#1: The common strategy was to prove (a) --> (b) --> (c) --> (a). The last implication is straight-forward.

For (a) --> (b), the key observation is that given a map f: S^1 --> X, a nullhomotopy H:S^1 x I --> X descends to a map F: D^2 --> X, since we can realize D^2 as the quotient of S^1 x I where we identify S^1 x {1} to a point.   The map is continuous because H(S^1 x I) = x \in X, where x is the point to which we're contracting f(S^1). Moreover, since F is just a quotient of H, it we have F(S^1 x {0}) = H(S^1 x {0}) = f(S^1), so F is an extension of f.

For (b) --> (c), if [f] in \pi_1(X), then any representative is a loop f:I --> X based at a point x \in X. Quotienting the endpoints of I, f descends to a map f:S^1 --> X, which can thus be extended to a map F:D^2 --> X. If we take i:S^1 --> D^2 to be the standard inclusion into to boundary of D^2, then f = F \circ i. Since the induced map F_*:\pi_1(D^2) --> \pi_1(X) is trivial and f_* = F_* \circ i_*, it follows that f_* sends the generator of g of \pi_1(S^1) to the identity. But f_*(g) = [f], so [f] is also trivial, as required.

For the "deduce" part, observe that if any two maps f,g:S^1 --> X are homotopic, then X must be path connected (and X is path connected if it's simply connected).   Moreover, any loop is thus homotopic to the constant map at its basepoint, so X is simply connected. On the other hand, if X is simply connected, then the concatenation \bar{h}\bar{g}*h*f is nullhomotopic, where h is a path between the basepoints of f and g, \bar{h} and \bar{g} are h,g traversed backwards (resp.), and the * denotes concatenation. But this is gives a homotopy between g and f.

#2: The key observation here is that the space in question is homotopic to the quotient of the 2-sphere S^2 by identifying 5 distinct points. Here's how this works:

- First, observe that the cube is homeomorphic to the sphere by taking the map x--> x/|x|, where |-| denotes the norm on R^3. Since the axes pass through the cube at the intersection points between this unit cube and the unit sphere, we can take the identity on the axes to get a homeomorphism between the cube unioned with the axes and the unit sphere union the axes.
- Now retract the portions of the coordinate axes outside of the sphere to points on these intersection points. This just leaves a graph attached to the inside of the sphere.
- Next, homotope the portions of the axes inside of the sphere into a small neighborhood of the unit sphere (i.e., away from the origin). Once we've done this, we can apply the reflection again x-->x/|x|, which will homeomorphically send the deformed partial axes inside of the sphere to the outside of the sphere.
- Finally, contract one of the edges of this graph to one of the attaching points on the sphere.
- Now we're in the same situation as HW 6 #11, so we get the homology calculations from there. The fundamental group is easily calculated using those same ideas and the Seifer-van Kampen theorem.

I borrowed these pictures from Raymond's test (between the 3rd and 4th pictures, one contracts the red arc):

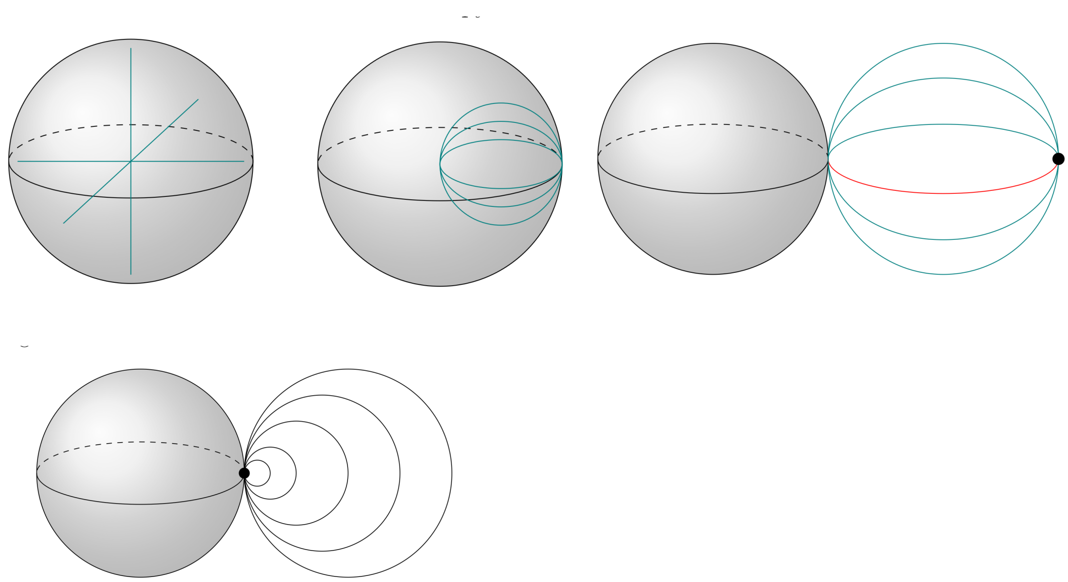

The strategy of these pictures is slightly different than what I described above, but hopefully you get the point.

#3: The idea here was to use cellular homology (which is isomorphic to singular homology) and the cellular boundary formula.

$H_i(X)$ is trivial for $i>2$ because X is a 2-dimensional CW complex. And $H_0(X)$ is Z because X is path connected. To compute $H_1(X)$ and $H_2(X)$, we'll need to understand the cellular boundary maps $d_2: H_2(X^2, X^1) \to H_1(X^1, X^0)$ and $d_1: H_1(X^1, X_0) \to H_0(X^0)$, where X^i denotes the i^th skeleton of X.

$H_2(X^2, X^1) \sim Z^2$ is generated by the two disks D_p, D_q we attached to S^1, $H_1(X^1, X^0) \sim Z$ is generated by the single 1-cell S (i.e., the copy of S^1), and $H_0(X^0) \sim Z$ is generated by the single vertex v that anchors everything.

Note that the map d_3 is trivial, so im $d_3 = 0$, so to compute $H_2(X)$ it suffices to compute ker d_2. For this, we'll use the cellular boundary formula (p. 140 in Hatcher). Since X has only one 1-cell, there are no other 1-cells to collapse, so the formula tells us that $d_2(D_p)$ is entirely determined by the attaching map for D_p, i.e. $d_2(D_p) = p S$. Similarly, $d_2(D_q) = q S$.

Since D_p and D_q generate $H_2(X^2, X^1)$, any element of $\ker(d_2)$ takes the form $aD_p + bD_q$, so that $0 = apS + bqS = (ap + bq)S$, and thus $ap + bq = 0$. Since p and q are coprime, this means that there is some c so that $a = cq$ and $b = -cp$. That is, $\ker(d_2)$ is generated by $qD_p + pD_q$, and thus is isomorphic to Z. Hence $H_2(X) \sim Z$.

Now we compute $H_1(X)$ by computing $\mathrm{im}(d_2)$ and $\ker(d_1)$. First observe that $d_1(S) = 0$, so $\ker(d_1) = H_1(X^1, X^0)$. It thus remains to understand $\mathrm{im}(d_2)$. Once again, since p and q are coprime, there exist integers x,y so that $xp + yq = 1$.

By our above calculation of d_2, we get $d_2(xD_p + yD_q) = xpS + yqS = S$. Hence $\mathrm{im}(d_2)$ is all of $H^1(X^1, X^0)$.

Thus $H_1(X)$ is trivial.

Remark: Note that if we only glued one of these disks onto S^1---say D_p---then $H_2(X) \sim Z$ and $H_1(X) \sim Z/Z_p$. By gluing on both disks, we've given a way for any curve to unwind itself into the interior of one of the disks, and thus be homologically trivial.

#4: For this problem, there were two key observations:

- Each filled-in surface R deformation retracts onto a wedge of circles. Here are pictures for genus 3:

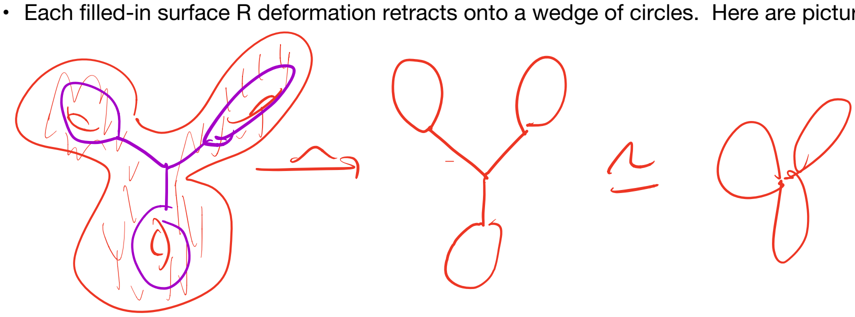

  - One can always homotope a genus g surface to have this rotational symmetry about a central point (as in the picture). Then one takes the set of loops around each genus and connects them via a star graph in the interior of R, giving you a new graph G. If one homotopes R to be nice and symmetric and chooses these loops and the star graph to be in the "central core" of R, then any cross-section transverse to G (except at the points where the circles attach to the star graph) will be a disk. The deformation retract is just radially collapsing these disks, though one has to be careful at the attaching points. A neighborhood of any of these attaching points is a thickened tripod, and hopefully you can convince yourself that this deformation retracts onto the tripod, in such a way that these parts glue up nicely to give a deformation retract of R onto G.

- The second key point what's happening in $\pi_1(R)$ vs. $\pi_1(M_g)$.
  - The first thing to observe is that any loop in R can be homotoped to a loop in its boundary M_g via the above deformation retraction. That is, one homotopes the loop into the graph G via the def. ret., and then pushes the graph G onto M_g in some way (e.g., up). Thus it suffices to understand what happens to $\pi_1(M_g)$ when we include M_g into R.
  - Let $\pi_1(M_g)$ be generated by $a_1,\dots, a_g$, and $b_1,\dots,b_g$, where the a_i's are loops that go through the genera, and the b_j are loops which go around the genera.
  - By filling in M_g to get R, notice that this allows us to fill in each of the a_i with a disk, and hence by #1, each of the a_i

    becomes nullhomotopic in R.
  - On the other hand, this is not possible for the b_j. Moreover, we see that the b_j generate the fundamental group of our core graph G, which is isomorphic to the fundamental group of R. Hence \pi_1(R) is a nonabelian free group on g generators, F_g, where the generators can be taken to be the b_j.

OK, so with this analysis in hand, we want to use Mayer-Vietoris. You can convince yourself that each copy of R in W admits a nice regular neighborhood which deformation retracts onto that copy of R, and so that the two neighborhoods intersect in an open neighborhood of the identified copies of M_g.

Using the MV sequence, we see that H_i(W) is trivial if i>3. Moreover, W is clearly path connected, so H_0(W) ~ Z.

For H_1(W), I think it's easiest to use Seifer-van Kampen and then abelianize. In particular, using SvK and the above open cover, we see that we've glued two copies of R along M_g using the identity. So in \pi_1(W) we have two copies of i_*(\pi_1(M_g)), where i: M_g --> R is the inclusion. By our above calculation, each of these is \pi_1(R) ~ F_g, and so we're identifying those two copies of F_g to a single copy.

  - In terms of curves, each copy of R has the (nontrivial) loops b_j. When we glue the two copies of R, we're identifying the two different copies of b_j into one curve, for each j.

Hence \pi_1(W) ~ F_g, and thus H_1(W) ~ Z^g.

For H_2(W), we need need to look at the MV sequence itself. For this part, we have

H_2(R) + H_2(R) --> H_2(W) --> H_1(M_g) --> H_1(R) + H_1(R) --> H_1(W)

Since H_2(R) is trivial (R is homotopic to our graph G!), this means that H_2(W) --> H_1(M_g) is injective. Since the sequence is exact, it suffices to compute the kernel of H_1(M_g) --> H_1(R) + H_1(R).

- For this, recall that the map H_1(M_g) --> H_1(R) + H_1(R) sends a curve [g] to ([g],-[g]), where the first [g] is the homology class of g in M_g, and the second [g]'s denote the homology class of g in R. Since H_1(M_g) is generated by the [a_i] and [b_j], the a_i become nullhomotopic in R, and ([b_j], -[b_j]) is nontrivial for each j, the kernel of the map is precisely the subgroup generated by the [a_i].
- Thus by exactness of the MV sequence, H_2(W) ~ Z^g.

Remark: This sort of construction is important in the theory of 3-manifolds. R is called a **handlebody**, and the gluing of the copies of R to get W is called a **Heegard splitting** of W.

  - Usually, instead of gluing these handlebodies R along their boundaries M_g by the identity, one takes a homeomorphism f: M_g --> M_g and glues using that. Since changing this homeomorphism f up to homotopy doesn't change the homotopy class of the resulting manifold W, each Heegard splitting actually determines an element of the **mapping class group**, i.e. the group of homeomorphisms of M_g up to homotopy.
  - Notice that gluing along by some other homeomorphism doesn't change the homology groups of W, but it does change \pi_1(W). When using the identity, we used SVK to compute \pi_1(W) as the quotient of the free product of two copies of \pi_1(R) amalgamated along the image of i_*(\pi_1(M_g)) in each copy. This left us with just one copy of \pi_1(R).
  - However, when we change the homeomorphism from id to f, we are now identifying i_*(\pi_1(M_g)) with i_*(f_*\pi_1(M_g)). The two resulting fundamental groups have isomorphic abelianizations, but in the latter case everything has been twisted by f. In particular, the curves a_i get identified with f(a_i), etc.

The mapping class group MCG is arguably the most important group in mathematics, and it is the main object of my research. MCG has many faces---it's the outer automorphism group of the fundamental group of the surface; it's the isometry group of all hyperbolic structures on a surface (same for singular flat structures, and conformal structures, and...); its subgroups parametrize all fiber bundles with surface fiber (up to bundle isomorphism); it's the (orbifold) fundamental group of the moduli space of curves; every closed hyperbolic 3-manifold admits a finite cover which is the mapping torus of some element of MCG; it has deep connections to number theory and algebraic geometry, etc.