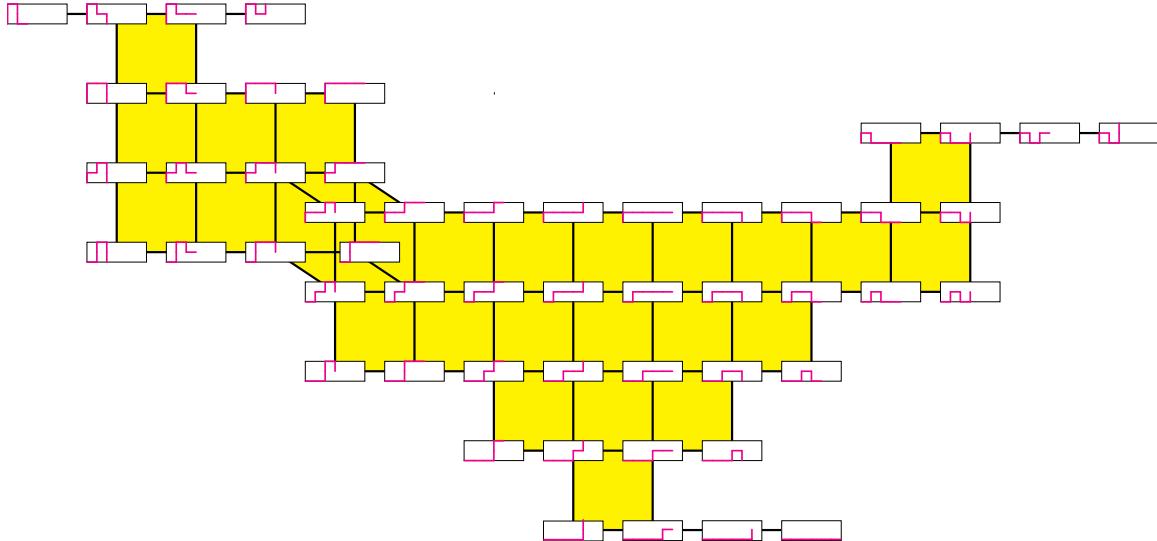


CAT(0) geometry, robots, and society.

Federico Ardila–Mantilla*



1 Moving objects optimally

There are many situations in which an object can be moved using certain prescribed rules, and many reasons – pure and applied – to solve the following problem.

Problem 1. *Move an object optimally from one given position to another.*

Without a good idea, this is usually very hard to do.

When we are in a city we do not know well, trying to get from one location to another quickly, most of us will consult a map of the city to plan our route. This simple, powerful idea is the root of a very useful approach to Problem 1: We build and understand the “map of possibilities”, which keeps track of all possible positions of the object; we call it the *configuration space*. This idea is pervasive in many fields of mathematics, which call such maps *moduli spaces*, *parameter spaces*, or *state complexes*.

This article seeks to explain that, for many objects that move discretely, the resulting “map of possibilities” is a *CAT(0) cubical complex*: a space of non-positive curvature made of unit cubes. When this is the case, we can use ideas from geometric group theory and combinatorics to solve Problem 1.

This approach is applicable to many different settings; but to keep the discussion concrete, we focus on the following specific example. For precise statements, see Section 3 and Theorems 5 and 9.

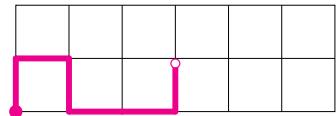


Figure 1: A pinned down robotic arm of length 6 in a tunnel of height 2. The figure above shows its configuration space.

Theorem 2. [?, ?] *The configuration space of a 2-D pinned down robotic arm in a rectangular tunnel is a CAT(0) cubical complex. Therefore there is an explicit algorithm to move this robotic arm optimally from one given position to another.*

2 Black Lives Matter

On July 4, 2016 we finished the implementation of our algorithm to move a discrete robotic arm. Three days later, seemingly for the first time in history, US police used a robot to kill an American citizen. Now, whenever I present this research, I also discuss this action.

*Professor of Mathematics, San Francisco State University. Profesor Adjunto, Universidad de Los Andes. federico@sfsu.edu. This work was supported by NSF CAREER grant DMS-0956178, grants DMS-0801075, DMS-1600609, DMS-1855610, Simons Fellowship 613384, and NIH grant 5UL1GM118985-03 awarded to SF BUILD.

The day started with several peaceful Black Lives Matter protests across the US, condemning the violence disproportionately inflicted on Black communities by the American state. These particular protests were prompted by the shootings of Alton Sterling and Philando Castile by police officers in Minnesota and Louisiana.

In Dallas, TX, as the protest was coming to an end, a sniper opened fire on the crowd, killing five police officers. Dallas Police initially misidentified a Black man – the brother of one of the protest organizers – as a suspect. They posted a photo of him on the internet and asked for help finding him. Fearing for his life, he turned himself in, and was quickly found innocent.

A few hours later, police identified US Army veteran Micah Johnson as the main suspect. After a chase, a standoff, and failed negotiations, they used a robot to kill him, without due process of law.

The organizers of the protest condemned the sniper's actions, and police officials believe he acted alone. The robot that killed Johnson cost about \$150,000; police said that the arm of the robot was damaged, but still functional after the blast [?]. The innocent man who was misidentified by the police continued to receive death threats for months afterwards.

Different people will have different opinions about the actions of the Dallas Police in this tragic event. What is certainly unhealthy is that the large majority of people I have spoken to have never heard of this incident.

Our mathematical model of a robotic arm is very simplified, and probably far from direct applications, but the techniques developed here have the potential to make robotic operations cheaper and more efficient. We tell ourselves that mathematics and robotics are neutral tools, but our research is not independent from how it is applied. We arrive to mathematics and science searching for beauty, understanding, or applicability. When we discover the power that they carry, how do we proceed?

Axiom. [?] *Mathematics is a powerful, malleable tool that can be shaped and used differently by various communities to serve their needs.*

Who currently holds that power? How do we use it? Who funds it and for what ends? With whom do we share that power? Which communities benefit from it? Which are disproportionately harmed by it?

For me these are the hardest questions about this work, and the most important. The second goal in writing this article – a central one for me – is to invite myself, and its readers, to continue to look for answers that make sense to us.

3 Moving robots

We consider a discrete 2-D robotic arm $R_{m,n}$ of length n moving in a rectangular tunnel of height m . The robot consists of n links of unit length, attached sequentially, facing up, down, or right. Its base is affixed to the lower left corner of the tunnel, as shown in Figure 1 for $R_{2,6}$.

The robotic arm may move freely, as long as it doesn't collide with itself, using two kinds of local moves:

- *Flipping a corner:* Two consecutive links facing different directions interchange directions.
- *Rotating the end:* The last link of the robot rotates 90° without intersecting itself.

This is an example of a *metamorphic robot* [?].

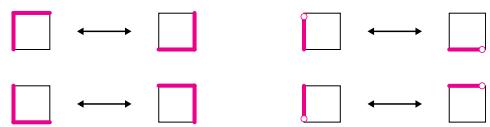


Figure 2: The local moves of the robotic arm.

How can we get the robot to navigate this tunnel efficiently?

Figure 3 shows two positions of the robot; suppose we want to move it from one position to the other. By trial and error, one will not have too much difficulty in doing it. It is not at all clear, however, how one might do this in the most efficient way possible.

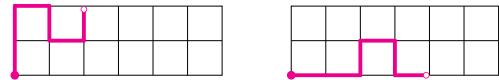


Figure 3: Two positions of the robot $R_{2,6}$.

4 Maps

To answer this question, let us build the “map of possibilities” of the robot. We begin with a *configuration graph*, which has a node for each position of the arm, and an edge for each local move between two positions. A small piece of this graph is shown below.

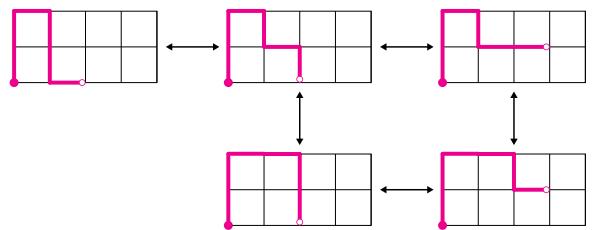


Figure 4: A part of the graph of possibilities of $R_{2,6}$.

As we see in the figure of page 1, the resulting graph looks a bit like the map of downtown San Francisco or Bogotá, with many square blocks lined up neatly. Such a

cycle of length 4 arises whenever the robot is in a given position, and there are two moves A and B that do not interfere with each other: if we perform move A and then move B, the result is the same as if we perform move B and then move A; see for example the 4-cycle of Figure 4. More generally, if the robot has k moves that can be performed independently of each other, these moves result in (the skeleton of) a k -dimensional cube in the graph.

This brings up an important point: If we wish to move the robot efficiently, we should let it perform various moves simultaneously. In the map, this corresponds to walking across the diagonal of the corresponding cube. Thus we construct the *configuration space* of the robot, by filling in the k -cube corresponding to any k moves that can be performed simultaneously, as illustrated in Figure 5; compare with Figure 4. The result is a *cubical complex*, a space made of cubes that are glued face-to-face.

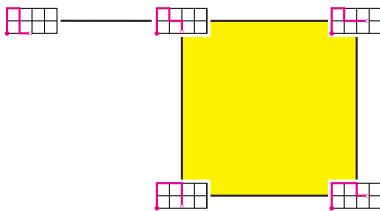


Figure 5: A part of the configuration space of $R_{2,6}$.

Definition 3. *The configuration space $\mathcal{C}(R)$ of the robotic arm R is the cubical complex with:*

- a vertex for each position of the robot,
- an edge for each local move between two positions,
- a k -dimensional cube for each k -tuple of local moves that may be performed simultaneously.

This definition applies much more generally to discrete situations that change according to local moves; see Section 9 and [?, ?].

In our specific example, Figure 6 shows the configuration space of the robot $R_{2,6}$ of length 6 in a tunnel of height 2. It is now clear how to move between two positions efficiently: just follow the shortest path between them in the map!

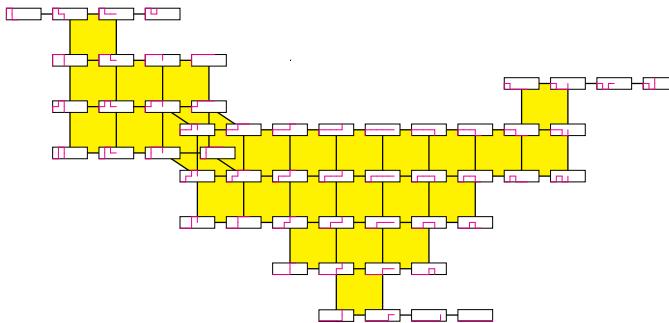


Figure 6: The configuration space of the robotic arm $R_{2,6}$.

5 What are we optimizing?

Is it so clear, just looking at a map, what the optimal path will be? It depends on what we are trying to optimize. In San Francisco, with its beautifully steep hills, the best route between two points can be very different depending on whether one is driving, biking, walking, or taking public transportation. The same is true for the motion of a robot.

For the configuration spaces we are studying, there are at least three reasonable metrics: ℓ_1 , ℓ_2 , and ℓ_∞ . In these metrics, the distance between points x and y in the same d -cube, say $[0, 1]^d$, is

$$\sqrt{\sum_{1 \leq i \leq d} (x_i - y_i)^2}, \quad \sum_{1 \leq i \leq d} |x_i - y_i|, \quad \max_{1 \leq i \leq d} |x_i - y_i|,$$

respectively. Figure 7 shows the two positions of the robot of Figure 3 in the configuration space, and shortest paths or *geodesics* between them according to these metrics.

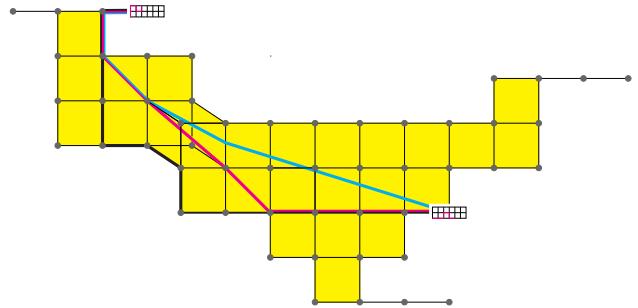


Figure 7: Some paths between two points in the configuration space $\mathcal{C}(R_{2,6})$. The black path is geodesic in the ℓ_1 metric, the magenta path is geodesic in ℓ_1 and ℓ_∞ , and the cyan path is geodesic in ℓ_1 , ℓ_2 , and ℓ_∞ .

If each individual move has a “cost” of 1, then performing d simultaneous moves – which corresponds to crossing a d -cube – costs \sqrt{d} , d , and 1 in the metrics ℓ_2 , ℓ_1 , and ℓ_∞ . Although the Euclidean metric is the most familiar, it seems unrealistic in this application; why should two simultaneous moves cost $\sqrt{2}$? For the applications we have in mind, the ℓ_1 and ℓ_∞ metrics are reasonable models for the cost and the time of motion:

Cost (ℓ_1): We perform one move at a time; there is no cost benefit to making moves simultaneously.

Time (ℓ_∞): We may perform several moves at a time, causing no extra delay.

These two metrics, studied in [?, ?], will be the ones that concern us in this paper. The Euclidean metric, which is useful in other contexts and significantly harder to analyze, is studied in [?, ?].

6 Morning routine

I write this while on sabbatical in a foreign city. Being the coffee enthusiast that I am, I carefully study a map several mornings in a row, struggling to find the best cafe on my way from home to my office. One morning, amused, my partner May-Li stops me on the way out:

– Fede, you know you don’t **always** have to take a geodesic, right?

Perhaps, instead of the most efficient paths, we should be looking for the most pleasant, or the greenest, or the most surprising, or the most beautiful.

7 CAT(0) cubical complexes

Our two most relevant algorithmic results are the explicit construction of cheapest (ℓ_1) and fastest (ℓ_∞) paths in the configuration space of the robot arm $R_{m,n}$. Still, the Euclidean metric (ℓ_2) turns out to play a very important role as well. Most configuration spaces that interest us exhibit non-positive curvature with respect to the Euclidean metric, and this fact is central in our construction of shortest paths in the cost and time metrics.

Let us consider a *geodesic metric space* (X, d) , where any two points x and y can be joined by a unique shortest path of length $d(x, y)$; such a path is known as a *geodesic*. Let T be a triangle in X whose sides are geodesics of lengths a, b, c , and let T' be the triangle with the same sidelengths in the plane. For any geodesic chord of length d connecting two points on the boundary of T , there is a comparison chord between the corresponding two points on the boundary of T' , say of length d' . If $d \leq d'$ for any such chord in T , we say that triangle T is *at least as thin as a Euclidean triangle*.



Figure 8: A chord in a triangle in X , and the corresponding chord in the comparison triangle in \mathbb{R}^2 . The triangle in X is *at least as thin as a Euclidean triangle* if $d \leq d'$ for all such chords.

Definition 4. A metric space X is **CAT(0)** if:

- between any two points there is a unique geodesic, and
- every triangle is at least as thin as a Euclidean triangle.

A (finite) *cubical complex* is a connected space obtained by gluing finitely many cubes of various dimensions along their faces. We regard it as a metric space with the Euclidean metric on each cube; all cubes necessarily have the same side length. Cubical complexes are

flat inside each cube, but they can have curvature where cubes are glued together, for example, by attaching three or five squares around a common vertex (obtaining positive and negative curvature, respectively), as shown in Figure 9. We invite the reader to check that the triangles in the left and right panel of Figure 9 are thinner and not thinner than a Euclidean triangle, respectively.

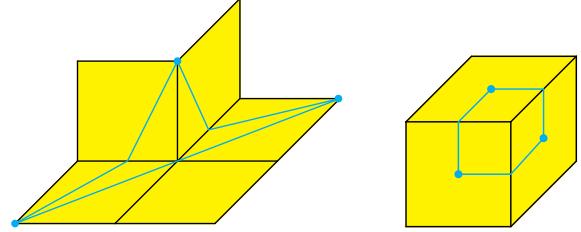


Figure 9: A CAT(0) and a non-CAT(0) cubical complex.

We have the following general theorem.

Theorem 5. [?, ?, ?, ?, ?, ?] Given two points x and y in a CAT(0) cubical complex, there are algorithms to find a geodesic from x to y in the Euclidean (ℓ_2), cost (ℓ_1), and time (ℓ_∞) metrics.

Thus a robot with a CAT(0) configuration space is easier to control: we have a procedure that automatically moves it optimally. We will see that this is the case for the robotic arm $R_{m,n}$.

8 How do we proceed?

Once I began to feel that this work, which started out in “pure” mathematics, could actually have real-life applications, I started getting anxious and selective about who I discussed it with. It is a strange feeling, to discover something you really like, and yet to hope that not too many people find out about it. When I was invited to write this article, I felt conflicted. I knew I did not want to only discuss the mathematics, but I am much less comfortable writing outside of the shared imaginary world of mathematicians, where we believe we know right from wrong. Still, I know it is important to listen, learn, discuss, and even write from this place of discomfort.

How should I tell this story? Should I do it at all? I have turned to many friends, colleagues, and students for their wisdom and advice.

Mario Sanchez, who thinks deeply and critically about the culture of mathematics and philosophy in our society, is wary of mathematical fashions: What if it becomes trendy for mathematicians to start working on optimizing robots, but not to think about what is being optimized, or whom that optimization benefits? He tells me, with his quiet intensity: “If you’re worried that your paper might have this effect, you should probably emphasize the human question pretty strongly.”

Laura Escobar just returned from a yoga retreat in Champaign-Urbana where a scholar of Indian literature taught them the story of Arjuna, a young warrior about to enter a rightful battle against members of his own family. Deeply conflicted about the great violence that will ensue, he turns to Krishna for advice. Oversimplifying his reply, Krishna says: “One should not abandon duties born of one’s nature, even if one sees defects in them. It is your duty as a warrior to uphold the Dharma, take action, and fight.” With her usual thoughtful laugh, she tells me about the distressed reactions of her peace-loving yoga classmates. Laura and I grew up in the middle of Colombia’s 60-year old civil war, which has killed more than 215,000 civilians and 45,000 combatants and has displaced more than 15% of the country’s population [?, ?]; it is hard for us to understand Krishna’s advice as well.¹ So we go to the bookstore and buy matching copies of the Bhagavad Gita.

Many of my friends who do not work in science are surprised by the lack of structural and institutional resources. They ask me: If a mathematician or a scientist is trying to understand or have some control over the societal impact of their knowledge and their expertise, what organizations they can turn to for support? I have been asking this question to many people. I have not found one, but I am collecting resources. Interdisciplinary organizations like the Union of Concerned Scientists, Science for the People, Data for Black Lives, and sections of the American Association for the Advancement of Science seek to use science to improve people’s lives and advance social justice. Our colleagues in departments of Science, Technology, and Society, Public Policy, History, Philosophy, and Ethnic Studies have been studying these issues for decades, even centuries. This has often taken place too far from science departments, and it must be said that my generation of scientists largely looked down on these disciplines as unrigorous, uninteresting, or unimportant. Governments, companies, and professional organizations assemble Ethics Committees, usually separate from their main operations, and give them little to no decision-making power.

How do we make these considerations an integral part of the practice and application of science? I am encouraged to see that the new generation of scientists understands their urgent role in society much more clearly than we do.

May-Li Khoe, whom I can always trust to be wise and direct, asks me: If you tell me that this model of mapping possibilities could be applicable in many areas, and you don’t trust the organizations that build the most powerful robots, why don’t you find other applications? She’s right. I’m looking.

¹We later learn that Robert Oppenheimer quoted Krishna when he and his team detonated the first nuclear bomb.

²We should approach them thoughtfully and critically; see Section 10.

9 Examples.

Just like any other cultural practice, mathematics respects none of the artificial boundaries that we sometimes draw, in an attempt to understand it and control it. This is evident for CAT(0) cubical complexes, a family of objects which appears in many seemingly disparate parts of (mathematical) nature. Let us discuss three sources of examples; each one raises different kinds of questions and offers valuable tools that have directly shaped this investigation.

Geometric group theory. This project was born in geometric group theory, which studies groups by analyzing how they act on geometric spaces. Gromov’s pioneering work in this field [?] led to the systematic study of CAT(0) cubical complexes. A concrete source of examples is due to Davis [?].

A *right-angled Coxeter group* $X(G)$ is given by generators of order 2 and some commuting relations between them; we encode the generators and commuting pairs in a graph G . For example, the graph of Figure 10.a encodes the group generated by a, b, c, d with relations $a^2 = b^2 = c^2 = d^2 = 1, ab = ba, ac = ca, bc = cb, cd = dc$.

The *Cayley graph* has a vertex for each element of $X(G)$ and an edge between g and gs for each group element g and generator s . This graph is the skeleton of a CAT(0) cube complex that G acts on, called the *Davis complex* $\mathcal{S}(G)$. It is illustrated in Figure 10.b. One can then use the geometry of $\mathcal{S}(G)$ to derive algebraic properties of $X(G)$. For example, one can easily solve the *word problem* for this group: given a word in the generators, determine whether it equals the identity. This problem is undecidable for general groups.

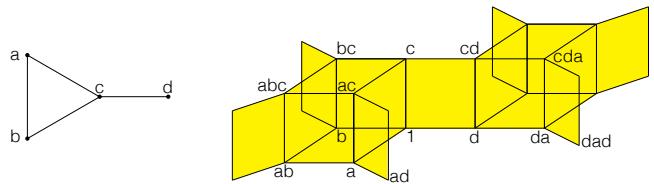


Figure 10: a. A graph G determining a right-angled Coxeter group $X(G)$, and b. part of its Davis complex $\mathcal{S}(G)$.

Phylogenetic trees. A central problem in phylogenetics is the following: given n species, determine the most likely evolutionary tree that led to them. There are many ways of measuring how different two species are²; but if we are given the $\binom{n}{2}$ pairwise distances between the species, how do we construct the tree that most closely fits that data?

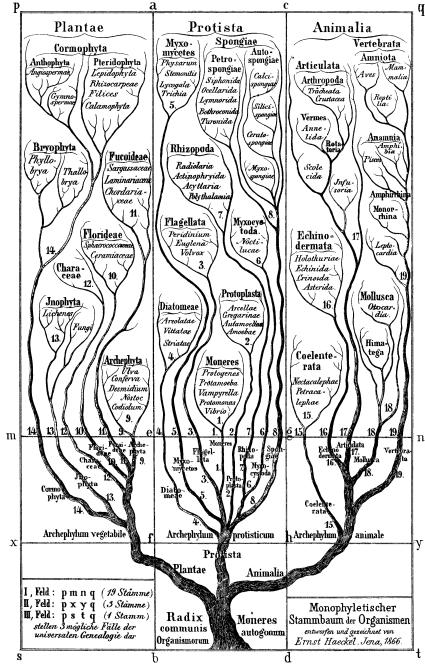


Figure 11: Ernst Haeckel's tree of life (1866).

Billera, Holmes, and Vogtmann [?] approached this problem by constructing the space of all possibilities: the *space of trees* \mathcal{T}_n . Remarkably, they proved that the space of trees \mathcal{T}_n is a CAT(0) cube complex. In particular, since it has unique geodesics, we can measure the distance between two trees, or find the average tree between them. This can be very helpful in applications: if 10 different algorithms propose 10 different phylogenetic trees, we can detect which proposed trees are close to each other, detect outlier proposals that seem unlikely, or find the average between different proposals. Owen and Provan showed how to do this in polynomial time [?].

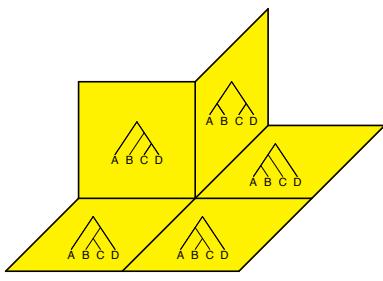


Figure 12: Five of the 15 squares in the space of trees \mathcal{T}_4 .

These results made us wonder whether one can similarly construct ℓ_2 -optimal paths in any CAT(0) cube complex. New complications arise, but it is possible. [?, ?].

Discrete systems: reconfiguration. Abrams, Ghrist, and Peterson introduced *reconfigurable systems* in [?, ?]. This very general framework models discrete objects that change according to local moves, keeping track of

which pairs of moves can be carried out simultaneously. Examples include discrete metamorphic robots moving around a space, particles moving around a graph without colliding, domino tilings of a region changing by *flips* $\square \longleftrightarrow \square$, and reduced words in the symmetric group changing by commutation moves $s_i s_j \leftrightarrow s_j s_i$ for $|i-j| \geq 2$ and braid moves $s_i s_{i+1} s_i \leftrightarrow s_{i+1} s_i s_{i+1}$.

Definition 3 associates a configuration space to any reconfigurable system. Such a configuration space is always **locally** CAT(0). It is often **globally** CAT(0), and when that happens Theorem 5 applies, allowing us to move our objects optimally.

10 Why do we map?

Math historian Michael Barany points out to me that, struck by the aesthetic beauty of the tree of life shown in Figure 11, I failed to notice another map that Haeckel drew: a hierarchical tree of nine human groups – which he regarded as different species – showing their supposed evolutionary distance from the ape-man. Modern biology shows this has no scientific validity, and furthermore, that there is no genetic basis for the concept of race. Haeckel's work is just one sample of the deep historical ties between phylogenetics and scientific racism, and between mapmaking and domination.

If we map from a different – an *other* – point of view [...] then mapping becomes a process of getting to know, connect, bring closer together in relation, remember, and interpret.

– Sandra Alvarez [?]

11 Characterizations

How does one determine whether a given space is CAT(0)? We surely do not want to follow Definition 4 and check whether **every** triangle is at least as thin as a Euclidean triangle; this is not easy to do, even for an example as small as Figure 9. Fortunately, this becomes much easier when the space in question is a cubical complex. In this case, Gromov showed that the CAT(0) property – a subtle metric condition – can be rephrased entirely in terms of topology and combinatorics; no measuring is necessary!

To state this, we recall two definitions. A space X is *simply connected* if there is a path between any two points, and every loop can be contracted to a point. If v is a vertex of a cubical complex X , then the *link* of v in X is the simplicial complex one obtains by intersecting X with a small sphere centered at v . A simplicial complex Δ is *flag* if it has no empty simplices: if A is a set of vertices and every pair of vertices in A is connected by an edge in Δ , then A is a simplex in Δ .

Theorem 6. [?] A cubical complex X is CAT(0) if and only if:

- X is simply connected, and
- the link of every vertex in X is flag.

In fact, one can also do without the topology: there is an entirely combinatorial characterization of CAT(0) cubical complexes. This is originally due to Sageev and Roller, and we rediscovered it in [?] in a different formulation that is more convenient for our purposes. Let a *pointed cubical complex* be a cubical complex with a distinguished vertex.

Definition 7. [?, ?] A poset with inconsistent pairs (PIP) $(P, \leq, \leftrightarrow)$ is a poset (P, \leq) together with a collection of inconsistent pairs, denoted $p \leftrightarrow q$ for $p \neq q \in P$, that is closed under \leq ; that is,

if $p \leftrightarrow q$ and $p \leq p'$, $q \leq q'$, then $p' \leftrightarrow q'$.

PIPs are also known as *prime event structures* in the computer science literature [?]. The *Hasse diagram* of a PIP $(P, \leq, \leftrightarrow)$ shows graphically the minimal relations that define it. It has a dot for each element of P , a solid line from p upward to q whenever $p < q$ and there is no r with $p < q < r$, and a dotted line between p and q whenever $p \leftrightarrow q$ and there are no $r \leq p$ and $s \leq q$ such that $r \leftrightarrow s$.

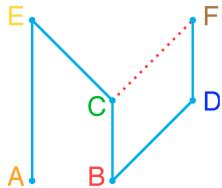


Figure 13: The Hasse diagram of a PIP: solid lines represent the poset, and dotted lines represent the (minimal) inconsistent pairs. Notice that $C \leftrightarrow F$ implies $E \leftrightarrow F$.

Theorem 8. [?, ?, ?] Pointed $CAT(0)$ cube complexes are in bijection with posets with inconsistent pairs (PIPs).

This rediscovery was motivated by the observation that CAT(0) cubical complexes look very much like distributive lattices. In fact, Theorem 8 is an analog of Birkhoff’s representation theorem, which gives a bijection between distributive lattices and posets. The proof is subtle and relies heavily on Sageev’s work [?], but the bijection is easy and useful to describe:

Pointed CAT(0) cubical complex \mapsto PIP: Let (X, v) be a CAT(0) cubical complex X rooted at vertex v . Every d -cube in X has d hyperplanes that bisect its edges. Whenever two cubes share an edge, let us glue the two hyperplanes bisecting it. The result is a system of *hyperplanes* associated to X [?]. Figure 14 shows an example.

The PIP corresponding to (X, v) keeps track of how one can navigate X starting from v . The elements of

the corresponding PIP are the hyperplanes. We declare $H < I$ if, starting from v , one must cross H before crossing I . We declare $H \leftrightarrow I$ if, starting from v , one cannot cross both H and I without backtracking. Remarkably, the simple combinatorial information stored in this PIP is enough to recover the space (X, v) .

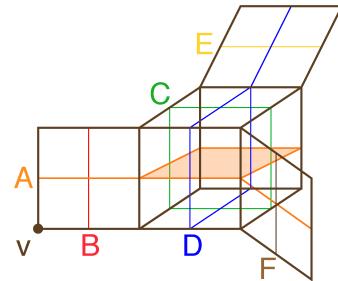


Figure 14: A rooted CAT(0) cubical complex with six hyperplanes. Its PIP is shown in Figure 13.

PIP \longmapsto rooted CAT(0) cubical complex: Let P be a PIP. An *order ideal* of P is a subset I closed under $<$; that is, if $x < y$ and $y \in I$ then $x \in I$. We say that I is *consistent* if it contains no inconsistent pair.

The vertices of the corresponding CAT(0) cubical complex $X(P)$ correspond to the consistent order ideals of P . Two vertices are connected if their ideals differ by a single element. Then we fill in all cubes whose edges are in this graph. The root is the vertex corresponding to the empty order ideal.

We invite the reader to verify that the PIP of Figure 13 corresponds to the rooted complex of Figure 14.

Theorem 8 provides a completely combinatorial way of proving that a cubical complex is CAT(0): one simply needs to identify the corresponding PIP!

12 Remote controls and geodesics

Intuitively, we think of the PIP P as a “remote control” to help an imaginary particle navigate the corresponding CAT(0) cubical complex X . If the particle is at a vertex of X , there is a corresponding consistent order ideal I of P . The hyperplanes that the particle can cross are the maximal elements of I and the minimal elements of $P - I$ consistent with I . We can then press the i th “button” of P if we want the point to cross hyperplane i .

This point of view is powerful because in practical applications, the configuration space X is usually very large, high dimensional, and combinatorially complicated, whereas the remote control P is much smaller and can be constructed in some cases of interest.

Theorem 5 provides algorithms to move optimally between any two points in a CAT(0) cubical complex in the ℓ_1, ℓ_2 , and ℓ_∞ metrics. We sketch the proof in the cases

that are relevant here: in the ℓ_1 and ℓ_∞ metrics, where the two points v and w are vertices.

Sketch of Proof of Theorem 5. To move from v to w , let us root the cube complex X at v , and let P be the corresponding PIP. Then w corresponds to an order ideal I of P ; these are the hyperplanes we need to cross.

Cost (ℓ_1) Metric: We simply cross the hyperplanes from v to w in non-decreasing order, with respect to the poset $I \subseteq P$: we first cross a minimal element $m_1 \in I$, then a minimal element $m_2 \in I - m_1$, and so on.

Time (ℓ_∞) Metric: We first cross all minimal hyperplanes M_1 in I simultaneously, then we cross all the minimal hyperplanes M_2 in $I - M_1$ simultaneously, and so on. This corresponds to Niblo and Reeves's *normal cube path* [?], where we cross the best available cube at each stage. \square

These algorithms show how to move a CAT(0) robot optimally and automatically.

13 Automation

Driving in San Francisco, I get stuck behind a terrible driver. They are going extremely slowly, hesitating at every corner, stalling at every speed bump. When I finally lose patience and decide to pass them, they swerve wildly towards me; I react quickly to avoid being hit. I turn to give the driver a nasty look, but I find there isn't one.

What happens if you are injured by an automated, self-driving vehicle or robot designed by well-meaning scientists and technologists? When you live this close to Silicon Valley, the question is not just philosophical.

14 Prototype: A robotic arm in a tunnel

If we wish to apply Theorem 5 to move an object optimally, our first hope is that the corresponding map of possibilities is a CAT(0) cubical complex. If this is true, we can prove it by choosing a convenient root and identifying the corresponding PIP. Tia Baker and Rika Yatchak pioneered this approach in their Master's theses [?].

For concreteness, let us consider our robotic arm of length n in a rectangular tunnel of height 1. Baker and Yatchak found that the number of states of the configuration space is the term F_{n+1} of the Fibonacci sequence. This seemed like good news, until we realized that these numbers grow exponentially! The dimension of the map is $n/3$, and its combinatorial structure is enormous and intricate. We cannot navigate this map by brute force.

Fortunately, by running the bijection of Theorem 8 on enough examples, Baker and Yatchak discovered that this robot has a very nice PIP: a triangular wedge T_n of a square grid with no inconsistent pairs, as shown in Figure 15. It is much simpler and only has about $n^2/4$

vertices. Indeed they proved that the map of possibilities of the robot $R_{1,n}$ is isomorphic to the cubical complex $X(T_n)$ corresponding to T_n . This implies that the map is CAT(0), and it allows us to use T_n as a remote control to move the robot optimally.

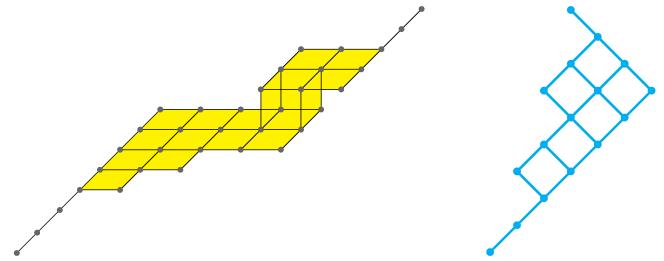


Figure 15: The map of the robot $R_{1,7}$ and its PIP T_7 .

More generally, we have the following.

Theorem 9. [?, ?] *The configuration space of the robotic arm $R_{m,n}$ of length n in a tunnel of height m is a CAT(0) cubical complex. Therefore, we have an algorithm to move the arm optimally from any position to any other.*

Naturally, as the height grows, the map becomes increasingly complex. After staring at many examples, getting stuck, and finally receiving a conclusive hint from the Pacific Ocean – a piece of coral with a fractal-like structure – we were able to describe the PIP of the robot $R_{m,n}$ for any m and n . It is made of triangular flaps like the one in Figure 15 recursively branching out in numerous directions.

This *coral PIP* serves as a witness that the map of possibilities of the robotic arm $R_{m,n}$ is a CAT(0) cubical complex. It can also be programmed to serve as a remote control, to help the arm explore the tunnel.

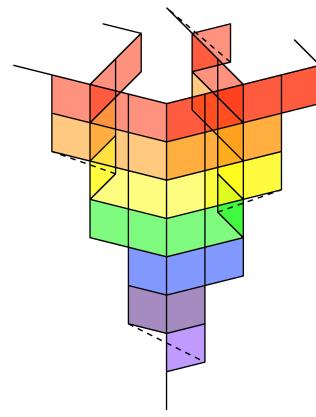


Figure 16: The coral PIPs of the robot $R_{2,9}$, which contains the PIPs of $R_{2,1}, \dots, R_{2,8}$, shown in different colors.

15 Implementation

The algorithms to navigate a CAT(0) space optimally – and hence move a CAT(0) robot, are described in [?]. We have implemented them in Python for the robotic arm in a tunnel [?]. Given two states, the program outputs the distance between the two states in terms of cost (ℓ_1) and time (ℓ_∞), and an animation moving the robot optimally between the two states. The downloadable code, instructions, and a sample animation are at <http://math.sfsu.edu/federico/robots.html>.

With the goal to broaden access to these tools, I joined my collaborator César Ceballos, who led a week-long workshop for young robotics enthusiasts, as part of the Clubes de Ciencia de Colombia. This program invites Colombian researchers to design scientific activities for groups of students from public high schools and universities across the country.

We proposed some discrete models of robotic arms, and our students successfully built their maps of possibilities. Extremely politely, they also pointed out that César and I really didn't know much about the mechanics of robots, and cleverly proposed several possible mechanisms. After the workshop, Arlys Asprilla implemented the design on CAD and built an initial prototype.



Figure 17: a. César Ceballos and students discuss configuration spaces during the Clubes de Ciencia de Colombia. b. Arlys Asprilla and one of his robotic arms.

16 Escuela de Robótica del Chocó

Arlys, his classmate Wolsey Rubio (on the right in Figure 17.a.), my partner May-Li Khoe, our friend Akil King, and I designed a similar workshop in Arlys and Wolsey's native Chocó. This region of the Colombian Pacific Coast is one of the most biodiverse in the world, and also one of the most neglected historically by our government. We partnered with the Escuela de Robótica del Chocó, led by Jimmy García, which seeks to empower local youth to develop their scientific and technological skills, in order to address the problems faced by their communities.

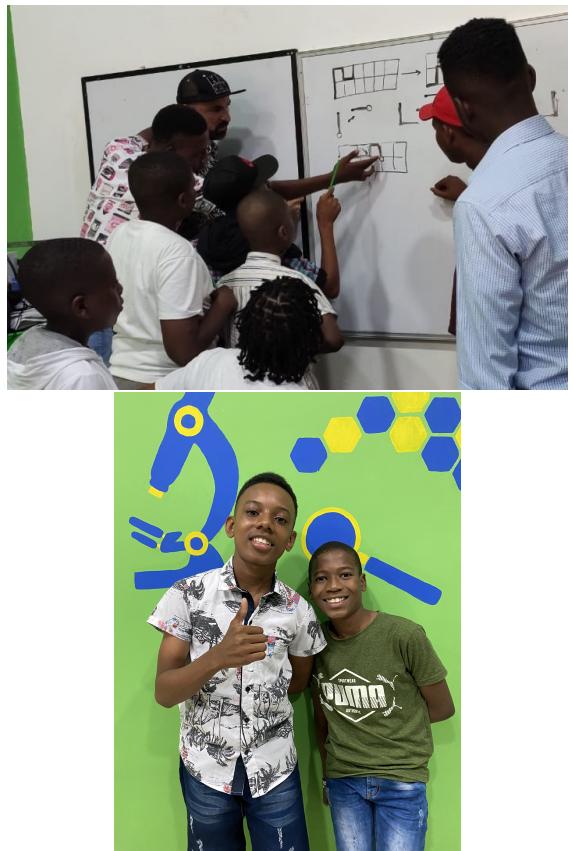


Figure 18: a. At the Escuela de Robótica del Chocó. b. Deison Rivas and Juan David Cuenta.

At the end of the workshop, we asked the students: What robot do you **really** want to design? Deison Rivas wants to build a firefighter robot; it will quickly and safely go in and out of houses – traditionally made of wood – and put out the fires that have razed entire city blocks in Quibdó in the past. Juan David Cuenta wants to design an agile rescue robot; it will help people stuck under the frequent landslides caused by illegal mining operations and by heavy rainfalls on the roads.

This theoretical exercise on robotic optimization immediately took on new meaning, thanks to the wisdom of these young people.

17 Our role as educators

As I put the finishing touches on this paper, I remember a book I inherited from my mom when she passed away ten years ago. Browsing through it I find this passage, highlighted by her:

The saddest achievement of our educational system is to produce [...] the most disastrous person of our society: the creative scientist who is at the same time enslaved to the military or industrial apparatus; someone who makes contributions, but has no interest in the way they will be used.

– Estanislao Zuleta [?]

18 What does it mean to do math ethically?

Six years ago, my student Brian Cruz asked me whether mathematicians have an ethical code, similar to the Hippocratic Oaths adopted by physicians. More than two decades into my mathematical career, I had never thought or heard of this specific suggestion.

Thanks to Brian, I did some research, gathered some resources with the help of my students³, and I now devote one day of each semester to discuss this question with them. Posing the question to them is surely more important than proposing an answer:

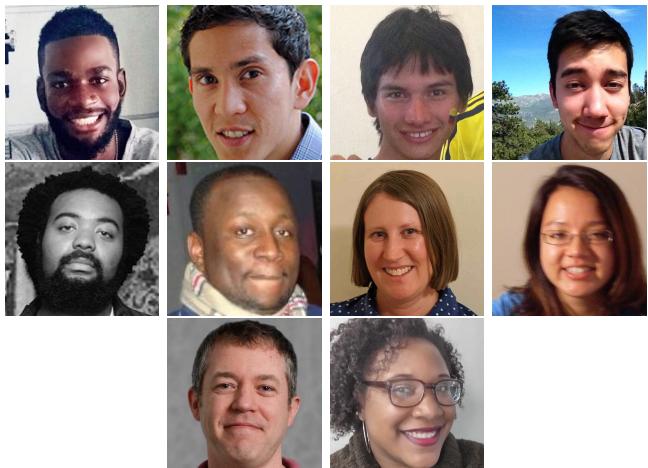
Writing assignment. What does “doing mathematics ethically” mean to you? This question is an invitation to recognize the power you carry as a mathematician, and the privilege and responsibility that comes with it. When you enter a scientific career, you do not leave yourself at the door. You can choose how to use that power. My

hope is that you will always continue to think about this in your work.

19 Acknowledgments

Este artículo está dedicado a mi mamá, Amparo Mantilla. Una de sus incontables enseñanzas es la importancia de pensar críticamente en el impacto de la ciencia en la sociedad.

I would like to extend my sincere gratitude to the students and colleagues who have collaborated with me on this research: Arlys Asprilla, César Ceballos, Hanner Bastidas, John Guo, Matthew Bland, Maxime Pouokam, Megan Owen, Rika Yatchak, Seth Sullivant, and Tia Baker.



I am also very grateful to the many people who encouraged me and helped me figure out how to tell this story.

³These resources are available at <http://math.sfsu.edu/federico/ethicsinmath.html>