Homotopy Excision Theorem. Let $(X;A,B)$ be an excisive triad and let $* \in C = A \cap B$. If (A,C) is $m-1$ connected and (B,C) is $n-1$ connected for $m \geq 2$ and $n \geq 1$, then the inclusion $(A,C) \hookrightarrow (X,B)$ induces an $n+m-2$ equivalence.

Our proof will follow the one given by J. P. May in *A Concise Course in Algebraic Topology.*

Proof. Define the q^{th} homotopy group of the triad $(X;A,B)$ as classes of maps

$$(I^q; I^{q-2} \times \{1\} \times I, I^{q-1} \times \{1\}, J^{q-2} \times I \cup I^{q-1} \times \{0\}) \to (X;A,B,*).$$

Here $J^{q-2} = \partial I^{q-2} \times I \cup I^{q-2} \times \{0\}$, so the last member of the tedrad constitutes the remaining faces of I^q. The following is then an exact sequence.

$$\cdots \to \pi_{q+1}(X;A,B) \to \pi_q(A,C) \to \pi_q(X,B) \to \pi_q(X;A,B) \to \cdots$$

Exactness at (A,C) can be seen in the following picture.

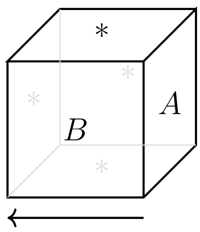

The right face coming from the cube is equivalent to the existence of a homotopy in (X,B) from the right face to the constant map on the left face.

The two pictures below show the two inclusions for exactness at (X,B).

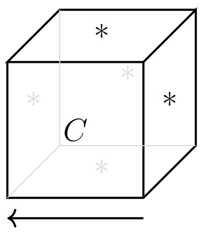

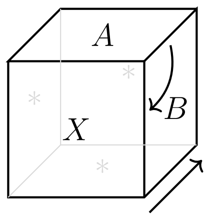

Viewing the first map from (A,C) as a map in $(X;A,B)$ (since $C \subset B$ on the front face and $* \in A$ on the right face), the null-homotopy is given by the arrow (note the interior of the cube maps to A). In the second picture, assuming the existence of the null-homotopy represented by the straight arrow gives the cube. The top face is then the preimage of the front face, and they are homotopic by the curved arrow.

Finally, the nontrivial inclusion to show exactness at $(X;A,B)$ can be seen in the following picture.

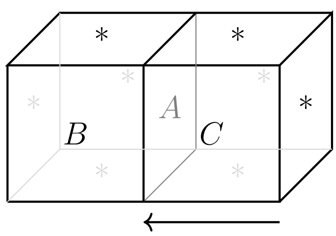

The assumption that the face labeled A is null-homotopic in (A, C) gives the right cube. Since $C \subset B$ the entire prism maps into (X, B). It is homotopic to the left cube by the arrow.

May notes in his proof that the sequence can be viewed as the long exact sequence of the pair $P(A,C) \subset P(X,B) = \{\gamma : I \to X : \gamma(0) = *, \gamma(1) \in B\}$ given the compact-open topology with the constant map as a base point. If one is comfortable with the picture below, then it can replace the arguments of all four pictures above.

$\pi_q(X; A, B) \ni$ 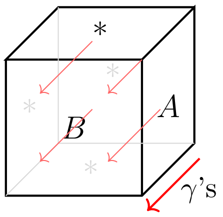

$=$ 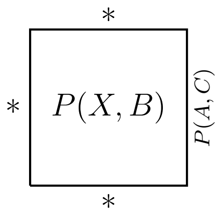

$\in \pi_{q-1}(P(X,B), P(A,C))$

In any case the claim of our theorem can be rewritten as $\pi_q(X; A, B) = 0$ for $2 \leq q \leq m + n - 2$. By taking a CW approximation of our triad and using the fact that I^q is compact, we may assume without loss of generality that (A, C) and (B, C) are finite relative CW complexes. Moreover, we may further reduce to the case where $A = C \cup D^m$ and $B = C \cup D^n$. Indeed, if the claim is true in this case then when B is built from C with more than one cell we can take an intermediary complex $C \subset B' \subset B$. Setting $X' = A \cup_C B'$ the inclusion map factors as $(A, C) \hookrightarrow (X', B') \hookrightarrow (X, B)$, and both maps induce isomorphisms by induction. For induction on the size of A we again take an intermediary complex $C \subset A' \subset A$ and set $X' = A' \cup_C B$. Then the long exact sequences for the triples (A, A', C) and (X, X', B) give us the following commutative diagram.

$$\begin{array}{ccccccccc}
\pi_{q+1}(A, A') & \longrightarrow & \pi_q(A', C) & \longrightarrow & \pi_q(A, C) & \longrightarrow & \pi_q(A, A') & \longrightarrow & \pi_{q-1}(A', C) \\
\downarrow & & \downarrow & & \downarrow & & \downarrow & & \downarrow \\
\pi_{q+1}(X, X') & \longrightarrow & \pi_q(X', B) & \longrightarrow & \pi_q(X, B) & \longrightarrow & \pi_q(X, X') & \longrightarrow & \pi_{q-1}(X', B)
\end{array}$$

Applying the Five Lemma for $2 \leq q < m+n-2$ and the Four Lemma for $q = m+n-2$ completes the inductive step. So we have reduced to the case $A = C \cup D^m$ and $B = C \cup D^n$.

Let D^m_r and D^n_r denote the open discs of radius r contained in the closed unit discs attached to C. Take points $x \in D^m_{1/2}$ and $y \in D^n_{1/2}$ and consider the following relationships among triads.

$$(X - \{y\}; X - \{y\}, X - \{x, y\}) \subset (X; X - \{y\}, X - \{x\}) \simeq (X; A, B)$$

The homotopy equivalence is given by retracting the punctured discs radially to their boundaries. Since the first triad has no nontrivial homotopy groups. A proof of this (though with different names of faces) can be seen in the first picture of our proof by allowing the right face to map to all of X rather than just A; use the arrow as the null-homotopy. So we will be done if we can homotope an arbitrary map into $(X; A, B)$ to a map into $(X - \{y\}; X - \{y\}, X - \{x, y\})$ with a homotopy that remains in $(X; X - \{y\}, X - \{x\})$. To that end, let f represent some element of $\pi_q(X; A, B)$. The homotopy we will produce consists of deforming I^q across some time coordinate and composing with f. The idea behind how we will manipulate the cube is illustrated in the following pictures of f in which π denotes the projection onto the first $q-1$ coordinates.

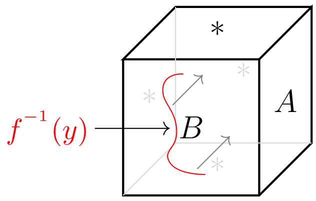

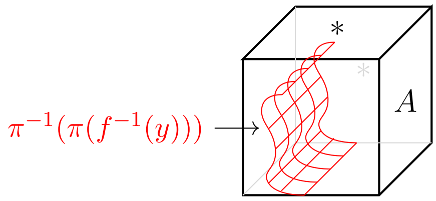

If the points in $f^{-1}(y)$ were removed from the cube then of course applying f to what remains would map into $X - \{y\}$ as desired. Removing points, however, does not constitute a continuous deformation of the cube. The next best thing is to push these points to a location where it is known that f maps nothing to y, like the back face as shown in the first picture. If we only drag $f^{-1}(y)$ back then by continuity other points will be left in its place defeating the purpose of dragging it. To avoid this, we push the entire prism $\pi^{-1}(\pi(f^{-1}(y)))$ from the front face straight to the back. This is a continuous deformation of the cube that leaves nothing in $f^{-1}(y)$. The reason this argument doesn't complete the proof is that there is an additional requirement on the front face - it must always avoid x. As we push the prism in the second picture from front to back we are indenting a slice of the front face and pulling it through the interior of cube, and we need to know that this slice avoids things that map to x since we require that the front face maps into $X - \{x\}$. So the following picture depicts the desired situation where the proof would essentially be complete. The projections of $f^{-1}(x)$ and $f^{-1}(y)$ onto the front face are disjoint so that the red slice may be pushed back without every passing through the blue.

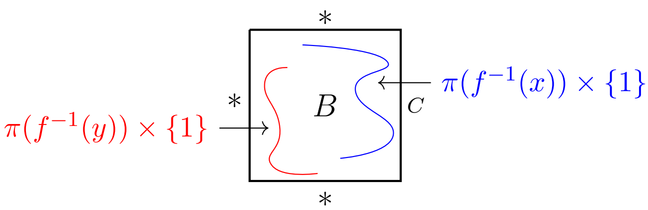

Our goal is to check that the requirement $m + n \geq q + 2$ makes the codomain large enough to guarantee the existence of a pair of points with the property shown in the picture above. We start by letting $U_r = f^{-1}(D_r^m \cup D_r^n)$. By the Whitney Approximation Theorem, there is a homotopy h between $f|_{U_{3/4}}$ and some smooth map $f' : U_{3/4} \to D_{3/4}^m \cup D_{3/4}^n$. Taking a partition of unity $\{\rho, \rho'\}$ subordinate to the cover $\{I^q - \overline{U}_{1/2}, U_{3/4}\}$ allows us to define a map $g = \rho f + \rho' f'$ where addition and scalar multiplication on $D_{3/4}^m \cup D_{3/4}^n$ is done via smooth charts as usual. Then g is smooth on $U_{1/2}$ and the map $(x, t) \mapsto \rho(x) f(x) + \rho'(x) h(x, t)$ shows that it is homotopic to f. Moreover, since $f(I^{q-2} \times \{1\} \times I) \subset A$ and $f(I^{q-1} \times \{1\}) \subset B$, by taking f' close to f we may assume that $g(I^{q-2} \times \{1\} \times I) \cap D_{1/2}^n = \emptyset$ and $g(I^{q-1} \times \{1\}) \cap D_{1/2}^m = \emptyset$. This way g will still represent an element of $\pi_q(X; X - \{y\}, X - \{x\})$ for any $x \in D_{1/2}^m$ and $y \in D_{1/2}^n$.

As an open subset of the smooth manifold I^{2q}, the set $V = g^{-1}(D^m_{1/2}) \times g^{-1}(D^n_{1/2})$ is a smooth submanifold. Consider the set $W = \{(v, v') \in V | \pi(v) = \pi(v')\}$. This is the zero set of the smooth submersion $(v, v') \mapsto (v_1 - v'_1, ..., v_{q-1} - v'_{q-1})$, so it is a smooth submanifold of codimension $q - 1$. This means that $g \times g : W \to D^m_{1/2} \times D^n_{1/2}$ is a smooth map between manifolds of dimension $q + 1$ and $m + n$. Since we are assuming $q + 1 < m + n$, it cannot be surjective so take some $(x, y) \notin \mathrm{im}(g \times g)$. We have found our points satisfying $\pi(g^{-1}(x)) \cap \pi(g^{-1}(y)) = \emptyset$ as desired. Now, since we also have $g(\partial I^{q-1} \times I) \cap D^n_{1/2} = \emptyset$ and $y \in D^n_{1/2}$, by Urysohn's Lemma for disjoint closed sets there is a map $\mu : I^{q-1} \to I$ satisfying

$$\mu(\pi(g^{-1}(x)) \cup \partial I^{q-1}) = 0 \qquad \text{and} \qquad \mu(\pi(g^{-1}(y))) = 1.$$

So finally we define the desired homotopy $H : I^{q+1} \to X$ by $H_t(r, s) = g(r, s - st\mu(r))$. Then $H_0 = g$ and it must be that H_1 avoids y since otherwise $y = H_1(r, s) = g(r, s(1 - \mu(r))) = g(r, 0) = *$ is a contradiction. It remains only to check that H_t is a map into $(X; X - \{y\}, X - \{x\})$ for all t. On $\partial I^{q-1} \times I$ and $I^{q-1} \times \{0\}$ we have $\mu(r) = 0$ and $t = 0$, respectively, so $H_t = g$ here. The last face is $I^{q-1} \times \{1\}$ whose image under H_t cannot contain x since otherwise $x = H_t(r, 1) = g(r, 1 - t\mu(r)) = g(r, 1) \notin D^m_{1/2}$ which is again a contradiction. This completes the proof. Comparing the formula for the homotopy to previous pictures, we can see that r, the first $q-1$ coordinates of I^q, is left unchanged by our manipulation of the cube, so no points are being moved left, right, up, or down. The requirement $\mu(\pi(g^{-1}(x))) = 0$ makes sure that nothing in the front face that would pass through $g^{-1}(x)$ is pushed back as t moves from 0 to 1. The requirement $\mu(\pi(g^{-1}(y))) = 1$ makes sure that everything that would map to y is pushed all the way to the back face by the time $t = 1$. □