorname{Spec} A$ is the union of the $D(f_i)$ with $i \in S$.

Now consider any open cover of $\operatorname{Spec} A = \bigcup_{i \in I} U_i$. Since the distinguished open sets are a base for $\operatorname{Spec} A$, we have $U_i = \bigcup_{j \in I_i} D(f_{i,j})$. Therefore we can write

$$\operatorname{Spec} A = \bigcup_{i \in I} \bigcup_{j \in I_i} D(f_{i,j}).$$

Therefore there is a finite set $S \subset \bigcup_{i \in I} I_i$ which covers $\operatorname{Spec} A$. The elements are

$$S = \{f_{i_1,j_1}, \ldots, f_{i_n,j_n}\}.$$

If we take those original open sets corresponding to $i_1, \ldots, i_n$, then we have

$$\operatorname{Spec} A = \bigcup_{k=1}^{n} D(f_{i_k,j_k}) \subset \bigcup_{k=1}^{n} U_{i_k} \subset \operatorname{Spec} A.$$

Therefore this choice is a finite subcover.

(b) The hint is to consider the non-Noetherian polynomial ring $A = k[x_1, \ldots]$ in countably many variables, and the maximal ideal $\mathfrak{m} = (x_1, \ldots)$. Then the claim is that $V(\mathfrak{m})^c$ is not quasicompact.

Consider the open cover by the distinguished open sets $D(x_i)$ for each $i \in \mathbb{N}$. Certainly we have $x_i \in \mathfrak{m}$ for each i, so we are not including the (only) closed point we must forbid. Further, consider any point (a_i) in $\mathbb{A}_k^\infty$ (for lack of better notation). This corresponds to the maximal ideal $(x_i - a_i)$, which is in each of the $D(x_i)$ such that $a_i \neq 0$. Since always some $a_i \neq 0$ (because we are ignoring the choice $(0, 0, \ldots)$), we have that $D(x_i)$ covers $V(\mathfrak{m})^c$. We know this is sufficient because any higher-dimensional points (e.g. curves) can be described in terms of their closed points.

This cover does not have a finite subcover, however. Suppose we remove (without loss of generality) $D(x_1)$. Then using the above reasoning, we know that the point $(1, 0, \ldots)$ does not lie in any other $D(x_i)$. Therefore we must keep all $D(x_i)$ in the original description. Graphically, what is going on is that we are removing exactly one coordinate axis in each of our open sets, and thus every open set is necessary.

□

Exercise 3.6.H.

(a) If X is a topological space that is a finite union of quasicompact spaces, show that X is quasicompact.

(b) Show that every closed subset of a quasicompact topological space is quasicompact.

Solution.

(a) Let $X = \bigcup_{i=1}^n X_i$, where the X_i are quasicompact. Then consider an open cover U_j of X for $j \in J$ some index set. Then certainly $\bigcup U_j$ is an open cover for X_i, whence we get a finite subcover $U_{i_1}, \ldots, U_{i_k}$. Doing this finitely many times, we get a collection

$$\{U_{1_1}, \ldots, U_{1_k}, U_{2_1}, \ldots, U_{n_k}\}.$$

In particular, this is a finite subcover of the original U_j which covers all the X_i. Therefore it also covers their union, so X is quasicompact.

(b) Let $Z \subset X$ be a closed subset, and let X be quasicompact. Then let U_i be an open cover of $Z \subset X$ for $i \in I$ some index set. Then we have an open cover of X of the form $\bigcup U_i \cup Z^c$. Therefore there exists a finite subcover $V_0, V_1, \ldots, V_n$ for X. We know that one of these V_i must be Z^c, as none of the U_i cover any points in $Z^c \subset X$. We take this in the labelling above to be V_0. Then $V_1, \ldots, V_n$ cover $(Z^c)^c = Z$. Since these V_i must come from the original U_i, we have our finite subcover, so X is quasicompact.

□

Exercise 3.6.I.
Show that the closed points of $\operatorname{Spec} A$ correspond to the maximal ideals.

Solution.
Let $\mathfrak{m} \subset A$ be a maximal ideal, and consider the closed set $V(\mathfrak{m}) = \{\mathfrak{p} \in \operatorname{Spec} A : \mathfrak{m} \subset \mathfrak{p}\}$. Since $\mathfrak{m}$ is a maximal ideal, $\mathfrak{m} \subset \mathfrak{p}$ implies that $\mathfrak{p} = \mathfrak{m}$ or $\mathfrak{p} = A$. Since A is not in $\operatorname{Spec} A$, we know that $V(\mathfrak{m}) = \{\mathfrak{m}\}$, i.e. this is a closed point.

Conversely, suppose that $\{\mathfrak{p}\}$ is not a maximal ideal. Then $\mathfrak{p} \subset \mathfrak{m}$ for some maximal ideal $\mathfrak{m}$ by Zorn's lemma. If $\{\mathfrak{p}\}$ were closed, then we would be able to write $\{\mathfrak{p}\} = V(I)$ for some ideal I. Since $\mathfrak{p} \subset \mathfrak{m}$, we know that $\mathfrak{p} \in V(I)$ implies that $\mathfrak{m} \in V(I)$, so any closed set containing $\mathfrak{p}$ must also contain $\mathfrak{m}$. Therefore $\{\mathfrak{p}\}$ is not closed. □

Exercise 3.6.J.

(a) Suppose that k is a field, and A is a finitely generated k-algebra. Show that closed points of $\operatorname{Spec} A$ are dense by showing that if $f \in A$ and $D(f)$ is a nonempty (distinguished) open subset of $\operatorname{Spec} A$, then $D(f)$ contains a closed point of $\operatorname{Spec} A$.

(b) Show that if A is a k-algebra that is not finitely generated, the closed points need not be dense.

Solution.

(a) First, we may suppose that $f \notin \mathfrak{N}$, otherwise $D(f) = \varnothing$. In this case, A_f (the localisation of A at the set $\{1, f, f^2, \ldots\}$ is nonzero, so it is itself a finitely generated k-algebra. Further, A_f has some maximal ideal $\mathfrak{p}A_f$, where $\mathfrak{p} \in D(f)$. We claim that this $\mathfrak{p}$ is actually a maximal ideal, which would complete the proof. Consider the quotient field
$$A_f/\mathfrak{p}A_f \cong (A/\mathfrak{p})_{\bar{f}},$$
where $\bar{f}$ is the image of f in $A/\mathfrak{p}$. This is a finite field extension of k. We know that $A/\mathfrak{p} \subset (A/\mathfrak{p})_{\bar{f}}$ is a domain since $\mathfrak{p}$ is prime. As such, $A/\mathfrak{p}$ must actually be a field, whence $\mathfrak{p}$ is a maximal ideal of A.

(b) The hint is to examine $\operatorname{Spec} k[x]_{(x)}$ as in Exercise 3.4.K. This is not finitely generated as a k-algebra, because we must invert (at least) all elements of the form $(x-a)$, which is an infinite set of k is an infinite field. It consists of two points (x) and (0), where (x) is closed but (0) is not closed. (x) is therefore the only closed points, and in particular the open set (0) does not contain a closed point, whence $\{(x)\} \subset \operatorname{Spec} k[x]_{(x)}$ is not dense.

□

Exercise 3.6.K.
Suppose k is an algebraically closed field, and $A = k[x_1, \ldots, x_n]/I$ is a finitely generated k-algebra with $\mathfrak{N}(A) = \{0\}$. Consider the set $X = \operatorname{Spec} A$ as a subset of $\mathbb{A}_k^n$. The space $\mathbb{A}_k^n$ contains the 'classical' points k^n. Show that functions on X are determined by their values on closed points.

Solution.
We want to show that for different functions f and g on X, then $f - g$ is nonzero on some closed point, and therefore we can distinguish different functions. We know that on the open set $D(f-g)$, $f-g$ is nonzero by definition. Further, since $\mathfrak{N}(A) = 0$, we know that $D(f-g)$ is nonempty for any choice of $f \neq g$. By part (a) of Exercise 3.6.J, there is some closed point (of $k[x_1, \ldots, x_n]$ perhaps) $\mathfrak{p} \in D(f-g)$. Therefore $f - g$ is nonzero at $\mathfrak{p}$. But this point $\mathfrak{p}$ can be viewed in $\operatorname{Spec} A$, so this proves what we would like. □

Exercise 3.6.L.
If $X = \operatorname{Spec} A$, show that $[\mathfrak{q}]$ is a specialisation of $[\mathfrak{p}]$ if and only if $\mathfrak{p} \subset \mathfrak{q}$. Hence show that $V(\mathfrak{p}) = \overline{\{[\mathfrak{p}]\}}$.

Solution.
We know that

$$\overline{U} = \bigcap_{\substack{U \subset V \\ V \text{ closed}}} V.$$

In our case, we have that

$$\overline{\{[\mathfrak{p}]\}} = \bigcap_{I \subset \mathfrak{p}} V(I).$$

By Exercise 3.4.C, we know that $\bigcap V(I_j) = V(\sum I_j)$, so we need only to get a handle on

$$\sum_{I \subset \mathfrak{p}} I = \mathfrak{a}$$

Certainly $\mathfrak{p} \subset \mathfrak{p}$, so $\mathfrak{p} \subset \mathfrak{a}$. Conversely, since elements of $\mathfrak{a}$ are linear combinations of elements in $\mathfrak{p}$, we have $\mathfrak{a} \subset \mathfrak{p}$, so there are equal. Therefore we have proved that $\overline{\{[\mathfrak{p}]\}} = V(\mathfrak{p})$.

As such, we have $\mathfrak{p} \subset \mathfrak{q}$ if and only if $\mathfrak{q} \in V(\mathfrak{p}) = \overline{\{[\mathfrak{p}]\}}$, and we are done. This also gives an easy proof of the fact that the closed points correspond to maximal ideals. □

Exercise 3.6.M.
Verify that $[(y - x^2)] \in \mathbb{A}^2$ is a generic point for $V(y - x^2)$.

Solution.
We know that this point is generic if $\overline{\{[(y - x^2)]\}} = V(y - x^2)$. But this is totally obvious from the previous problem. □

Exercise 3.6.N.
Show that every point x of a topological space X is contained in an irreducible component of X.

Solution.
We have the description of irreducible components as maximal among the irreducible closed subsets of X. Let $\mathcal{F}$ be the family of irreducible closed subsets of X containing x. Then to apply Zorn's lemma (and thus conclude that x is in some irreducible component), we need to show that $\mathcal{F}$ is nonempty and that chains in $\mathcal{F}$ have an upper bound in $\mathcal{F}$.

We know that x is contained in $\overline{\{x\}}$, which is irreducible. If it were reducible, then there would be some proper closed $Z \subset \overline{\{x\}}$ with $x \in Z$, a contradiction. Therefore $\mathcal{F}$ is nonempty. Therefore consider a chain $Z_1 \subset Z_2 \subset \cdots$ with $Z_i \in \mathcal{F}$. Then

$$Z = \overline{\bigcup_{i=1}^{\infty} Z_i}$$

is still in $\mathcal{F}$. To see this, suppose that Z were reducible for some $A \cup B$. Then consider $A_i = A \cap Z_i$ and $B_i = B \cap Z_i$. These are closed sets such that $Z = A_i \cup B_i$, whence we have $A_i = \varnothing$ or $B_i = \varnothing$. Since this must hold for all i, we see that $A = \varnothing$ or $B = \varnothing$. Therefore Z is not actually reducible. Therefore this chain has an upper bound in $\mathcal{F}$, so Zorn's lemma tells us that x is in some irreducible component. □

Exercise 3.6.O.
Show that $\mathbb{A}^2_{\mathbb{C}}$ is a Noetherian topological space.

Solution.
We need to show that $\mathbb{A}^2_{\mathbb{C}}$ satisfies the descending chain condition on closed subsets. Recall that $V(I) \subset V(J)$ if and only if $I \subset J$. Hence given a sequence of closed subsets $Z_1 \supset Z_2 \supset \cdots$, we may rewrite this

$$V(I_1) \supset V(I_2) \supset \cdots \iff I_1 \subset I_2 \subset \cdots .$$

Therefore we ask if $\mathbb{C}[x, y]$ satisfies the ascending chain condition on ideals, i.e. is it a Noetherian ring. It is: if R is a Noetherian ring, then $R[x]$ is Noetherian. Since $\mathbb{C}$ is Noetherian (because it has only finitely many ideals), $\mathbb{C}[x]$ and hence $\mathbb{C}[x][y] = \mathbb{C}[x, y]$ are both Noetherian. (You can also do this explicitly, but why bother when you can preempt the proceeding discussion about Noetherian rings?) □

Exercise 3.6.P.
Show that every connected component of a topological space X is the union of irreducible components. Show that any subset of X that is simultaneously open and closed must be the union of the connected components of X. If X is a Noetherian topological space, show that the union of any subset of the connected components of X is always open and closed in X.

Solution.
We know that irreducible topological space is connected from Exercise 3.6.D. Therefore the connected component of any element x must contain every irreducible component it belongs to, say some Z_i. Then if we let $Z_{x,i}$ be the irreducible components containing x in some index set I_x, we have

$$U = \bigcup_{x \in U} \bigcup_{i \in I_x} Z_{x,i}.$$

Therefore U is the union of these irreducible components.

Now, suppose that U is an open and closed set, and for any $x \in U$ let C_x be the connected component of x. We have the following two subsets of C_x: $A = C_x \cap U$ and $B = C_x \setminus A$. Since C_x is closed and U is closed, A is closed, hence B is open. Further, since we can write $B = C_x \cap U^c$, and U^c is closed as well, then A is open. Therefore we have a disconnection of C_x unless $A = \varnothing$ or $B = \varnothing$. By assumption, $x \in A$ so we must have $B = \varnothing$, and hence $C_x \subset U$. Therefore we can write

$$U = \bigcup_{x \in U} C_x.$$

Now, if X is a Noetherian topological space, then we know we can write

$$X = Z_1 \cup \cdots \cup Z_n$$

for its (finitely many) irreducible components Z_i. We know that each of these Z_i are both open and closed. Since we have shown that every connected component is a union of irreducible components, then every connected component is open and closed (since these must be finite unions). Finally, every union of connected components is actually a finite union, so it is still open and closed.

□

Exercise 3.6.Q.
Show that a ring A is Noetherian if and only if every ideal of A is finitely generated.

Solution.
Suppose that A is Noetherian. Suppose $\mathfrak{a} \subset A$ is an ideal that can only be generated by infinitely many elements, say $a_1, a_2, \ldots$. Then we have an ascending chain of ideals

$$(a_1) \subsetneq (a_1, a_2) \subsetneq (a_1, a_2, a_3) \subsetneq \cdots$$

that does not stabilise. This is a contradiction.

Conversely, if every ideal A is finitely generated, consider a chain

$$\mathfrak{a}_1 \subset \mathfrak{a}_2 \subset \cdots$$

of ideals. Then $\mathfrak{a} = \sum \mathfrak{a}_i$ is also an ideal, hence is finitely generated by some elements $a_1, \ldots, a_n$. We know that $a_i \in \mathfrak{a}_{m_i}$ for some indices m_i, and so let $m = \max m_i$. Then $\mathfrak{a} = (a_1, \ldots, a_n) \subset \mathfrak{a}_m$, and since $\mathfrak{a} \supset \mathfrak{a}_m$ as well, we must have an equality. Therefore this chain stabilises at $\mathfrak{a}_m$, so A is Noetherian. □

Exercise 3.6.R.
Show that if A is Noetherian, then so is

$$A[[x]] := \varprojlim A[x]/x^n$$

the ring of power series in x.

Solution.
We take the possible hint. Let $\mathfrak{a} \subset A[[x]]$ be an ideal. Let $\mathfrak{a}_n \subset \mathfrak{a}$ be the coefficients of x^n which appear in the elements of $\mathfrak{a}$. Then we claim that $\mathfrak{a}_n$ is an ideal in A. Indeed, if $a, b \in \mathfrak{a}_n$ appear as coefficients of x^n in $\mathfrak{a}$, then they correspond to some power series $f(x)$ and $g(x)$. Then the coefficient of x^n in $(f+g)(x)$ (which is still a power series) is $a+b$, so $a+b \in \mathfrak{a}_n$. Similarly, if $r \in A$ is a scalar, then since $r \cdot f(x) \in \mathfrak{a}$, we know that the coefficient of x^n in this new power seris is $r \cdot a \in \mathfrak{a}_n$. Therefore $\mathfrak{a}_n$ is an ideal.

Now, it is clear that $\mathfrak{a}_n \subset \mathfrak{a}_{n+1}$ in the following way: since $\mathfrak{a}$ is closed under left multiplication, we know that $x\mathfrak{a} \subset \mathfrak{a}$. As such, any coefficient of x^n in $f(x)$ is also a coefficient of x^{n+1} in the power series $x \cdot f(x)$. Further, we know that the ideals $\mathfrak{a}_i$ completely determine $\mathfrak{a}$: we have that

$$\mathfrak{a} = \bigcup_{i=0}^{\infty} \mathfrak{a}_i x^i.$$

Now, since $\mathfrak{a}_0 \subset \mathfrak{a}_1 \subset \cdots$ is an ascending chain in A, it stabilises at some $\mathfrak{a}_n$. Further, each of the $\mathfrak{a}_i$ are finitely generated for some $a_{i,1}, \ldots, a_{i,m_i}$. Hence we see that $\mathfrak{a}$ is generated by

$$\{a_{0,1}, \ldots, a_{0,m_0}, a_{1,1}x, \ldots, a_{1,m_1}, \ldots, a_{i,j_i}x^i, \ldots, a_{n,m_n}x^n\}.$$

In particular, $\mathfrak{a}$ is finitely generated. Since $\mathfrak{a}$ was an arbitrary ideal, this shows that $A[[x]]$ is Noetherian. $\square$

Exercise 3.6.S.
If A is Noetherian, show that $\operatorname{Spec} A$ is a Noetherian topological space. Describe a ring A such that $\operatorname{Spec} A$ is not a Noetherian topological space.

Solution.
I think we have shown this question sufficiently. Since descending chains of closed subsets correspond to ascending chains of ideals, the stabilisation of one implies the other. The ring $\mathbb{C}[x_1, \ldots]$ will not give a Noetherian topological space because the ascending chain of ideals $(x_1) \subset (x_1, x_2) \subset \cdots$ corresponds to a nontrivial descent of closed subspaces in $\operatorname{Spec} \mathbb{C}[x_1, \ldots]$. $\square$

Exercise 3.6.T.
Show that every open subset of a Noetherian topological space is quasicompact. Hence if A is Noetherian, every open subset of $\operatorname{Spec} A$ is quasicompact.

Solution.
Let $V \subset U$ be an open subset, where U is a Noetherian topological space. Then consider any chain of closed subsets $V_1 \supset V_2 \supset \cdots$ in V. Then since V^c is closed in U, we see that $U_i := V_i \cup V^c$ is closed in U. Hence we have a descending chain of closed subsets $U_1 \supset U_2 \supset \cdots$ in U. Since U is Noetherian, this terminates at some U_n. That is, for any $N > n$, we have $U_n = U_N$, i.e. $V^c \cup V_n = V^c \cup V_N$. Therefore

$$V \cap (V^c \cup V_n) = V_n = V_N = (V^c \cup V_N) \cap V$$

for all $N > n$. Therefore the chain in V also stabilises, so V is Noetherian as well. $\square$

Exercise 3.6.U.
Show that if M is a Noetherian A-module, then any submodule of M is a finitely generated A-module.

Solution.
This is the statement equivalent to 'A is Noetherian if and only if all its ideals are finitely generated'.

Consider $N \subset M$ a submodule that is not finitely generated. Let $n_1 \in N$ be any element. Then $An_1 \subsetneq N$ is a proper ideal, so we have some $n_2 \in N$ so that $An_1 \subsetneq An_1 + An_2 \subsetneq N$ is a proper chain of submodules. We may continue this process infinitely to obtain an ascending chain of submodules in M that never terminates, else N would be finitely generated. This is contradiction. Hence no such N can exist, and we have proved what we want. □

Exercise 3.6.V.
If

$$0 \to M' \to M \to M'' \to 0$$

is exact, show that M' and M'' are Noetherian if and only if M is Noetherian.

Solution.
Suppose that M is Noetherian. Then any ascending chain of submodules in M' injects to an ascending chain of submodules in M, which must terminate. Hence M' is Noetherian. Further, any ascending chain of submodules in M'' pulls back to an ascending chain of submodules in M, which must terminate. Hence submodules and quotient modules of Noetherian modules are Noetherian.

Now, suppose that M' and M'' are Noetherian. Let $N_1 \subset N_2 \subset \cdots$ be an ascending chain in M. We will reconsider this chain in the following way. Since $M/M' \cong M''$, we know that the submodules $N_i' := N_i \cap M'$ are viewable in M' and $N_i'' := N_i/N_i'$ are viewable in M''. These chains both terminate, say at some N_k' and N_m''.

This shows the following: only finitely many of our original N_i are contained in M', and this first part of the sequence terminates. The second part of the sequence can be realised by $N_i = N_i'' + M'$, and this sequence also terminates. Therefore ascending chain N_i must stabilise somewhere. □

Exercise 3.6.W.
Show that if A is a Noetherian ring, then $A^{\oplus n}$ is a Noetherian A-module.

Solution.
By the above exercise, we have

$$0 \to A \to A \oplus A \to A \to 0$$

is a (split) exact sequence (viewed as A-modules), hence $A^{\oplus 2}$ is Noetherian. By induction we have $A^{\oplus n}$ is Noetherian for any n. □

Exercise 3.6.X.
Show that if A is a Noetherian ring and M is a finitely generated A-module, then M is a Noetherian module. Hence by Exercise 3.6.U, any submodule of a finitely generated module over a Noetherian ring is finitely generated.

Solution.
If M is finitely generated, suppose that it can be generated by n elements $m_1, \dots, m_n$. Then we have an exact sequence

$$0 \to \ker\phi \to A^{\oplus n} \xrightarrow{\phi} M \to 0,$$

where $\phi(a_1, \dots, a_n) = a_1 m_1 + \cdots + a_n m_n$. Then by the previous two exercises, $A^{\oplus n}$ is Noetherian, hence its quotient object M is Noetherian. $\square$

3.7 The function $I(-)$, taking subsets of $\operatorname{Spec} A$ to ideals of A

Exercise 3.7.A.
Let $A = k[x, y]$. If $S = \{[(x)], [(x-1, y)]\}$, then $I(S)$ consists of those polynomials vanishing on the y-axis, and at the point $(1, 0)$. Give generators for this ideal.

Solution.
We know by definition that $I(S) = (x) \cap (x-1, y)$. Since these ideals are comaximal, we have $(x) \cap (x-1, y) = (x) \cdot (x-1, y)$. Hence $I(S)$ is generated by $x(x-1)$ and xy. $\square$

Exercise 3.7.B.
Suppose $S \subset \mathbb{A}^3_{\mathbb{C}}$ is the union of the three axes. Give generators for the ideal $I(S)$.

Solution.
Since $\mathbb{A}^3_{\mathbb{C}} \cong \operatorname{Spec} \mathbb{C}[x, y, z]$, it will be easiest to prove it using that. We know that the x-axis corresponds to $y = z = 0$, so it is generated by y and z. Similarly, the y-axis is generated by x and z and the z-axis by x and y. We therefore have

$$I(S) = (y, z) \cap (x, z) \cap (x, y) = (xy, xz, yz).$$

First, we see that $(xy, xz, yz) \subset I(S)$. Second, suppose that we had any polynomial in only one variable as a generator, say $f(x)$. Then f cannot vanish on the entire x-axis by basic properties of polynomials, so this is a contradiction. All generators must be in at least two variables, and indeed what we have is sufficient. $\square$

Exercise 3.7.C.
Show that $V(I(S)) = \overline{S}$.

Solution.
Going strictly from the definitions, we know that

$$V(I(S)) = V\left(\bigcap_{\{\mathfrak{p}\} \in S} \mathfrak{p} \subset A\right) = \left\{\mathfrak{q} \in \operatorname{Spec} A : \bigcap_{\{\mathfrak{p}\} \in S} \mathfrak{p} \subset \mathfrak{q}\right\}.$$

$V(I(S))$ is closed by definition and certainly $S \subset V(I(S))$, so $\overline{S} \subset V(I(S))$. Now suppose that $a \notin \overline{S}$. Then because $\operatorname{Spec} A \setminus \overline{S}$ is open, so we can construct a function f that is nonzero at a but is zero on $\overline{S}$. Then $f \in I(S)$, so $a \notin V(I(S))$. Therefore $\overline{S} = V(I(S))$. $\square$

Exercise 3.7.D.
Prove that if $\mathfrak{a} \subset A$ is an ideal, then $I(V(\mathfrak{a})) = \sqrt{\mathfrak{a}}$.

Solution.
The huge hint is Exercise 3.4.J, which says that f vanishes on $V(\mathfrak{a})$ if and only if $f \in \sqrt{\mathfrak{a}}$. But that is all we want to show, so we've done it already. $\square$

Exercise 3.7.E.
Show that $V(-)$ and $I(-)$ give a bijection between irreducible closed subsets of $\operatorname{Spec} A$ and prime ideals of A. From this conclude that in $\operatorname{Spec} A$ there is a bijection between points of $\operatorname{Spec} A$ and irreducible closed subsets of $\operatorname{Spec} A$ (where a point determines an irreducible closed subset by taking the closure). Hence *each irreducible closed subset of* $\operatorname{Spec} A$ *has precisely one generic point.*

Solution.
The theorem preceding this exercise gives a bijection between closed subsets and radical ideals. Now, suppose that we take a prime ideal $\mathfrak{p} \subset A$. Then $V(\mathfrak{p}) = \{\mathfrak{q} \in \operatorname{Spec} A : \mathfrak{p} \subset \mathfrak{q}\}$ is irreducible because we have shown $V(\mathfrak{p}) = \overline{\{[\mathfrak{p}]\}}$. If we could write $V(\mathfrak{p}) = V(\mathfrak{a}) \cup V(\mathfrak{b})$, then we would have $[\mathfrak{p}] \in V(\mathfrak{a})$ (without loss of generality) and hence $\overline{\{[\mathfrak{p}]\}} \subset V(\mathfrak{a})$ and thus $V(\mathfrak{a}) = V(\mathfrak{p})$. Therefore $V(\mathfrak{p})$ is irreducible.

Conversely, let Z be a closed irreducible subset. Then $I(Z) = \bigcap_{[\mathfrak{p}] \in Z} \mathfrak{p} \subset A$. If this ideal were not prime, then we have $ab \in I(Z)$ such that both $a, b \notin I(Z)$. Then we claim

$$Z = (V(a) \cap Z) \cup (V(b) \cap Z),$$

which makes Z reducible, a contradiction. Indeed, we know that $V(a) \cap Z$ and $V(b) \cap Z$ are both nontrivial since $ab \in I(Z)$, but they cannot be all of Z since $a, b \notin I(Z)$. Therefore we have our contradiction.

The bijection between points of $\operatorname{Spec} A$ and irreducible closed subsets is clear. $\square$

Exercise 3.7.F.
A prime of a ring A is a *minimal prime* if it is minimal with respect to inclusion. If A is any ring, show that the irreducible components of $\operatorname{Spec} A$ are in bijection with the minimal primes of A. In particular, $\operatorname{Spec} A$ is irreducible if and only if A has only one minimal prime ideal.

Solution.
We showed earlier that a point $[\mathfrak{q}]$ is a specialisation of a point $[\mathfrak{p}]$ if and only if $\mathfrak{p} \subset \mathfrak{q}$. We know too that the sets $\overline{\{\mathfrak{p}\}}$ determine all irreducible closed *subsets* (not components) of $\operatorname{Spec} A$. Therefore given any chain of prime ideals $\mathfrak{p}_1 \subset \mathfrak{p}_2 \subset \cdots$, the irreducible *component* corresponding to all of them is determined by the bottom of the chain, i.e. the minimal prime $\mathfrak{p}_1$. In particular, given a minimal prime $\mathfrak{p}$, any prime $\mathfrak{q}$ containing it will have its corresponding irreducible closed subset contained in the irreducible component given by $\overline{\{[\mathfrak{p}]\}}$. This is the bijection we would like. $\square$

Exercise 3.7.G.
What are the minimal primes of $k[x, y]/(xy)$, where k is a field?

Solution.
The primes of $k[x,y]/(xy)$ are those primes of $k[x,y]$ containing xy. We know the primes of $k[x,y]$ are either the closed points $(x-a, y-b)$, $(f(x,y))$ for an irreducible $f \in k[x,y]$, and the generic point (0). Now, since $xy = 0$ in $k[x,y]/(xy)$, any prime ideal must contain either x or y, i.e. $(x) \subset \mathfrak{p}$ or $(y) \subset \mathfrak{p}$. Since (x) and (y) are themselves prime ideals in $k[x,y]$ that contain xy, these are prime ideals in $k[x,y]/(xy)$. There are no other candidates for minimal primes, for any nonzero $\mathfrak{p} \subset (x)$ (without loss of generality) would need to contain x or y. $y \in \mathfrak{p}$ is a contradiction and $x \in \mathfrak{p}$ implies $\mathfrak{p} = (x)$. Therefore (x) and (y) are the minimal primes. □

4 The structure sheaf, and the definition of schemes in general

4.1 The structure sheaf of an affine scheme

Exercise 4.1.A.
Show that the natural map $A_f \to \mathcal{O}_{\operatorname{Spec} A}(D(f))$ is an isomorphism.

Solution.
The given hint is that $D(f) \subset D(g)$ if and only if g is invertible in A_f, which is Exercise 3.5.E. This is great for us. Since we have defined $\mathcal{O}_{\operatorname{Spec} A}(D(f))$ to be the localisation at all $g \in A$ with $D(f) \subset D(g)$, the natural map is clearly an isomorphism. □

Exercise 4.1.B.
Make tiny changes to the proof of Theorem 4.1.2 to show base identity for any distinguished open $D(f)$.

Solution.
We see that $D(f) = \operatorname{Spec} A_f$ pretty readily. We have $D(f) = \{\mathfrak{p} \in \operatorname{Spec} A : f \notin \mathfrak{p}\}$. Primes in A_f are precisely those in A_f avoiding f. Therefore we just need to apply the above argument to the ring A_f, the main consequence being that $\operatorname{Spec} A_f$ is quasicompact (which is not true for a general open subset of $\operatorname{Spec} A$!). The same argument then applies. □

Exercise 4.1.C.
Alter the argument appropriately to show base gluability for any distinguished open set $D(f)$.

Solution.
I think this works the same was as the previous exercise. The replacement of A by A_f should make everything follow pretty reasonably. □

Exercise 4.1.D.
Suppose M is an A-module. Show that the following construction describes a sheaf $\widetilde{M}$ on the distinguished base. Define $\widetilde{M}(D(f))$ to be the localisation of M at the multiplicative set of all functions that do not vanish outside $V(f)$. Define restriction maps $\operatorname{res}_{D(f),D(g)}$ in the analogous way to $\mathcal{O}_{\operatorname{Spec} A}$. Show that this defines a sheaf on the distinguished base, and hence a sheaf on $\operatorname{Spec} A$. Show that this is an $\mathcal{O}_{\operatorname{Spec} A}$-module.

Solution.
We hope to describe $\widetilde{M}(D(f))$ in a more coherent way. Indeed, we should have a natural isomorphism of A-modules $M_f \to \widetilde{M}(D(f))$, where

$$M_f := \{m/f^n : m \in M, n \in \mathbb{N}\}.$$

This is the usual localised module. We should able to prove identity and gluability on $\widetilde{M}(\operatorname{Spec} A)$ and let the specialisation to each $D(f)$ follow as they did in the previous two exercises. Consider $\operatorname{Spec} A = \bigcup D(f_i)$ and the finite refinement $\operatorname{Spec} A = D(f_1) \cup \cdots \cup D(f_n)$ (i.e. $(f_1, \ldots, f_n) = A$). Suppose $s \in \widetilde{M}(\operatorname{Spec} A)$ such that $\operatorname{res}_{\operatorname{Spec} A, D(f_i)} s = 0$ for all f_i. Since $s = 0$ in M_{f_i} implies that there is some $N \in \mathbb{N}$ such that $f_i^N s = 0$ for all $i \in \{1, \ldots, n\}$. Because $(f_1^N, \ldots, f_n^N) = A$, we have $r_i \in A$ so that $\sum_{i=1}^n r_i f_i^N = 1$, so that

$$s = \left(\sum_{i=1}^n r_i f_i^N \right) s = \sum_{i=1}^n r_i (f_i^N s) = 0.$$

This is basically the same argument, because nothing depended on A being a ring, just an A-module. The same is true for the gluability axioms, which I don't feel the need to re-type.

Now, we need to see why this is an $\mathcal{O}_{\operatorname{Spec} A}$-module. However, this is perfectly clear: since for each distinguished open set $D(f)$ we have $\mathcal{O}_{\operatorname{Spec} A}(D(f)) \cong A_f$ and $\widetilde{M}(D(f)) \cong M_f$, we have a natural structure of M_f and as A_f-module. This defines the appropriate $\mathcal{O}_{\operatorname{Spec} A}$-module structure. □

4.3 Definition of schemes

Exercise 4.3.A.
Describe a bijection between the isomorphisms $\operatorname{Spec} A \to \operatorname{Spec} A'$ and the ring isomorphisms $A' \to A$.

Solution.
Abbreviate the structure sheaves by $\mathcal{O}_A$ and $\mathcal{O}_{A'}$. Let $\pi : \operatorname{Spec} A \to \operatorname{Spec} A'$ be an isomorphism of schemes. This comes with a homeomorphism of topological spaces (which we denote π again) and an isomorphism $\mathcal{O}_{A'} \to \pi_* \mathcal{O}_A$. In particular, we know that we can recover our rings from a global section: $\Gamma(D(1'), \mathcal{O}_{A'}) = A'$ (where $1' \in A'$ is the multiplicative identity). π gives us an isomorphism

$$\Gamma(D(1'), \mathcal{O}_{A'}) \to \Gamma(D(1'), \pi_* \mathcal{O}_A)$$

where $D(1') \subset \operatorname{Spec} A'$. But what is the righthand side? By definition,

$$\Gamma(D(1'), \pi_* \mathcal{O}_A) = \Gamma(\pi^{-1}(D(1')), \mathcal{O}_A) = \Gamma(D(1), \mathcal{O}_A).$$

Therefore π gives us an isomorphism $\pi^\sharp : A' \to A$ on these global sections. This is one direction of the correspondence.

We now need to show that the scheme isomorphism corresponding to the ring isomorphism corresponding to the scheme isomorphism is what we started with; i.e., π yields a ring isomorphism, and that ring isomorphism yields π right back. A priori, let ρ be the scheme isomorphism coming from the ring isomorphism $\pi^\sharp$ corresponding to π. We need to show that $\rho = \pi$ on points of $\operatorname{Spec} A$. Let $\mathfrak{p}$ be a prime ideal in A. Then we need to show that $\rho(\mathfrak{p}) = \pi(\mathfrak{p})$. Certainly $\rho(\mathfrak{p}) = (\pi^\sharp)^{-1}(\mathfrak{p})$. Additionally, if we consider the stalk of $\mathcal{O}_A$ corresponding to $\mathfrak{p}$, π must have yielded an isomorphism of stalks $\mathcal{O}_{A',\mathfrak{p}} \to \pi_*\mathcal{O}_{A,\mathfrak{p}} \cong \mathcal{O}_{A,\pi^{-1}(\mathfrak{p})}$. In particular, the point $\pi^{-1}(\mathfrak{p})$ (where we view π as a homeomorphism of topological spaces) should be the same point as $(\pi^\sharp)^{-1}(\mathfrak{p})$. Hopefully this is sufficient to prove $\pi = \rho$ on points.

(FINISH LATER)

□

Exercise 4.3.B.
Suppose $f \in A$. Show that under the identification of $D(f)$ in $\operatorname{Spec} A$ with $\operatorname{Spec} A_f$, there is a natural isomorphism of ringed spaces $(D(f), \mathcal{O}_{\operatorname{Spec} A}|_{D(f)}) \cong (\operatorname{Spec} A_f, \mathcal{O}_{\operatorname{Spec} A_f})$.

Solution.
The fact that $D(f) \cong \operatorname{Spec} A_f$ as topological spaces is already clear. We just need to show that we have the appropriate isomorphism of sheaves. Consider a distinguished open set $D(g) \subset \operatorname{Spec} A_f$. Then since $\operatorname{Spec} A_f$ is open in $\operatorname{Spec} A$, $D(g)$ is a distinguished open subset of $\operatorname{Spec} A$. As such, $\Gamma(D(g), \mathcal{O}_{\operatorname{Spec} A}|_{D(f)}) \cong \Gamma(D(g), \mathcal{O}_{\operatorname{Spec} A_f})$. But that's all we wanted to show. I'm not sure what further details are necessary here. □

Exercise 4.3.C.
If X is a scheme, and U is *any* open subset, prove that $(U, \mathcal{O}_{X|U})$ is also a scheme.

Solution.
We can attack this problem locally. It's enough to consider the problem on affine schemes. Further, we know that this type of restriction works for $U = D(f)$ for some distinguished open set. Since any open subset U is just a union of $D(f_i)$ for $i \in I$ some index set, we can glue together the reasoning above for $D(f)$ to make $(U, \mathcal{O}_{X|U})$ a scheme. □

Exercise 4.3.D.
Show that if X is a scheme, then the affine open sets form a base for the Zariski topology.

Solution.
Let $U \subset X$ be open. We know that every $x \in U$ has a neighbourhood U_x that is isomorphic to $\operatorname{Spec} A_x$ for some ring A_x. We know that this U_x (and its appropriate rings) is an affine open subscheme of X, and we can realise

$$U = \bigcup U_x.$$

Therefore every open set is a union of these U_x. Therefore the set of all affine open subschemes (which may be larger than this collection or all U) is also a base for the topology. □

Exercise 4.3.E.

(a) Show that the disjoint union of a *finite* number of affine schemes is also an affine scheme.

(b) Show that an infinite disjoint union of (nonempty) affine schemes is not an affine scheme.

Solution.

(a) The hint is to use Exercise 3.6.A, in which we proved that

$$\bigsqcup_{i=1}^{n} \operatorname{Spec} A_i \cong \operatorname{Spec} \prod_{i=1}^{n} A_i.$$

So indeed, the disjoint union of finitely many affine schemes is also an affine scheme.

(b) We proved in Exercise 3.6.G(a) that affine schemes are quasicompact, so it is sufficient to show that

$$\bigsqcup_{i=1}^{\infty} \operatorname{Spec} A_i = X$$

is not quasicompact. Indeed, we basically did this already. We know that $\operatorname{Spec} A_i \subset X$ is open for each i purely topologically, and thus we have an open cover of X

$$X = \bigcup_{i=1}^{\infty} \operatorname{Spec} A_i.$$

Omitting any one of these will yield something less than X, so it cannot be refined to a finite cover, whence X is not quasicompact. Therefore it cannot be an affine scheme.

□

Exercise 4.3.F.
Show that the stalk of $\mathcal{O}_{\operatorname{Spec} A}$ at the point $[\mathfrak{p}]$ is the local ring $A_{\mathfrak{p}}$.

Solution.
It suffices to build up the stalk by looking at the base of the topology. We know that $[\mathfrak{p}] \in D(f)$ if and only if $f \notin \mathfrak{p}$, and that $\mathcal{O}_{\operatorname{Spec} A}(D(f))$ is A localised at all those $g \in A$ such that $D(f) \subset D(g)$. Then we are inverting all elements that eventually are identified by $g \notin \mathfrak{p}$, which is just the local ring $A_{\mathfrak{p}}$. This can probably be made more specific by saying that we have a representative for each germ for every element not in $\mathfrak{p}$. □

Exercise 4.3.G.

(a) If f is a function on a locally ringed space X, show that the subset of X where f doesn't vanish is open.

(b) Show that if f is a function on a locally ringed space that vanishes nowhere, then f is invertible.

Solution.

(a) Taking the hint, we would like to show that if f is a function on a ringed space X, then the subset of X on which the germ of f is invertible is open. Indeed, if the germ of f is invertible at a point p, then since we know that f is actually some (f, U) for $p \in U$, we know that f is invertible on all of U, making this an open condition.

If X is now a locally ringed space, we know that the value of f at p is its residue in $\kappa(p) = \mathcal{O}_{X,p}/\mathfrak{m}_p$. f being nonzero at p is equivalent to $f \notin \mathfrak{m}_p$. In the case of a local ring, we know that $\mathcal{O}_{X,p}^\times = \mathcal{O}_{X,p} \setminus \mathfrak{m}_p$, which means f is nonvanishing at p if and only if its germ is invertible in $\mathcal{O}_{X,p}$. By the first paragraph, we are done.

(b) If f vanishes nowhere, then we know that locally at any point p that f is invertible (by the above). We have done something like this much earlier in this document as well. As such, we should be able to glue together the inverses for f in around each point p into a global inverse for f by the general properties of schemes.

□

4.4 Three examples

Exercise 4.4.A.
Show that you can glue an arbitrary collection of schemes together. That is, given a collection of schemes X_i for $i \in I$ some index set, open subschemas $X_{ij} \subset X_i$ with $X_{ii} = X_i$, isomorphisms (of schemes) $f_{ij} : X_{ij} \to X_{ji}$ with $f_i i = \mathrm{id}_{X_i}$, and the cocycle condition, show that there is a unique scheme X up to unique isomorphism with open subsets isomorphic to X_i which respects this gluing data.

Solution.
This is an absolutely miserable question to ask anyone. But if Vakil says it is essential, so be it. We know that by Exercise 2.7.D, we can perform this exact process on sheaves. In particular, we can do it on the structure sheaves of the schemes, so we probably don't have to do a whole lot of work on that end. For the topological space that goes along with this, however, we know that we want it to be

$$X = \left. \bigsqcup_{i \in I} X_i \middle/ (X_{ij} = X_{ji})_{i,j \in I} \right. .$$

This X is certainly a unique choice. We know that given our data, we can define each sheaf $\mathcal{O}_{X_i}$ on the topological space X, which is induced by the $X_i \subset X$ as an open subset. Throwing Exercise 2.7.D at this should do all the hard work for us, so we have the unique up to unique isomorphism scheme that we wanted all along. □

Exercise 4.4.B.
Show that the affine line with doubled origin is not affine.

Solution.
Let X be $\mathbb{A}^1$ with doubled origin. Then if it were affine, we would have $X \cong \operatorname{Spec} A$ for some A, and thus

$$A = \Gamma(X, \mathcal{O}_X)$$

We need to figure out what the ring of global sections is. We construct X from two copies of the affine line from $k[t]$ and $k[u]$ via the gluing $k[t, 1/t] \cong k[u, 1/u]$ where $u \leftrightarrow t$. We know that we can identify $\operatorname{Spec} k[u] \subset X$ with $D(t)$ and $\operatorname{Spec} k[t] \subset X$ with $D(u)$. Hence $X = D(u) \cup D(t)$, and so global sections on X are those on $D(u)$ and $D(t)$ that agree on $D(u) \cap D(t)$. Therefore really all we are looking at are the global sections on one of the $D(u)$, i.e.

$$\Gamma(X, \mathcal{O}_X) \cong \Gamma(D(u), \mathcal{O}_{X|D(u)}) = k[u].$$

Therefore if X were affine, it would be the spectrum of $k[u]$, which is nonsense. □

Exercise 4.4.C.
Do the same construction with $\mathbb{A}^1$ replaced by $\mathbb{A}^2$. Describe two affine open subsets of this scheme whose intersection is not an affine open subset.

Solution.
Well, should get that $X = \mathbb{A}^2$ with doubled origin can be written $\operatorname{Spec} k[x, y] \sqcup \operatorname{Spec} k[s, t]$, where we identify $k[x, y, 1/x, 1/y] \cong k[s, t, 1/s, 1/t]$ via $x \leftrightarrow s$ and $y \leftrightarrow t$. Hence we have $D(s, t) = \operatorname{Spec} k[x, y]$ and $D(x, y) = \operatorname{Spec} k[s, t]$, and as above we should be able to calculate the ring of global sections to be $k[x, y]$ as with the affine plane with no origin.

We further see that $D(x, y) \cap D(s, t)$ is isomorphic to the affine plane with no origin, so we have the intersection of affine open subsets which is not affine, since Vakil has proven that the affine plane with no origin is not affine itself. □

Exercise 4.4.D.
Check that the gluing information in the construction of $\mathbb{P}^n_k$ works on triple intersections.

Solution.
I think what we have to do is the following. We are trying to glue together open sets $U_0, \ldots, U_n$ which are made by

$$U_i := \{[x_0, x_1, \ldots, x_n] : x_i \neq 0\} \text{ viewed like } \{(x_{0/i}, \ldots, \widehat{x_{i,i}}, \ldots, x_{n/i})\},$$

where $x_{i/j} = x_i/x_j$. We know this works when $x_i \neq 0$, and really the ratios are all that matter in the projective case. We can identify $U_i = \operatorname{Spec} k[x_{0/i}, \ldots, x_{n/i}/(x_{i/i} - 1)$ if we view these as proper variables but we don't really care about $x_{i/i}$ since it will equal 1 for practical purposes. The gluing in question is gluing (somewhat augmented for technical reasons) U_i to U_j along $D(x_{j/i}) \cong D(x_{i/j})$ via $x_{k/i} = x_{k/j}/x_{i/j}$ and $x_{k/i} = x_{k/j}/x_{j/i}$.

Now we're expected to write this down for triple overlaps. First, what is the triple intersection $U_i \cap U_j \cap U_k$? It is those values such that $x_i, x_j, x_k \neq 0$. Given all the dividing we're doing, it should just looks like the spectrum of $k[x_0, \ldots, \widehat{x_i}, \widehat{x_j}, \widehat{x_k}, \ldots, x_n]$. Therefore any gluing information should do fine in the intersection as it is an affine scheme. This is what Vakil claims, but I think this construction is nice enough that writing more is counterproductive. □

Exercise 4.4.E.
Show that the only functions on $\mathbb{P}^n_k$ are constants ($\Gamma(\mathbb{P}^n_k, \mathcal{O}) \cong k$), and hence that $\mathbb{P}^n_k$ is not affine if $n > 0$.

Solution.
It looks like we're supposed to mostly talk our way through this. We know that we need to look at functions defined on each U_i that agree on all intersections. But because of what we've been doing this whole time, any rational functions we might thing to include have to exclude one variable for each of the U_i, which ultimately exhausts all the variables at our disposal. Hence $\Gamma(\mathbb{P}^n_k, \mathcal{O})$ is $k[x_0, \dots, x_n]$ but without any polynomials containing any variables, i.e. just k itself. □

Exercise 4.4.F.
Show that if k is algebraically closed, the closed points of $\mathbb{P}^n_k$ may be interpreted in the traditional way: points are of the form $[a_0, \dots, a_n]$ where not all a_i are zero and $[a_0, \dots, a_n]$ is identified with $[\lambda a_0, \dots, \lambda a_n]$ for all $\lambda \in k^\times$.

Solution.
We have glued together $n+1$ copies of affine space $\mathbb{A}^n_k$ in a particular way to make $\mathbb{P}^n_k$, so the only closed points in there should have been closed in affine space to begin with. By the Nullstellensatz, the closed points in $\mathbb{A}^n_k$ for k algebraically closed are what we expected. The gluing that we did identifies the orbits of the $k^\times$ action on this union of points. Therefore our intuition is sound. □

4.5 Projective schemes, and the Proj construction

Exercise 4.5.A.
Consider $\mathbb{P}^2_k$ with projective coordinates x_0, x_1, x_2. Think through how to define a scheme that should be interpreted as $x_0^2 + x_1^2 - x_2^2 = 0$ "in $\mathbb{P}^2_k$".

Solution.
It would be great to have this question asked more firmly, but it is what it is. The hint tells us that if we restrict to the affine scheme $U_2 \subset \mathbb{P}^2_k$, this scheme will be cut out by $x_{0/2}^2 + x_{1/2}^2 - 1 = 0$, whence it should be the scheme $\operatorname{Spec} k[x_{0/1}, x_{0/2}]/(x_{0/2}^2 + x_{1/2}^2 - 1)$. Similarly, we should have the restriction to U_1 we should get $\operatorname{Spec} k[x_{0/1}, x_{2/1}]/(x_{0/1}^2 - 1 + x_{2/1}^2)$ and for U_0 we should get $\operatorname{Spec} k[x_{1/0}, x_{2/0}]/(1 + x_{1/0}^2 - x_{2/0}^2)$. This gives us something that looks vaguely circular in two copies of affine space and something that looks hyperbolic in the third copy.

I have been advised not to deal with questions of gluing as much as I can, so I will avoid this question. □

Exercise 4.5.B.
Consider $\mathbb{P}^n_A$ with projective coordinates $x_0, \dots, x_n$. Given a collection of homogeneous polynomials $f_i \in A[x_0, \dots, x_n]$, make sense of the scheme "cut out in $\mathbb{P}^n_A$ by the f_i".

Solution.
The hint is to use the same procedure as 4.4.D, which (of course) we didn't really do. We can realise the scheme cut out by these functions by restricting to the appropriate notion in the affine spaces $U_0, \dots, U_n$ and note that we can still glue them together.

To try to make this precise, given a single function $f(x_0, \dots, x_n)$, we can look at the functions $f(x_{0/i}, \dots, x_{n/i})$ which live in $k[x_{0/i}, \dots, x_{n/i}]$. We then can take the quotient of this ring by $(x_{i/i} - 1)$ as in 4.4.9 to get us back to affine space, and now we know how to look at the vanishing set of $f(x_{0/i}, \dots, \widehat{x_{i/i}}, \dots, x_{n/i})$ in this space. We then go through the same effort of gluing everything together.

This is a pretty sketchy explanation, but it should suffice for the moment. □

Exercise 4.5.C.

(a) Show that an ideal I is homogeneous if and only if it contains the degree n piece of each of its elements for each n.

(b) Show that homogeneous ideals are closed under sum, product, intersection, and radical.

(c) Show that a homogeneous ideal $I \subset S_\bullet$ is prime if $I \neq S_\bullet$, and for any *homogeneous* $a, b \in S$, if $ab \in I$, then $a \in I$ or $b \in I$.

Solution.

(a) First, assume I is homogeneous. By definition, it is generated by homogeneous elements $\{a_i\}$ for $i \in A$ some index set. First, consider an element of the form $s \cdot a$ for one of the generators a. Assume that $a \in S_m$ and $s = \sum_{j \in \mathbb{Z}} s_j$ is the decomposition of s into homogeneous elements. Then the degree n piece of $s \cdot a$ is $s_{n-m} \cdot a$. This is just a multiple of one of the generators, so it is contained in I.

An arbitrary element $a \in I$ is of the form $a = \sum_{i=1}^{\ell} s_i \cdot a_i$, where $a_i \in S_{m_i}$ and $s_i = \sum_{j \in \mathbb{Z}} s_{i,j}$ is the decomposition into homogeneous elements. Then the degree n part of a is $\sum_{i=1}^{\ell} s_{i,n-m_i} \cdot a_i$, which is contained in I as it is again a sum of multiples of the generators. This proves the forward direction.

For the converse, suppose that I contains the degree n piece of each of its elements. Consider a generating set $\{a_i\}$ of I for $i \in A$ some index set. Then the degree n part of each of the a_i is also contained in I for each $n \in \mathbb{Z}$, so we can construct a new generating set

$$\left\{ a_{i,n} : i \in A, a_{i,n} \in S_n, a_i = \sum_{n \in \mathbb{Z}} a_{i,n} \right\}.$$

This is just taking each generator to its homogeneous decomposition. This is still a generating set for I, so I is generated by homogeneous elements, i.e. I is homogeneous.

(b) Let I, J be homogeneous ideals. Let I be generated by $\{a_i\}$ and J by $\{b_j\}$. Then $I + J$ is generated by $\{a_i, b_j\}$ all of which are homogeneous so $I + J$ is homogeneous. $I \cdot J$ is generated by $\{a_i \cdot b_j\}$, all of which are still homogeneous, so $I \cdot J$ is homogeneous.

Now, consider $I \cap J$. If $a \in I \cap J$, then the degree n part of a is contained in both I and J as well, so $a_n \in I \cap J$ (where $a = \sum_{n \in \mathbb{Z}} a_n$ in the usual way). By part (a), this means that $I \cap J$ is homogeneous. Similarly, we know that $\sqrt{I} = \{a \in S : a^m \in I \text{ for some } m \in \mathbb{N}\}$. If $a^m \in I$, then we know that the degree n part of a^m is contained in I, which we denote $(a^m)_n$. It would be great to say $(a^m)_n = (a_n)^m$, but this is patently false. However, let a_k be the highest degree piece of a. Then we know that $(a^m)_{mk} = (a_k)^m$, so $a_k \in \sqrt{I}$. Therefore $a - a_k \in \sqrt{I}$ as well. By induction, this shows that every homogeneous piece of a is also in $\sqrt{I}$.

(c) Suppose that I satisfies $ab \in I$ implies $a \in I$ or $b \in I$ for homogeneous elements. Let $a, b \in S_\bullet$ be two arbitrary elements such that $ab \in I$. Since I is homogeneous, we know that the homogeneous pieces of ab are all contained in I, which we denote $(ab)_n$. If we write $a = \sum a_n$ and $b = \sum b_n$ as the individual decompositions into homogeneous elements, we know that

$$(ab)_n = \sum_{i+j=n} a_i \cdot b_j.$$

If $a \notin I$, then there exists some component a_d of a that is not in I, and similarly if $b \notin I$ there is some $b_e \notin I$. Without loss of generality, we may take d and e to be the maximal degree satisfying this property.

Now, consider

$$(ab)_{d+e} = \sum_{i+j=d+e} a_i \cdot b_j = a_d \cdot b_e + \sum_{i \neq d, j \neq e} a_i \cdot b_j.$$

Then in the rightmost sum, we have $i > d$ or $j > e$ in each term, so each term is a multiple of an element of I. Therefore

$$(ab)_{d+e} - \sum_{i \neq d, j \neq e} a_i \cdot b_j \in I.$$

But this is precisely $a_d \cdot b_e$. However, we assumed that $a_d, b_e \notin I$, which violates the assumption on I. Hence we contradict that $a \notin I$ and $b \notin I$, so $a \in I$ or $b \in I$, so I is prime.

□

Exercise 4.5.D.
Note: all graded rings are now assumed to contain *no elements of negative degree* henceforth, i.e. they are all $\mathbb{Z}^{\geq 0}$-graded.

(a) Show that a graded ring $S_\bullet$ over A is a finitely generated graded ring if and only if $S_\bullet$ is a finitely generated graded A-algebra, i.e. generated over $A = S_0$ by a finite number of homogeneous elements of positive degree.

(b) Show that a graded ring $S_\bullet$ over A is Noetherian if and only if $A = S_0$ is Noetherian and $S_\bullet$ is a finitely generated graded ring.

Solution.

(a) If $S_\bullet$ is a finitely graded ring, then the ideal $S_+ = \bigoplus_{n>0} S_n$ is finitely generated. Let $s_1, \ldots, s_k$ be a generating set for this ideal. Without loss of generality, we can assume that all these elements are homogeneous, with $s_i \in S_{m_i}$. We know further that $m_i > 0$ for all i. The claim is that this set generates $S_\bullet$ as an A-algebra.

Let $s \in S_+$. Then we can write

$$s = \sum_{i=1}^{k} c_i \cdot s_i$$

for some elements $c_i \in S_\bullet$. For simplicity, we can assume that s is also a homogeneous element, because we can always take the original problem and break it up into homogeneous pieces for consideration. Suppose $s \in S_n$. Then we can choose each c_i to be homogeneous of degree $n - m_i$. This is because of c_i is not homogeneous, then $c_i \cdot s_i$ will have support in a degree not equal to n, which must be zeroed by the rest of the elements. If we just eliminate the parts of the c_i which are not in degree $n - m_i$ for each i, then we do not change the net result.

If $m_i > n$, then $c_i = 0$, and if $m_i = n$ then $c_i \in A$. If we have $n > m_i$, then $c_i \in S_+$, so in turn we may write

$$c_i = \sum_{j=1}^{\ell} d_i \cdot s_i.$$

We now have the same situation: each d_i is (after trivial adjustment) homogeneous of degree $(n - m_i) - m_i$. If $2m_i > n$, then $d_i = 0$, and if $2m_i = n$ then $d_i \in A$. By induction, since $m_i > 0$, we eventually express every coefficient as an A-linear combination of products of the s_i. This is contained in the A-algebraic span of the s_i, so this set s_i suffices to generate S_+ as an A-algebra. Since we can get 1_A by the empty sum of the s_i, we can see that s_i span $S_\bullet$ as an A-algebra.

The converse is trivial. If $S_\bullet$ is finitely generated as an A-algebra, then any element of $S_+ \subset S_\bullet$ may be written

$$s = \sum_{i=1}^{k} c_i \cdot s_i$$

for some generating set $s_1, \ldots, s_k$ and elements $c_i \in A$. This means that $s_1, \ldots, s_k$ generate S_+, so this ideal is finitely generated.

(b) If $S_\bullet$ is Noetherian, then S_+ is an ideal of $S_\bullet$, so is Noetherian. $A = S_0 = S_\bullet / S_+$, so it is a quotient object of a Noetherian ring, hence is Noetherian. Further, if $S_\bullet$ is Noetherian, all its ideals are finitely generated, so S_+ is finitely generated, hence $S_\bullet$ is finitely generated as a ring.

Conversely, if $S_\bullet$ is finitely generated as a ring and A is Noetherian, by part (a) we know that $S_\bullet$ is finitely generated as an A-algebra. We have shown before that a finitely generated algebra over a Noetherian ring is still Noetherian, so $S_\bullet$ is Noetherian.

□

Exercise 4.5.E.
Suppose $f \in S_+$ is homogeneous.

(a) Give a bijection between the primes of $((S_\bullet)_f)_0$ and the homogeneous prime ideals of $(S_\bullet)_f$.

(b) Interpret the set of prime ideals of $((S_\bullet)_f)_0$ as a subset of $\operatorname{Proj} S_\bullet$.

Solution.

(a) Vakil immediately tells us to prove something else, so we will. Suppose A is a $\mathbb{Z}$-graded ring with a homogeneous invertible element f in positive degree. Then we should show there is a bijection between prime ideals of A_0 and homogeneous prime ideals of A. This is actually what we want to do, but we are making our notation more polite.

Let $\phi : A_0 \to A$ be the inclusion ring homomorphism. Then we know that if $\mathfrak{p} \subset A$ is a homogeneous prime ideal, then $\phi^{-1}(\mathfrak{p})$ is a prime ideal of A_0. We now need to reverse the correspondence. We follow Vakil's advice. Let $\mathfrak{p}_0 \subset A_0$ be a prime ideal. Define the subset $\mathfrak{p} \subset A$ (which we'd like to be a homogeneous prime ideal eventually) as the direct sum of $\mathfrak{q}_i$, where

$$\mathfrak{q}_i = \{a \in A_i : a^{\deg f}/f^i \in \mathfrak{p}_0\}.$$

In particular, $\mathfrak{q}_0 = \mathfrak{p}_0$, so if we are successful, $\phi^{-1}(\mathfrak{p}) = \mathfrak{p}_0$, which would prove that we have a bijection.

We will prove properties as Vakil suggests them: suppose that $a \in \mathfrak{q}_i$. We claim that $a^2 \in \mathfrak{q}_{2i}$. Indeed, $a \in \mathfrak{q}_i$ means that $a^{\deg f}/f^i \in \mathfrak{p}_0$. Then

$$a^{\deg f}/f^i \cdot a^{\deg f}/f^i = a^{2\deg f}/f^{2i} = (a^2)^{\deg f}/f^{2i} \in \mathfrak{p}_0.$$

Conversely, if $a^2 \in \mathfrak{q}_{2i}$, then

$$\left(a^{\deg f}/f^i\right)^2 \in \mathfrak{p}_0 \implies a^{\deg f}/f^i \in \mathfrak{p}_0$$

because prime ideals are radical.

Now, given $a_1, a_2 \in \mathfrak{q}_i$, we know that

$$(a_1 + a_2)^2 = a_1^2 + 2a_1a_2 + a_2^2 \in \mathfrak{q}_{2i}.$$

The square elements we have just shown are in $\mathfrak{q}_{2i}$, and the mixed element satisfies

$$(a_1a_2)^{\deg f}/f^{2i} = a_1^{\deg f}/f^i \cdot a_2^{\deg f}/f^i \in \mathfrak{p}_0.$$

Therefore $a_1 + a_2 \in \mathfrak{q}_i$. All this shows that $\mathfrak{q}_i \subset A_i$ are ideals. As such, $\mathfrak{p} = \oplus \mathfrak{q}_i$ is a homogeneous ideal of A. We claim that it satisfies the homogeneous prime ideal criterion we have been using for a while. Suppose that $a, b \in A$ are homogeneous elements, so without loss of generality $a \in A_m$, $b \in A_n$. Further suppose that $ab \in \mathfrak{p}$. In particular, $ab \in \mathfrak{q}_{n+m}$. Hence

$$(ab)^{\deg f}/f^{m+n} = a^{\deg f}/f^m \cdot b^{\deg f}/f^n \in \mathfrak{p}_0.$$

Since $\mathfrak{p}_0$ is a prime ideal, either $a^{\deg f}/f^m$ or $b^{\deg f}/f^n$ is in $\mathfrak{p}_0$. Therefore $a \in \mathfrak{q}_m$ or $b \in \mathfrak{q}_n$, so in particular $a \in \mathfrak{p}$ or $b \in \mathfrak{p}$. This shows that $\mathfrak{p}$ is a homogeneous prime ideal, so we have our correspondence.

(b) The points of $\operatorname{Proj} S_\bullet$ are those homogeneous prime ideals of $S_\bullet$ which do not contain S_+. We know that prime ideals of $((S_\bullet)_f)_0$ correspond to homogeneous prime ideals of $(S_\bullet)_f$. We know that (homogeneous) prime ideals of $(S_\bullet)_f$ correspond to (homogeneous) prime ideals of $S_\bullet$ which do not contain f. An ideal not containing f certainly cannot contain S_+, so this is the subset we are looking for.

□

Exercise 4.5.F.
Show that $D(f)$ "is" (or more precisely, "corresponds to") the subset $\operatorname{Spec}((S_\bullet)_f)_0$ from 4.5.E(b).

Solution.
We have defined $D(f) = \operatorname{Proj} S_\bullet \setminus V(f)$, where $V(f)$ is the set of homogeneous prime ideals of $S_\bullet$ containing f but not containing S_+. Therefore $\operatorname{Proj} S_\bullet \setminus V(f)$ is the set of homogeneous prime ideals of $S_\bullet$ which do not contain f (nor S_+). By the above, this correspondence is clear. □

Exercise 4.5.G.
Verify that the projective distinguished open sets $D(f)$ (as f runs through the homogeneous elements of S_+) form a base of the Zariski topology.

Solution.
We want to show that any open set $U \subset \operatorname{Proj} S_\bullet$ is a union of the $D(f)$. We know that the complement of U is some closed set $V(I)$, where I is a homogeneous ideal contained in S_+. We know I is generated by some positive degree homogeneous elements a_i for $i \in A$ some index set. We know that

$$V(I) = \bigcap_{i\in A} V(a_i),$$

because any ideal containing all the a_i will contain I and vice versa. Hence

$$U = \operatorname{Proj} S_\bullet \setminus V(I) = \operatorname{Proj} S_\bullet \setminus \bigcap_{i\in A} V(a_i) = \bigcup_{i\in A} \operatorname{Proj} S_\bullet \setminus V(a_i) = \bigcup_{i\in A} D(a_i).$$

This completes the proof. □

Exercise 4.5.H.
Fix a graded ring $S_\bullet$.

(a) Suppose I is any homogeneous ideal of $S_\bullet$ contained in S_+, and f is a homogeneous element of positive degree. Show that f vanishes on $V(I)$ (i.e. $V(I) \subset V(f)$) if and only if $f^n \in I$ for some n.

(b) If $Z \subset \operatorname{Proj} S_\bullet$, define $I(Z) \subset S_+$. Show that it is a homogeneous ideal of $S_\bullet$. For any two subsets, show that $I(Z_1 \cup Z_2) = I(Z_1) \cap I(Z_2)$.

(c) For any subset $Z \subset \operatorname{Proj} S_\bullet$, show that $V(I(Z)) = \overline{Z}$.

Solution.

(a) We have done this in the affine case, so we just need to bring it over to our new situation. In the affine case, which was Exercise 3.4.J, we showed that f vanishes on $V(I)$ if and only if $f \in \sqrt{I}$. In our case, we know that $\sqrt{I}$ is still the intersection of all prime ideals containing I (since that result was general). Therefore if $\mathfrak{p}$ is an appropriate prime ideal containing I, it will also contain $\sqrt{I}$, i.e. $V(I) = V(\sqrt{I})$. Since $V(-)$ still reverses inclusion, $V(\sqrt{I}) = V(I) \subset V(f)$ if and only if $\{f\} \subset \sqrt{I}$.

(b) Analogous to the affine case, we should have $I(Z) \subset S_+$ defined by

$$I(Z) = \bigcap_{[\mathfrak{p}] \in Z} \mathfrak{p}.$$

That is, it should be the intersection of homogeneous prime ideals not containing S_+ on which Z vanishes. We know that homogeneity is preserved under intersection, and the proof above (Exercise 4.5.C) works for arbitrary intersection, so this ideal is homogeneous. Now,

$$I(Z_1 \cup Z_2) = \bigcap_{[\mathfrak{p}] \in Z_1 \cup Z_2} \mathfrak{p} = \left(\bigcap_{[\mathfrak{p}] \in Z_1} \mathfrak{p} \right) \cap \left(\bigcap_{\mathfrak{q} \in Z_2} \mathfrak{q} \right) = I(Z_1) \cap I(Z_2).$$

This is all we wanted to show.

(c) We have shown this in the affine case (Exercise 3.7.C). We know that $V(I(Z))$ is a closed set which contains Z and so contains $\overline{Z}$. Further, if $a \notin \overline{Z}$, then because $\operatorname{Proj} S_\bullet \setminus Z$ is open, we can find some function f which is nonzero at a but zero on $\overline{Z}$, so $f \in I(Z)$ and hence $a \notin V(I(Z))$. This gives the required double inclusion, so $\overline{Z} = V(I(Z))$.

This reasoning should still work even though we don't know we're working with a scheme yet.

□

Exercise 4.5.I.
Fix a graded ring $S_\bullet$ and a homogeneous ideal I. Show that the following are equivalent.

(a) $V(I) = \varnothing$.

(b) For any f_i (as i runs through some index set) generating I, $\bigcup D(f_i) = \operatorname{Proj} S_\bullet$.

(c) $\sqrt{I} \supset S_+$.

Solution.
We have already shown $V(I) = V(\sqrt{I})$. If $V(I) = \varnothing$, then there is no prime ideal which contains $\sqrt{I}$ but does not contain S_+. This is obviously the case if $\sqrt{I} \supset S_+$, which shows (c) implies (a).

If we assume (a), then we have shown above that if f_i generate I, then

$$\bigcup D(f_i) = \operatorname{Spec} S_\bullet \setminus V(I).$$

Since $V(I) = \varnothing$, we have (b).

Finally, if (b) holds, then we know $V(I) = 0$. This shows that any prime ideal containing I contains S_+. Therefore we have

$$\sqrt{I} = \bigcap_{I \subset \mathfrak{p}} \mathfrak{p} \supset S_+.$$

This completes the proof. □

Exercise 4.5.J.
Suppose some homogeneous $f \in S_+$ is given. Via the inclusion

$$D(f) = \operatorname{Spec}((S_\bullet)_f)_0 \hookrightarrow \operatorname{Proj} S_+$$

of Exercise 4.5.F, show that the Zariski topology on $\operatorname{Proj} S_\bullet$ restricts to the Zariski topology on $\operatorname{Spec}((S_\bullet)_f)_0$.

Solution.
It is easy to check that they have the same distinguished open sets. Consider $g \in ((S_\bullet)_f)_0$ and the open set $D'(g)$ in $\operatorname{Spec}((S_\bullet)_f)_0$. We want to show this is compatible with $D(g)$. $D'(g)$ will be the set of prime ideals not containing g in $((S_\bullet)_f)_0$, which corresponds to the set of homogeneous prime ideals of $(S_\bullet)_f$ not containing g. But this is just $D(g) \cap D(f)$, so we have the appropriate restriction. □

Exercise 4.5.K.
If $f, g \in S_+$ are homogeneous and nonzero, describe an isomorphism between $\operatorname{Spec}((S_\bullet)_{fg})_0$ and the distinguished open subset $D(g^{\deg f}/f^{\deg g})$ of $\operatorname{Spec}((S_\bullet)_f)_0$.

Solution.
We know that $\operatorname{Spec}((S_\bullet)_{fg})_0$ is the set of homogeneous prime ideals of $S_\bullet$ not containing fg. Suppose $\mathfrak{p}$ is a homogeneous prime ideal not containing f. Then $\mathfrak{p}$ will also not contain fg if and only if it does not contain g by the properties of primes. To describe "$D(g)$" in $\operatorname{Spec}((S_\bullet)_f)_0$ requires changing the degree of g to zero, i.e. by examining instead $g^{\deg f}/f^{\deg g}$ so that $\deg(g^{\deg f}/f^{\deg g}) = \deg g \cdot \deg f - \deg f \cdot \deg g = 0$. This at least gives us a bijection on points. It should extend to an isomorphism of topological spaces trivially. □

Exercise 4.5.L.
By checking that these gluings behave well on triple overlaps, finish the definition of the scheme $\operatorname{Proj} S_\bullet$.

Solution.
Absolutely not. See Exercise 2.7.D if you'd like to torture yourself. □

Exercise 4.5.M.
(Re)interpret the structure sheaf of $\operatorname{Proj} S_\bullet$ in terms of compatible stalks.

Solution.
(FINISH LATER) □

Exercise 4.5.N.
Check that the definition $\mathbb{P}^n_A := \operatorname{Proj} A[x_0, \dots, x_n]$ agrees with our earlier construction.

Solution.
It should agree fine. The distinguished open sets we will glue together are $D(x_i)$, which are just homogeneous prime ideals that avoid the various x_i. They glue together here just as they glued together before. The question Vakil gives is that how do we know this covers everything? By construction before, we had homogenised our issue, so we know we need only look at the homogeneous prime ideals. □

Exercise 4.5.O.
Suppose that k is an algebraically closed field. We know from Exercise 4.4.F that the closed points of $\mathbb{P}_k^n$ are in bijection with the points of "classical" projective space. By Exercise 4.5.N, the scheme $\mathbb{P}_k^n$ is isomorphic to $\operatorname{Proj} k[x_0, \ldots, x_n]$. Therefore, each point $[a_0, \ldots, a_n]$ of classical projective space corresponds to a homogeneous prime ideal of $k[x_0, \ldots, x_n]$. Which homogeneous prime ideal is it?

Solution.
In the affine case, we had the correspondence

$$(a_1, \ldots, a_n) \leftrightarrow (x_1 - a_1), \ldots, (x_n - a_n).$$

Of course these ideals are not homogeneous. Instead, if we assume we have chosen $a_0 \neq 0$, we can associate

$$(a_0, \ldots, a_n) \leftrightarrow \left(x_1 - \frac{a_1}{a_0}x_0\right), \ldots, \left(x_n - \frac{a_n}{a_0}x_0\right).$$

These are homogeneous prime ideals and should give us what we want. □

Exercise 4.5.P.
If $S_\bullet$ is generated in degree 1, and if $f \in S_+$ is homogeneous, explain how to define $V(f)$ "in" $\operatorname{Proj} S_\bullet$, the *vanishing scheme of f*. Hence define $V(I)$ for any homogeneous ideal I of S_+.

Solution.
I have no idea what Vakil wants here. It should be the set of points on which f vanishes. If you imagine this with classical points, $V(f)$ should look like all those classical points such that $f(a_0, \ldots, a_n) = 0$, albeit with more information.

To define $V(I)$, take a generating set f_i for I and then take unions. □

Exercise 4.5.Q.
Suppose k is algebraically closed. Let V be a $(n+1)$-dimensional k vector space. Describe a natural bijection between one-dimensional subspaces of V and the closed points of $\mathbb{P}V$.

Solution.
We define $\mathbb{P}V$ to be $\operatorname{Proj}(\operatorname{Sym}^\bullet V^\vee)$, where $\operatorname{Sym}^\bullet V^\vee$ is the symmetric algebra of $V^\vee$. We are told that $\operatorname{Sym}^\bullet V^\vee = k[x_0, \ldots, x_n]$ (if we view $x_0, \ldots, x_n$ as a dual basis to V) which is convenient for us.

We know that the closed points of $\mathbb{P}V$ will then be defined by homogeneous coordinates $[a_0, \ldots, a_n]$. We can view this as a unique basis vector for a one-dimensional subspace of V, since no two different closed points will be scalar multiples. This is our bijection. □

5 Some properties of schemes

5.1 Topological properties

Exercise 5.1.A.
Show that $\mathbb{P}_k^n$ is irreducible.

Solution.
Suppose $\mathbb{P}_k^n$ is reducible, so we could write $\mathbb{P}_k^n = Y \cup Z$. Then we know $Y = V(I)$ and $Z = V(J)$ for two homogeneous ideals I and J of elements of positive degree. Then since $[(0)] \in Y$ (without loss of generality), we know that $I \subset (0)$ so $I = (0)$. But this means that $Y = V((0)) = \mathbb{P}_k^n$, which is a contradiction. □

Exercise 5.1.B.
Show that there is a bijection between irreducible closed subsets and points for general schemes.

Solution.
We showed this in Exercise 3.7.E for affine schemes. Let $X = \bigcup U_i$, where the structure sheaf restricted to each U_i is isomorphic to $\operatorname{Spec} A_i$ for some ring A_i. Since each point of X lies in some U_i, we know that to each point $p \in X$ we can associate and irreducible closed subset of U_i by taking the closure of p in U_i. Let Z_i be this set.

LATER LATER LATER.

□

Exercise 5.1.C.
We need to show that X satisfies the descending chain property on closed subsets. Suppose that $Z_1 \supset Z_2 \supset \cdots$ is a descending chain. Then we can write

$$Z_j = \bigcup_{i=1}^{n} Z_j \cap \operatorname{Spec} A_i.$$

Denote $Z_j \cap \operatorname{Spec} A_i = Y_{i,j}$. Then $Y_{i,j}$ for $j = 1, 2, \ldots$ is a descending chain of (relatively) closed subsets of $\operatorname{Spec} A_i$, so it stabilises at some Y_{i,m_j}. Let $m = \max\{m_1, \ldots, m_n\}$. Then we know that $Y_{i,\ell} = Y_{i,m}$ for all $\ell > m$ and for all i. Hence

$$Z_\ell = \bigcup_{i=1}^{n} Y_{i,\ell} = \bigcup_{i=1}^{n} Y_{i,m} = Z_m$$

for all $m > \ell$, so our original sequence stabilises. Hence X is Noetherian.

Exercise 5.1.D.
Show that a scheme X is quasi compact if and only if it can be written as a finite union of affine open subschemes.

Solution.
We know that any scheme X can be written $X = \bigcup U_i$, where $U_i \cong \operatorname{Spec} A_i$ are affine open subsets. If X is quasicompact, then this open cover admits a finite subcover, hence X is covered by finitely many affine open subschemes.

Conversely, an affine scheme is quasicompact. Therefore if X is covered by finitely many affine open subschemes $U_1, \ldots, U_n$, consider an open cover V_i of all of X. Then the set $V_i \cap U_j$ is an open cover of each U_j, hence admits a finite subcover. Therefore from our initial cover V_i, take only those V_i appearing as $V_i \cap U_j$ in one of these finite subcovers. This is a finite collection of open sets which must cover X. Hence X is quasicompact. □

Exercise 5.1.E.
Show that if X is a quasicompact scheme, then every point has a closed point in its closure. Show that every nonempty closed subset of X contains a closed point of X. In particular, every nonempty quasicompact scheme has a closed point.

Solution.
From Exercise 5.1.B, we know that there is a bijection between irreducible closed subsets and points, in the sense that a point determines an irreducible closed subset by taking the closure. Therefore we should show that every irreducible closed subset contains a closed point.

If X is quasicompact, then it is a finite union of open affine subschemes. Let Z be an irreducible closed subset of X.

□