Expository Paper on the Fundamental Group

Jonah Saks

January 2020

1 Abstract

The fundamental group of a topological space is one of the most important constructions in algebraic topology. It also happens to be one of the simplest. We provide the context required to define the fundamental group, define it, and show some examples. Then we proceed to explore some elegant applications and introduce van Kampen's Theorem.

2 Homotopy

Definition 2.1. A **path** in a space X is a continuous map $f : I \to X$ where I is the unit interval.

Definition 2.2. A **homotopy** of paths in X is a family of paths $f_t : I \to X$, $0 \leqslant t \leqslant 1$, such that:

1. The endpoints $f_t(0) = x_0$ and $f_t(1) = x_1$ are independent of t.
2. The associated map $F : I \times I \to X$ defined by $F(s,t) = f_t(s)$ is continuous.

When two paths f_0 and f_1 are connected in this way by a homotopy f_t, they are said to be homotopic, and we write $f_0 \simeq f_1$

Definition 2.3. A map $f : X \to Y$ is called a **homotopy equivalence** if there is a map $g : Y \to X$ such that $fg \simeq gf \simeq \mathbb{1}$ (the identity map). We then say that X and Y are **homotopy equivalent**.

Example 2.4. Any two paths $f_0(s)$ and $f_1(s)$ in $\mathbb{R}^n$ with the same endpoints are homotopic, by the homotopy:

$$f_t(s) = (1-t)\, f_0(s) + t\, f_1(s)$$

This is illustrated in the continuous deformation displayed on the top of the next page, where $t = 0$ in the leftmost picture and $t = 1$ in the rightmost picture.

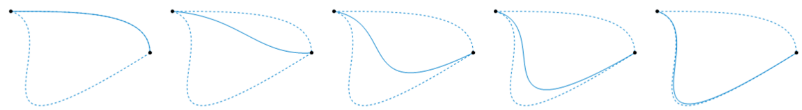

Indeed, using this construction, it is not difficult to show that any convex subset $X \subset \mathbb{R}^n$ satisfies the previous property. Namely, any two paths f_0 and f_1 in X convex with the same endpoints are homotopic. By the previous construction of $f_t(s)$, it suffices to show that this homotopy lies in X, but this is necessarily the case since f_0 and f_1 lie in X and X is convex.

It is not entirely obvious that there exist topological spaces for which this property does not hold. As an example of such a space, consider the 2-torus, denoted $\mathbb{T}^2$. In the illustration below, two paths in $\mathbb{T}^2$ are displayed in red with the same endpoints (each path actually has exactly one endpoint). Denote the two paths $h(s)$ and $h'(s)$. We claim that there does not exist a homotopy $f_t(s)$ with $f_0(s) = h(s)$ and $f_1(s) = h'(s)$ such that $F(s,t)$ is continuous.

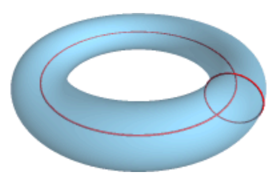

Intuitively, this makes sense. Any homotopy sending the small circular path to the larger one must involve "breaking" the smaller one at some point. We will return to this example several times later in this paper once we have established more tools to investigate this topological space.

Proposition 2.5. *The relation of homotopy on paths with fixed endpoints in any space is an equivalence relation. We denote this relation by* $\simeq$.

Proof. $f \simeq f$ is achieved by taking the constant homotopy $f_t = f$. For symmetry, if we have $f_0 \simeq f_1$ by the homotopy f_t, then $f_1 \simeq f_0$ by the homotopy f_{1-t}. Finally, if $f_0 \simeq f_1$ by the homotopy g_t and $f_1 \simeq f_2$ by the homotopy g'_t, then we define the homotopy $h_t = H(s,t)$ in terms of its associated map as:

$$H(s,t) = \begin{cases} G(s,2t) & 0 \leqslant t \leqslant \frac{1}{2} \\ G'(s,2t-1) & \frac{1}{2} \leqslant t \leqslant 1 \end{cases}$$

Clearly, by continuity of G and G', we have that H is continuous on $I \times [0, \frac{1}{2}]$ and $I \times [\frac{1}{2}, 1]$. Thus, since a function defined on the union of two closed sets is continuous if it is continuous on each of these closed sets separately, then H is continuous on $I \times I$. Then $f_0 \simeq f_2$ via the homotopy h_t, satisfying transitivity. □

Example 2.6. In light of this equivalence relation, we know from Example 2.4 that there is only one homotopy equivalence class of paths in any convex subset $X \subset \mathbb{R}^n$. Much more subtle is the case of $\mathbb{T}^2$. All we have shown up to now is that there are at least two homotopy equivalence classes of paths on $\mathbb{T}^2$, and so $\mathbb{T}^2$ is not homotopy equivalent to $\mathbb{R}^n$ for every $n \in \mathbb{N}$.

In the next section, we define the fundamental group, a natural and elegant extension of homotopy equivalence classes.

3 Fundamental Group

Definition 3.1. Let $f, g : I \to X$ be two paths such that $f(1) = g(0)$, then the product path $f \cdot g$ which first traverses f then traverses g is defined by

$$f \cdot g(s) = \begin{cases} f(2s) & 0 \leqslant s \leqslant \frac{1}{2} \\ g(2s-1) & \frac{1}{2} \leqslant s \leqslant 1 \end{cases}$$

Note that f and g are traversed twice as quickly in order for the product path to be completed in unit time. Also, notice that this path product respects homotopy equivalence. Indeed, if we let f_0, f_1, f_2, g_t and g'_t be as in the proof of Proposition 2.5, then $f_0 \cdot f_1$ and $f_1 \cdot f_2$ are well-defined, so $f_0 \cdot f_1 \simeq f_1 \cdot f_2$ by the homotopy $g_t \cdot g'_t$. Henceforth, we will restrict our attention to paths $f : I \to X$ with the same start and end point (like in the example of the 2-torus). We call such paths **loops**, and we refer to the endpoint as the **basepoint**.

Theorem 3.2. *Let $\pi_1(X, x_0)$ be the set of all homotopy equivalence classes of loops with basepoint $x_0 \in X$. $\pi_1(X, x_0)$ is a group with operation $[f][g] = [f \cdot g]$.*

Proof. By restricting our attention to loops, we are ensured that $f \cdot g$ is well-defined given two such loops f and g. Also, we have already observed that the path product respects homotopy equivalence. Thus, the operation of the group is well-defined. It remains to prove the three group axioms.

First, we define the **reparameterization** of a path f to be the composition $f\phi$, where $\phi : I \to I$ is some continuous map such that $\phi(0) = 0$ and $\phi(1) = 1$. We claim that reparameterization preserves homotopy (ie. $f\phi \simeq f$). To see this, consider the family of paths $\phi_t(s) = (1-t)\phi(s) + ts$ so that $\phi_0(s) = \phi(s)$ and $\phi_1(s) = s$. Notice that $\phi_t(s)$ lies between $\phi(s)$ and s, so it is in I for all t. Then the composition $f\phi_t(s)$ is a homotopy from $f\phi$ to f.

Suppose that we are given paths f, g, h such that $f(1) = g(0)$ and $g(1) = h(0)$. Then $f \cdot (g \cdot h)$ and $(f \cdot g) \cdot h$ are both defined. In the first product, f is traversed in half time, and g and h are both traversed in one quarter of the time. In the second product, f and g are traversed in quarter time whereas h is traversed in half time. Thus, choosing $\phi(s) = 2s$ on the interval $0 \leqslant s \leqslant \frac{1}{4}$, $\phi(s) = s$ on the interval $\frac{1}{4} \leqslant s \leqslant \frac{1}{2}$, and finally $\phi(s) = \frac{s}{2}$ on the interval $\frac{1}{2} \leqslant s \leqslant 1$, we see that $((f \cdot g) \cdot h)\phi = f \cdot (g \cdot h)$, hence we have that $f \cdot (g \cdot h) \simeq (f \cdot g) \cdot h$.

Now for identity, given a path $f : I \to X$, let c be the constant path such that $c(s) = f(1)$. Then it is clear that $f \cdot c$ is a reparameterization of f. Similarly, if we let c be the constant path such that $c(s) = f(0)$, then we have that $c \cdot f$ is a reparameterization of f. Since we have restricted our attention to loops at basepoint x_0, then in fact if we let c be the constant path at x_0, then $c \cdot f \simeq f \simeq f \cdot c$ for all loops f at basepoint x_0. Thus, we have proved the existence of an identity element in the set $\pi_1(X, x_0)$.

Finally, we define the **inverse path** $\overline{f} := f(1-s)$, and we claim that this is indeed the inverse of the element $[f] \in \pi_1(X, x_0)$. Let f_t be equal to f on the interval $[0, 1-t]$ and constant at $f(1-t)$ on the interval $[1-t, 1]$, and let g_t be the inverse path of f_t. Then let $h_t = f_t \cdot g_t = f_t \cdot \overline{f}_t$. Then since $f_0 = f$ and f_1 is the constant path c at x_0, then h_t is a homotopy from $f \cdot \overline{f}$ to $c \cdot \overline{c}$, so $f \cdot \overline{f} \simeq c \cdot \overline{c} = c$. Replacing f by $\overline{f}$, we get that $\overline{f} \cdot f \simeq c$. Thus, $[\overline{f}]$ is indeed the two-sided inverse of the element $[f] \in \pi_1(X, x_0)$.

□

If X is path-connected, the group $\pi_1(X, x_0)$ is independent of the choice of basepoint x_0, up to isomorphism. In this case, we may write $\pi_1(X, x_0)$ as $\pi_1(X)$. Next, we compute one of the simpler examples of a non-trivial fundamental group; that of a circle. To do so, it is necessary to introduce a few more concepts which are, in fact, quite relevant to any further study of the concepts presented in this paper. We introduce these necessary concepts and state two lemmas without proof. The proof of these lemmas is beyond the scope of this paper.

Theorem 3.3. *The fundamental group of the circle is an infinite cyclic group generated by the homotopy class of the loop $w(s) = (\cos(2\pi s), \sin(2\pi s))$ based at $x_0 := (1, 0)$, or equivalently:*

$$\pi_1(S^1, x_0) \cong \mathbb{Z}$$

As explained above, we begin with some preliminary definitions before the proof.

Definition 3.4. Given a topological space X, a **covering space** of X is a topological space $\widetilde{X}$ and a map $p : \widetilde{X} \to X$ such that the following condition is satisfied:

For each point $x \in X$, there is an open neighbourhood U of X such that $p^{-1}(U)$ is a union of disjoint open sets, each of which is mapped homeomorphically onto U by p.

Definition 3.5. Given a path $f : I \to X$ and a covering space $\widetilde{X}$ with the associated map $p : \widetilde{X} \to X$, we say that the path $\widetilde{f} : I \to \widetilde{X}$ is a **lift** of the path f if they satisfy $p\widetilde{f} = f$.

In order to prove the theorem above, we only need the following two lemmas:

Lemma 3.6. *For each path $f : I \to X$ starting at $x \in X$ and each $\widetilde{x} \in p^{-1}(x)$, there exists a unique lift $\widetilde{f} : I \to \widetilde{X}$ starting at $\widetilde{x}$.*

Lemma 3.7. *For each homotopy of paths $f_t : I \to X$ starting at $x \in X$ and each $\widetilde{x} \in p^{-1}(x)$, there exists a unique lifted homotopy $\widetilde{f}_t : I \to \widetilde{X}$ of paths starting at $\widetilde{x}$.*

Proof of Theorem 3.3

Proof. We will use $\mathbb{R}$ as a covering space of S^1, via the map $p : \mathbb{R} \to S^1$ given by $p(s) = (\cos(2\pi s), \sin(2\pi s))$. Geometrically, this map can be visualized as embedding $\mathbb{R}$ into $\mathbb{R}^3$ as the helix parameterized by $s \to (\cos(2\pi s), \sin(2\pi s), s)$, and then restricting to just the first two coordinates: $(\cos(2\pi s), \sin(2\pi s))$.

Define $w_n(s) = (\cos(2\pi ns), \sin(2\pi ns))$ for $n \in \mathbb{Z}$. Note that $[w]^n = [w_n]$, thus the theorem is equivalent to the fact that every loop in S^1 based at $x_0 = (1, 0)$ is homotopic to w_n for some unique $n \in \mathbb{Z}$.

Let $f : I \to S^1$ be a loop with basepoint x_0. By Lemma 3.6, there is a unique lift $\widetilde{f} : I \to R$ starting at 0. This path $\widetilde{f}$ ends at some integer $n \in \mathbb{R}$ since $p\widetilde{f}(1) = f(1) = x_0$, and since $p^{-1}(x_0) = \mathbb{Z} \subset \mathbb{R}$. Another path from 0 to n in $\mathbb{R}$ is $\widetilde{w}_n$. Notice that $\widetilde{f} \simeq \widetilde{w}_n$, by the linear homotopy $(1-t)\widetilde{f} + t\widetilde{w}_n$. We can compose this homotopy with p, thereby constructing a homotopy from f to w_n. Thus, we have that $f \simeq w_n$, or $[f] = [w_n]$.

It remains to show that n is uniquely determined by $[f]$. Suppose that $w_m \simeq f \simeq w_n$. Let f_t be a homotopy from $w_m = f_0$ to $w_n = f_1$. By Lemma 3.7, this homotopy lifts to a homotopy $\widetilde{f}_t$ of paths starting at 0. By uniqueness of the lift of f_0 and f_1, we have that $\widetilde{f}_0 = \widetilde{w}_m$ and $\widetilde{f}_1 = \widetilde{w}_n$. By definition of homotopy, we know that the endpoint of $\widetilde{f}_t(1)$ is independent of t. For $t = 0$, the endpoint is $\widetilde{f}_0(1) = \widetilde{w}_m(1) = m$. For $t = 1$, the endpoint is $\widetilde{f}_1(1) = \widetilde{w}_n(1) = n$. Therefore, $m = n$. □

It is possible to derive several interesting, yet seemingly unrelated, results from the theorem above. For example, the Fundamental Theorem of Algebra and the Brouwer fixed point theorem in dimension 2 can be proven as corollaries of Theorem 3.3. For proofs of these applications and for the proofs of Lemma 3.6 and Lemma 3.7, we refer the reader to [1]. Finally, we prove one more proposition which will render Theorem 3.3 even more powerful still.

Proposition 3.8. *If X and Y are two path-connected topological spaces, then* $\pi_1(X \times Y) \cong \pi_1(X) \times \pi_1(Y)$

Proof. Let $g : Z \to X$ and $h : Z \to Y$ be continuous, then $f : Z \to X \times Y$ defined by $f(z) = (g(z), h(z))$ is continuous. Thus, a loop f in $X \times Y$ based at (x_0, y_0) is equivalent to a pair of loops $g \in X$ and $h \in Y$ based at $x_0 \in X$ and $y_0 \in Y$ respectively. Analogously, a homotopy f_t of a loop in $X \times Y$ is equivalent to a pair of homotopies g_t and h_t of corresponding loops in X and Y respectively. Thus, we obtain the desired bijection $\pi_1(X \times Y) \approx \pi_1(X) \times \pi_1(Y)$. This is clearly a group homomorphism, so we have the desired isomorphism. □

Example 3.9. The unit cylinder is the topological product of S^1 with I. Therefore, by Proposition 3.8, the fundamental group of the unit cylinder is

$$\pi_1(S^1 \times I, (x_0, y_0)) \cong \pi_1(S^1, x_0) \times \pi_1(I, y_0) \cong \mathbb{Z} \times \{e\} \cong \mathbb{Z}$$

Example 3.10. Let us return to the example of the 2-torus. $\mathbb{T}^2 = S^1 \times S^1$, thus by Proposition 3.8, we can compute its fundamental group:

$$\pi_1(S^1 \times S^1, (x_0, y_0)) \cong \pi_1(S^1, x_0) \times \pi_1(S^1, y_0) \cong \mathbb{Z} \times \mathbb{Z}$$

Note that we have shown that the fundamental group of $\mathbb{T}^2$ is abelian. Intuitively, this means that a closed path which circles the torus 'longitudinally' then 'latitudinally' can be continuously deformed into a closed path which circles the torus first 'latitudinally' then 'longitudinally'.

The last proposition was elegant, but there are plenty of relatively simple topological spaces which cannot be easily described as a topological product of simpler spaces.

Example 3.11. The graph illustrated below can be easily expressed as the **wedge sum** (definition in the next section) of two circles, denoted $S^1 \vee S^1$, but does not admit a simple description in terms of a topological product.

It becomes clear that we are in need of a much more powerful tool to compute fundamental groups of more general topological spaces. Van Kampen's theorem will serve as this crucial tool, and will be introduced in the coming section. If we denote the circles on the right and left in the above illustration by A and B respectively, we know that $\pi_1(A) = \mathbb{Z} = \pi_1(B)$. Certainly, arbitrary products of a and b (the generators of A and B respectively) should be elements of the fundamental group of the graph above. For example, aaa and bb should be elements of our group, but so should $abaa$ and $aba^{-1}bab$. Intuitively, the fundamental group of the graph should be some *combination* of $\pi_1(A)$ and $\pi_1(B)$. In fact, van Kampen's theorem will confirm this intuition and will make this *combination* precise.

4 Van Kampen's Theorem

Definition 4.1. Given groups $G_1, G_2, ..., G_n$, a **word** in $\bigcup_{i=1}^n G_i$ is a product $s_1 s_2 ... s_m$ such that, for all j, there exists some i such that $s_j \in G_i$.

Definition 4.2. A word is called **reduced** if for all j, s_j is not equal to the identity of some group G_i, and if consecutive elements of the product are in different groups, (ie. $\nexists i$ such that $s_j, s_{j+1} \in G_i$).

Definition 4.3. The **free product** of $G_1, G_2, ...G_n$ is the set of reduced words in $\bigcup_{i=1}^n G_i$, where the operation is concatenation followed by reduction. We denote this product by $*_{i=1}^n G_i$, or simply as $*_i G_i$.

It is quite clear that the free product of groups forms a group, where the empty word acts as the identity element. For a proof of this, we refer the reader to [1]. We are now ready to present van Kampen's theorem.

Theorem 4.4. *If X is the union of path connected open sets A_α each containing the basepoint $x_0 \in X$ and if each intersection $A_\alpha \cap A_\beta$ is path connected, then the homomorphism $\Phi : *_\alpha \pi_1(A_\alpha) \to \pi_1(X)$ is surjective. If each intersection $A_\alpha \cap A_\beta \cap A_\gamma$ is also path connected, then the kernel of Φ is the normal subgroup $N = \{\langle i_{\alpha\beta}(w) i_{\beta\alpha}(w^{-1}) \rangle : w \in \pi_1(A_\alpha \cap A_\beta)\}$, where $i_{\alpha\beta} : \pi_1(A_\alpha \cap A_\beta) \to \pi_1(A_\alpha)$ is the homomorphism induced by the inclusion $A_\alpha \cap A_\beta \hookrightarrow A_\alpha$. Hence, we have the induced isomorphism $\pi_1(X) \cong *_\alpha \pi_1(A_\alpha) \,/\, N$.*

The proof of this theorem is beyond the scope of this paper. Again, we refer the curious reader to [1] for the proof. For a categorical approach to this theorem, we refer the reader to [2].

Often, this theorem will be applied to a union of two spaces $A_\alpha \cup A_\beta$, making the extra condition on the path connectedness of triples $A_\alpha \cap A_\beta \cap A_\gamma$ superfluous. In this case, if $A_\alpha \cap A_\beta$ is path connected, then we obtain the isomorphism $\pi_1(X) \cong \pi_1(A_\alpha) * \pi_1(A_\beta) \,/\, N$. Next, we return to the example of the wedge sum of circles as displayed in Example 3.11.

Example 4.5. The pair of circles joined together at one point can be expressed as the wedge sum of two circles. Let x_0 and y_0 be points on each circle and denote the graph X, then $X = S^1 \vee S^1 = \{(S^1 \sqcup S^1) \,/\, \sim\}$, where $\sim$ is the equivalence closure of the relation $\{(x_0, y_0)\}$. If we denote the circles A and B, then Theorem 4.4 tells us that $\pi_1(X) = \pi_1(A \vee B) \cong \pi_1(A) * \pi_1(B) \,/\, N$, where $N = \{\langle i_{AB}(w) i_{BA}(w^{-1}) \rangle : w \in \pi_1(A \cap B)\}$. However, $A \cap B = \{x_0\}$, so it must be that $\pi_1(A \cap B) = 0$, (since there is only one possible map from I to x_0, then there cannot be two different paths, so the fundamental group of any point is trivial). Thus, N is trivial, meaning:

$$\pi_1(A \vee B) \cong \pi_1(A) * \pi_1(B) = \pi_1(S^1) * \pi_1(S^1) \cong \mathbb{Z} * \mathbb{Z}$$

Indeed, this result confirms our intuition presented in Example 3.11. We have shown more generally that given topological spaces X_α, it is the case that:

$$\pi_1(\bigvee_\alpha X_\alpha) \cong *_\alpha \pi_1(X_\alpha)$$

As a final example, we perform a slightly more elaborate computation. We attempt to compute the fundamental group of a genus-2 surface, illustrated on the next page.

Example 4.6. We begin our computation with the observation that our desired genus-2 surface, denoted Y, can be decomposed into two tori minus an open disk. Thus, let $U = V = \mathbb{T}^2 - D^2$, then our surface can be constructed from U and V by connecting them along their missing disks. We claim that U is homotopy equivalent to the figure eight loop displayed in Example 3.11. One way to see this is to modify the common construction of a torus. The common construction of a torus from a rectangle is as follows: starting with a rectangle, glue one pair of opposite edges together, creating a cylinder, then glue opposite ends of the cylinder to construct the torus. If, instead, the initial rectangle has a hole in it, then we end up with our desired punctured torus. If we imagine the hole on the initial rectangle growing very large as to only leave the boundary of the rectangle left, then it becomes clear that the procedure described above constructs two loops joined at a point, as desired. Then by our prior computation $\pi_1(S^1 \vee S^1) \cong \mathbb{Z} * \mathbb{Z}$, we know that $\pi_1(\mathbb{T}^2 - D^2) \cong \mathbb{Z} * \mathbb{Z}$. The following illustrations from left to right, are the genus-2 surface, the initial rectangle in the case of the regular torus construction, and the punctured rectangle in the punctured torus construction.

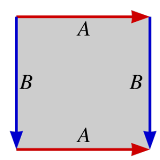

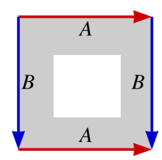

Clearly, $U \cap V = S^1$, so it is path-connected. Then, van Kampen's theorem applied to the union of two spaces gives $\pi_1(Y) \cong \pi_1(U) * \pi_1(V)/N$ where the normal subgroup $N = \{\langle i_u(w) i_v(w^{-1}) \rangle : w \in \pi_1(U \cap V) = \pi_1(S^1)\}$. Here, $i_u : \pi_1(S^1) \to \pi_1(U)$ and $i_v : \pi_1(S^1) \to \pi_1(V)$. We know that $\pi_1(U \cap V)$ is singly-generated by Theorem 3.3, so let g denote the generator, which is just a single loop around $U \cap V$. To compute N, it suffices to compute $i_u(g)$ and $i_v(g^{-1})$. i_u is induced by the inclusion $U \cap V \hookrightarrow U$, so $i_u(g)$ is the equivalence class of the loop around $U \cap V$ in terms of the generators of $\pi_1(U) \cong \mathbb{Z} * \mathbb{Z}$. Let $\pi_1(U) = \langle a, b \rangle$ denote the generators. If we think about this loop as existing on the initial punctured rectangle as opposed to on the punctured torus itself, it becomes clear that the loop around the hole is homotopic to the word $aba^{-1}b^{-1}$. Thus, $i_u(g) = aba^{-1}b^{-1}$. Analogously, if we denote the generators of $\pi_1(V) = \langle c, d \rangle$, then we find that $i_v(g^{-1}) = c^{-1}d^{-1}cd$, so our final result is as follows:

$$\begin{aligned}\pi_1(Y) &\cong \pi_1(U) * \pi_1(V) \,/\, \{\langle i_u(w) i_v(w^{-1}) \rangle : w \in \pi_1(U \cap V) = \pi_1(S^1)\} \\ &\cong \mathbb{Z} * \mathbb{Z} * \mathbb{Z} * \mathbb{Z} \,/\, \{\langle i_u(g) i_v(g^{-1}) \rangle : g \in \pi_1(S^1) = \langle g \rangle\} \\ &\cong \langle a, b, c, d : aba^{-1}b^{-1}c^{-1}d^{-1}cd = 1 \rangle\end{aligned}$$

5 Closing Remarks

It is hoped that the elegance of these concepts has been properly conveyed to the reader. The purpose of this paper was for it to serve as a brief introduction into the landscape of algebraic topology. If this paper has successfully seduced the reader to learn more about the ideas presented here and what comes next, we provide some resources on the next page.

Hatcher's *Algebraic Topology* [1] was followed closely in the writing of this paper. For the presentation of concepts in a manner similar in style to this paper,

[1] is strongly recommended. For a presentation of van Kampen's theorem with a more categorical flavour, we recommend looking into [2]. Some concepts in the field which follow naturally from those presented in this paper are cell complexes, covering spaces, and homology.

References

[1] Allen Hatcher. *Algebraic Topology.* Cambridge University Press, 2002. ISBN: 9780521795401.

[2] J. P. May. *A Concise Course in Algebraic Topology.* The University of Chicago Press, 1999. ISBN: 9780226511832.