AFFINE TRANSFORMATIONS OF THE PLANE

1. Introduction

We are interested here in the geometry of the two-dimensional plane, $\mathbb{R}^2$. There are many concepts commonly thought of as "geometric", e.g., lines, circles, distances, angles, areas, etc. In 1872, Felix Klein proposed his *Erlangen program*, an organizing framework for different geometric notions:

- There is not a single "geometry", but rather multiple different "geometries", depending on which basic notions one is allowed to talk about.
- Each geometry is characterized by its allowed "symmetries". For example, the symmetries of Euclidean geometry include rotation, reflection, and translation.
- The meaningful notions in a particular geometry should be precisely those which are preserved by all the symmetries. For example, the symmetries of Euclidean geometry preserve lines and angles, but do not preserve, say, the notion of *vertical* line (which can be rotated), which is therefore not a meaningful notion in Euclidean geometry.

Of course, 1872 was before the advent of mathematical logic. In modern terminology, the various possible planar "geometries" are given by different structures (first-order or otherwise) one can put on the underlying set $\mathbb{R}^2$; while the "symmetries" of each "geometry" are the automorphisms of that structure. The last bullet point above is then an instance of the general fact that the notions definable in a structure are precisely those preserved by all automorphisms (at least if one is willing to use higher-order or infinitary logic; see Exercise 2.84 in the notes on first-order logic).

2. Affine geometry

Roughly speaking, affine geometry is the geometry one gets by regarding lines as the only primitive notion. This is a more primitive geometry than Euclidean geometry: for example, scaling by a constant factor is an affine automorphism, which shows that "distance" is not an affine notion.

We can encode "lines" into a first-order structure by talking about collinearity of three points. Define the ternary relation on $\mathbb{R}^2$

$$\mathrm{Coll}(A, B, C) \;:\Longleftrightarrow\; A, B, C \text{ are collinear.}$$

Note that here, each of A, B, C is an ordered pair of real numbers. Note also that we do not necessarily require A, B, C to be distinct; if two of the points (or even all three) are equal, then they are trivially collinear. Note, finally, that we do not say anything about the *order* in which A, B, C appear on a line; for example, the following triples (A, B, C) both satisfy Coll:

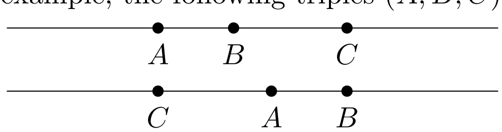

The following more quantitative notion allows us to talk about not just the ordering of points on a line, but even the precise position where a point appears on a line through two other points. An **affine combination** of two points $A = (x_A, y_A)$ and $B = (x_B, y_B)$ is a point which can be written as a weighted average of them:

$$C = (1 - t)A + tB, \quad \text{for some } t \in \mathbb{R}.$$

Here, $tB := (tx_B, ty_B)$ denotes the coordinatewise multiple, and similarly for $(1-t)A$; while $+$ denotes coordinatewise addition. For example, when $t = 0$ we get A, while $t = 1$ yields B; as t varies from 0 to 1, we move "at constant speed" from A to B.

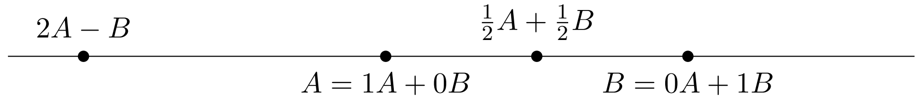

Note that we also allow t outside the interval $[0, 1]$, in which case we get points on the line AB but not the line segment; e.g., when $t = -1$, we get $2A - B$, which can be written more intelligibly as $A + (A - B)$ (start at A, then shift by the vector from B to A). Note, finally, that we again do not require $A \neq B$; if $A = B$, then all affine combinations of them are the same.

In order to express affine combinations in first-order logic, we introduce, for each $t \in \mathbb{R}$, the binary operation

$$\begin{aligned} \mathrm{AC}_t : (\mathbb{R}^2)^2 &\longrightarrow \mathbb{R}^2 \\ (A, B) &\longmapsto (1-t)A + tB. \end{aligned}$$

Thus we get an uncountable signature $\mathcal{A}_{\mathsf{aff}}$, consisting of all the binary function symbols AC_t for each t (similarly to the example $\mathcal{A}_{\mathsf{vec}}$ of vector spaces), together with a natural $\mathcal{A}_{\mathsf{aff}}$-structure on $\mathbb{R}$.

It is intuitively clear that the $\mathcal{A}_{\mathsf{aff}}$-structure on $\mathbb{R}$ is richer than merely the $\{\mathrm{Coll}\}$-structure: if we can say "C appears on the line AB at position t", for each t, then we can say that C appears at *some* position along the line. Formally, this "some" would be an infinite disjunction over all $t \in \mathbb{R}$. We also have to be slightly careful about the degenerate case where $A = B$; as noted above, in this case their only affine combination will be the same point $A = B$, even though any C will be trivially collinear with A, B, C (we can take the line AC instead). In summary, the infinitary formula

$$(A = B) \vee \bigvee_{t \in \mathbb{R}} (\mathrm{AC}_t(A, B) = C) \in \mathcal{L}_{\mathrm{form}}^{\{A,B,C\}}(\mathcal{A}_{\mathsf{aff}}) \tag{2.1}$$

shows that the ternary relation Coll is positive-existential definable from the $\mathcal{A}_{\mathsf{aff}}$-structure on $\mathbb{R}^2$.

Exercise 2.2. More generally, an **affine combination** of any finite number of points $A_1, \dots, A_n \in \mathbb{R}^2$ is a point of the form

$$a_1 A_1 + \cdots + a_n A_n, \qquad \text{for some } a_1, \dots, a_n \in \mathbb{R} \text{ adding to } 1.$$

For example, when $n = 3$ and $a_1 = a_2 = a_3 = \frac{1}{3}$, we get the **centroid** or **barycenter** of a triangle:

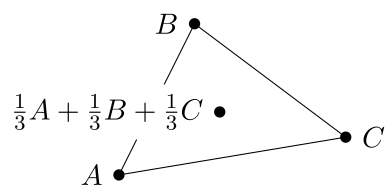

Show that for each n and $a_1, \dots, a_n$ adding to 1, the n-ary operation

$$\begin{aligned} \mathrm{AC}_{a_1,\dots,a_n} : (\mathbb{R}^2)^n &\longrightarrow \mathbb{R}^2 \\ (A_1, \dots, A_n) &\longmapsto a_1 A_1 + \cdots + a_n A_n \end{aligned}$$

is definable by a *term* from the $\mathcal{A}_{\mathsf{aff}}$-structure on $\mathbb{R}^2$. [Hint: when $n \geq 2$, not all the a_i can be 1.]

Exercise 2.3. Show that for any three points $A, B, C \in \mathbb{R}^2$, the following are equivalent:

(a) every point can be written as an affine combination of A, B, C in at least one way;
(b) every point can be written as an affine combination of A, B, C in at most one way;
(c) A, B, C are not collinear, i.e., none of them is an affine combination of the other two.

If these hold, we say A, B, C are an **affine basis** for $\mathbb{R}^2$.

3. Affine transformations

An $\mathcal{A}_{\mathsf{aff}}$-homomorphism $T : \mathbb{R}^2 \to \mathbb{R}^2$ is called an **affine transformation**. It follows that an affine transformation preserves all operations or relations positive-existential definable from the operations AC_t, possibly in infinitary logic. This includes the collinearity relation (by (2.1)), as well as general n-ary affine combinations as in Exercise 2.2.

Exercise 3.1. Using the affine basis $(0,0),(1,0),(0,1)$ for $\mathbb{R}^2$, show that affine transformations of the plane are precisely all functions of the form

$$\begin{aligned} T : \mathbb{R}^2 &\longrightarrow \mathbb{R}^2 \\ (x,y) &\longmapsto (ax+by+c, dx+ey+f) \end{aligned}$$

where $a,b,c,d,e,f \in \mathbb{R}$ are arbitrary. Moreover, these parameters a,b,c,d,e,f are uniquely determined by T, so that we have a bijection between affine transformations and $\mathbb{R}^6$.

[These 6 parameters a,b,c,d,e,f are commonly called the **matrix** of T, written as

$$\begin{pmatrix} a & b & c \\ d & e & f \end{pmatrix}.$$

The left 2×2 submatrix gives a linear transformation $\mathbb{R}^2 \to \mathbb{R}^2$, called the **linear part** of T; T is this linear part followed by the translation $(x,y) \mapsto (x,y)+(c,f)$.]

A $\{\mathrm{Coll}\}$-homomorphism $T : \mathbb{R}^2 \to \mathbb{R}^2$ is called a **collineation**. Thus, all affine transformations are collineations. Since Coll only allows us to express "being on the same line", but doesn't say anything quantitative about positions along a line, we might intuitively expect collineations to be much "looser" than affine transformations; for example, *a priori*, a collineation could arbitrarily permute the order of points on a line. This intuition is partially correct:

Exercise 3.2.
(a) Give an example of a collineation $T : \mathbb{R}^2 \to \mathbb{R}^2$ which is not affine.
(b) Conclude that affine combinations are not positive-existential definable from the collinearity relation (even using infinitary logic).
(c) (for those who know some set theory) Show that the cardinality of the set of collineations $T : \mathbb{R}^2 \to \mathbb{R}^2$ is strictly bigger than that of the set of affine transformations.

By an **invertible affine transformation**, respectively, **invertible collineation**, we mean an automorphism of $\mathbb{R}^2$ equipped with the respective structure. Note that since $\mathcal{A}_{\mathsf{aff}}$ consists only of function symbols, an invertible affine transformation is the same thing as a bijective one. The same happens to be true for collineations, even though Coll is a relation:

Exercise 3.3.
(a) Show that $A,B,C \in \mathbb{R}^2$ are *not* collinear iff they are distinct, and every point in $\mathbb{R}^2$ lies on a line through two distinct points each of which lies on one of the lines AB, BC, CA.
(b) Conclude that the inverse of every bijective collineation is a collineation. [See HW6 Q4.]

Now comes the surprise:

Theorem 3.4. Every invertible collineation $T : \mathbb{R}^2 \to \mathbb{R}^2$ is affine.

In other words, keeping in mind the correspondence between automorphisms and definability (in infinitary logic), this says that the purely qualitative notion of collinearity (the relation Coll) can in fact be used to define positions along a line (the operations AC_t)! Indeed, we will prove this theorem by giving such an explicit definition of AC_t from Coll. This will be done via a series of lemmas, each showing that progressively more "quantitative" notions can be defined from Coll.

Lemma 3.5. The quaternary relation

$$\mathrm{Para}(A, B, C, D) \;:\Longleftrightarrow\; A \neq B \text{ and } C \neq D \text{ and the lines } AB \text{ and } CD \text{ are parallel}$$

can be defined from Coll.

Proof. Two lines in the plane are parallel iff they don't intersect, or they are the same line; thus Para is defined by the {Coll}-formula

$$\neg(A = B) \wedge \neg(C = D) \wedge (\neg\exists E\,(\mathrm{Coll}(A, B, E) \wedge \mathrm{Coll}(C, D, E)) \vee (\mathrm{Coll}(A, B, C) \wedge \mathrm{Coll}(A, B, D))). \square$$

Lemma 3.6. The quaternary relation

$$\mathrm{Pgram}(A, B, C, D) \;:\Longleftrightarrow\; ABCD \text{ is a non-collinear parallelogram}$$

can be defined from Coll and Para, hence from just Coll.

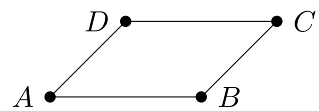

Proof. We basically just need to say that opposite sides are parallel; for non-collinearity, it is enough to say that A, B, C are not collinear (which will in particular force all four vertices to be distinct, in agreement with the first two clauses in the definition of Para above):

$$\neg\mathrm{Coll}(A, B, C) \wedge \mathrm{Para}(A, B, C, D) \wedge \mathrm{Para}(A, D, B, C). \qquad \square$$

At this stage, we have already extracted a fair amount of quantitative information from the Coll relation. Indeed, note that parallelograms essentially allow us to express vector addition:

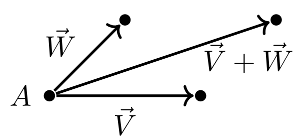

Informally speaking, all that is still needed in order to express arbitrary affine combinations is scalar multiplication of vectors. Scaling by an integer amount can be expressed via repeated addition; scaling by a rational a/b can then be expressed by saying that scaling one vector by a is equal to scaling another by b. Of course, there is the technical annoyance that we need to add *parallel* vectors here, while above we restricted to *non-collinear* parallelograms. This is easily worked around:

Lemma 3.7. The binary operation

$$\begin{aligned} \mathrm{AC}_2 : (\mathbb{R}^2)^2 &\longrightarrow \mathbb{R}^2 \\ (A, B) &\longmapsto -A + 2B = B + (B - A) \end{aligned}$$

can be defined from Pgram, hence from Coll.

Proof. Recall that to say that the operation AC_2 is definable means that its ternary graph relation "$C = B + (B - A)$" is definable. Indeed, we have

$$C = B + (B - A) \iff \exists D\,\exists E\,(\mathrm{Pgram}(A, B, E, D) \wedge \mathrm{Pgram}(B, C, E, D)).$$

This expresses the following situation:

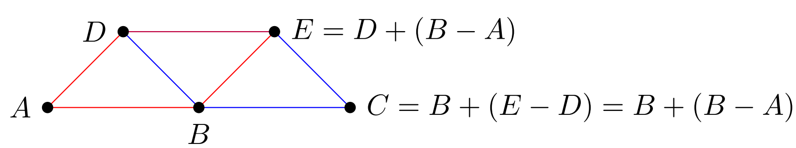

$\square$

Lemma 3.8. For any rational $t \in \mathbb{Q}$, the binary operation $\mathrm{AC}_t : (\mathbb{R}^2)^2 \to \mathbb{R}^2$ is definable from Coll.

Proof. We split the proof into several stages, for increasingly general t.

First, consider the case $t = n \in \mathbb{N}$. We use induction on n. For $n = 0$ or 1, we have

$$\mathrm{AC}_0(A, B) = A, \qquad\qquad \mathrm{AC}_1(A, B) = B.$$

For $n \geq 2$, by considering the picture

we are led to the calculation

$$\begin{aligned}\mathrm{AC}_n(A, B) &= (1 - n)A + nB \\ &= (2 - n)A + (n - 1)B + ((2 - n)A + (n - 1)B - (3 - n)A - (n - 2)B) \\ &= \mathrm{AC}_2(\mathrm{AC}_{n-2}(A, B), \mathrm{AC}_{n-1}(A, B));\end{aligned}$$

since AC_2 is definable from Coll by the preceding lemma, while $\mathrm{AC}_{n-1}, \mathrm{AC}_{n-2}$ are definable by the IH, we get that AC_n is definable from Coll. This proves the case $t = n \in \mathbb{N}$.

For negative $t \in \mathbb{Z}$, we have

$$\mathrm{AC}_t(A, B) = (1 - t)A + tB = \mathrm{AC}_{1-t}(B, A);$$

since $t < 0$, $1 - t \geq 0$, so this reduces to the previous case. So we have proved all integer cases $t \in \mathbb{Z}$.

Finally, for a rational $t = a/b$ where $a, b \in \mathbb{Z}$, we have

$$\begin{aligned}\mathrm{AC}_t(A, B) = C &\iff (1 - \tfrac{a}{b})A + \tfrac{a}{b}B = C \\ &\iff \tfrac{a}{b}(B - A) = C - A \\ &\iff a(B - A) = b(C - A) \\ &\iff (1 - a)A + aB = (1 - b)A + bC \\ &\iff \mathrm{AC}_a(A, B) = \mathrm{AC}_b(A, C);\end{aligned}$$

since $\mathrm{AC}_a, \mathrm{AC}_b$ are definable from Coll by the previous cases, so is AC_t. $\square$

In order to complete the proof, we just need to extend from rational t to arbitrary real $t \in \mathbb{R}$. This turns out to be *much* harder than any of the preceding lemmas. Intuitively, the idea is that we want to "approximate" AC_t for an arbitrary real t by AC_r for rational $r \approx t$, and then take a limit as $r \to t$. We thus need to define "limit" from Coll. We can formulate this as follows:

Lemma 3.9. The ternary relation

$$\begin{aligned}\mathrm{Between}(A, B, C) \;&:\iff B \text{ is on the line } \textit{segment } AC \\ &\iff B = \mathrm{AC}_t(A, C) \text{ for some } t \in [0, 1]\end{aligned}$$

is definable from Coll. [Such affine combinations with $t \in [0, 1]$ are called **convex combinations**.]

Once we know this lemma, we can approximate $t \in \mathbb{R}$ by rationals $r < t < s$ on either side of it, and then demand that $\mathrm{AC}_t(A, B)$ be between $\mathrm{AC}_r(A, B)$ and $\mathrm{AC}_s(A, B)$ for all such approximations; by taking r, s to be closer and closer to t, a "squeezing" argument then shows that $\mathrm{AC}_t(A, B)$ must be what it should be. Assuming Lemma 3.9 for now, we can thus complete the

Proof of Theorem 3.4. We claim that for any $t \in \mathbb{R}$,

$$(*) \qquad \mathrm{AC}_t(A, B) = C \iff \textstyle\bigwedge_{r,s \in \mathbb{Q}; r < t < s} \mathrm{Between}(\mathrm{AC}_r(A, B), C, \mathrm{AC}_s(A, B)).$$

This is enough, since $\mathrm{AC}_r, \mathrm{AC}_s$ are definable (in first-order logic) from Coll by Lemma 3.8, and Between is definable by Lemma 3.9, whence AC_t is definable (in infinitary logic) from Coll, hence preserved by all {Coll}-automorphisms (by Exercise 2.62(c) in the notes on first-order logic).

To see (∗): if $\mathrm{AC}_t(A,B) = C$, then for all $r < t < s$, we have

$$\begin{aligned} C &= (1-t)A + tB \\ &= A + t(B-A) \\ &= A + (r + \tfrac{t-r}{s-r}(s-r))(B-A) \\ &= A + ((1-\tfrac{t-r}{s-r})r + \tfrac{t-r}{s-r}s)(B-A) \\ &= (1-\tfrac{t-r}{s-r})(A + r(B-A)) + \tfrac{t-r}{s-r}(A + s(B-A)) \\ &= (1-\tfrac{t-r}{s-r})((1-r)A + rB) + \tfrac{t-r}{s-r}((1-s)A + sB) \\ &= \mathrm{AC}_{\frac{t-r}{s-r}}(\mathrm{AC}_r(A,B), \mathrm{AC}_s(A,B)) \end{aligned}$$

where $\frac{t-r}{s-r} \in [0,1]$ because $r < t < s$, whence C is a convex combination of $\mathrm{AC}_r(A,B), \mathrm{AC}_s(A,B)$. Conversely, suppose that for all $r, s \in \mathbb{Q}$ with $r < t < s$, C is between $\mathrm{AC}_r(A,B), \mathrm{AC}_s(A,B)$. Since

$$\begin{aligned} \mathrm{AC}_s(A,B) - \mathrm{AC}_r(A,B) &= ((1-s)A + sB) - ((1-r)A + rB) \\ &= (s-r)(B-A), \end{aligned}$$

the distance between $\mathrm{AC}_r(A,B), \mathrm{AC}_s(A,B)$ is $s-r$ times the distance between A, B. Since C is between $\mathrm{AC}_r(A,B), \mathrm{AC}_s(A,B)$, as is $\mathrm{AC}_t(A,B)$ (as shown above), the distance between $C, \mathrm{AC}_t(A,B)$ is thus at most $s-r$ times the distance between A, B. Since we may choose r, s so that $s-r$ is arbitrarily small, the distance between $C, \mathrm{AC}_t(A,B)$ must thus be zero. □

So it remains to prove Lemma 3.9. We can make one more easy reduction:

Lemma 3.10. The ternary relation

$$\begin{aligned} \mathrm{Ray}(A,B,C) &:\iff B \text{ is on the ray } AC \\ &\iff B = \mathrm{AC}_t(A,C) \text{ for some } t \in [0,\infty) \end{aligned}$$

is definable from Coll.

Proof of Lemma 3.9. $\mathrm{Between}(A,B,C) \iff \mathrm{Ray}(A,B,C) \wedge \mathrm{Ray}(C,B,A)$. □

Proof of Lemma 3.10. If $A = C$, then $B = \mathrm{AC}_t(A,C)$ again just means $A = B = C$. So we may restrict attention to the case $A \neq C$. Consider the following picture:

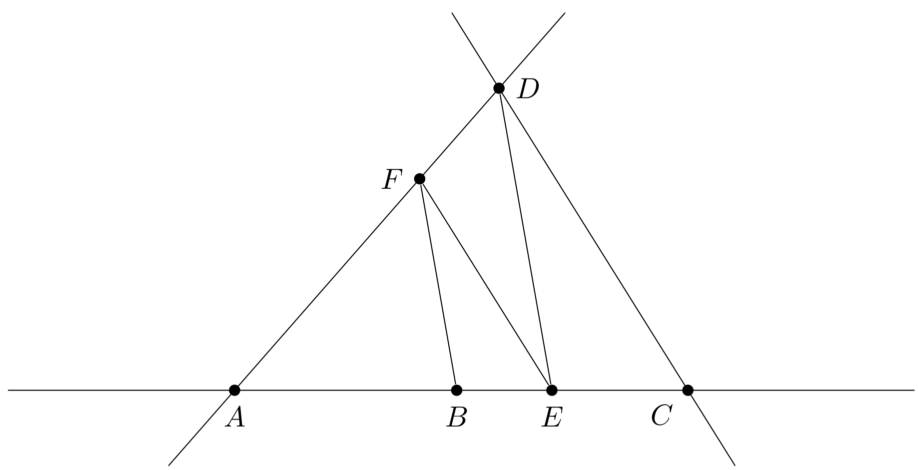

Here D is any point not on the line AC, while E is any point on the line AC, so that

$$E = (1-t)A + tC$$

for some $t \in \mathbb{R}$. The unique line through E parallel to CD intersects the line AD at a unique point F, since CD is not parallel to AD (since A, C, D are not collinear). We must have

$$F = (1-t)A + tD,$$

in order to ensure that the vector $F - E = t(D - C)$ is parallel to the line CD. Now similarly to before, the unique line through F parallel to DE intersects the line AC at a unique point B, since DE is not parallel to AC (since these two lines intersect at a unique point E); and for the same reason as for F, we must have

$$\begin{aligned} B &= (1-t)A + tE \\ &= (1-t)A + t((1-t)A + tC) \\ &= (1+t)(1-t)A + t^2 C \\ &= (1-t^2)A + t^2 C. \end{aligned}$$

The key thing to note is that the coefficient t^2 must be ≥ 0. In other words, what we have shown is that, starting from any point E on the *line* AC, and then constructing F and B uniquely as above, we end up with a point B on the *ray* AC; and this point B may be an arbitrary point $B = (1-s)A + sC$ on the ray AC, for any $s \geq 0$, since we may choose E above with $t := \sqrt{s}$. Putting everything together, we may define Ray from Coll and Para as follows:

$$\mathrm{Ray}(A, B, C) \iff (A = B) \vee \left(\neg(A = C) \wedge \exists D\, \exists E\, \exists F \begin{pmatrix} \neg\mathrm{Coll}(A, C, D) \wedge \mathrm{Coll}(A, C, E) \wedge \\ \mathrm{Para}(C, D, E, F) \wedge \mathrm{Coll}(A, D, F) \wedge \\ \mathrm{Para}(D, E, F, B) \wedge \mathrm{Coll}(A, B, C) \end{pmatrix} \right).$$

(The clause $A = B$, aside from covering the case $A = C$, also takes care of the case $E = A$, in which case $E = F = B = A$ and so the relations $\mathrm{Para}(C, D, E, F)$ and $\mathrm{Para}(D, E, F, B)$ fail to describe the above construction of F and B.) $\square$

Exercise 3.11. Draw the above picture starting with E on the other side of A, to see that B still ends up on the right side of A.

We started out trying to prove Theorem 3.4 by pointing out that it amounts to defining quantitative positions on a line from the seemingly non-quantitative notion of collinearity. This is indeed what all the formulas in the proofs of the preceding lemmas achieve. However, the details of how we achieved this are a bit surprising: as we noted above, the relatively simple proofs of Lemmas 3.5 to 3.8 already extracted most of the quantitative information from Coll; whereas the *hardest* step of passing from rational to real t is embodied by the still seemingly non-quantitative Lemma 3.10, where we had to work quite hard in order to bridge the trivial-looking gap between *lines* and *rays*.

In fact, there is a precise sense in which Lemma 3.10 really is the single hardest part of the proof, and not just because the proof we happened to give was so complicated: while the rest of the proof also holds for affine geometry over planes $\mathbb{F}^2$ with coordinates from more general fields $\mathbb{F}$, Lemma 3.10 depends crucially on the fact that the only field automorphism of $\mathbb{R}$ is the identity (which we exploited via the closely related fact that positive reals have square roots).

Exercise 3.12. Let $\mathbb{F}$ be an arbitrary field, and define collinearity of points in $\mathbb{F}^2$ via affine combinations as in $\mathbb{R}^2$, but taking coefficients from $\mathbb{F}$.

(a) Show that for any field automorphism $h : \mathbb{F} \cong \mathbb{F}$, the function $T : \mathbb{F}^2 \to \mathbb{F}^2$ defined by applying h in each coordinate (i.e., $T(x, y) := (h(x), h(y))$) is an invertible collineation, but is not an affine transformation unless $h = \mathrm{id}_{\mathbb{F}}$.

(b) Conclude that $\mathbb{Q}[\sqrt{2}]^2 \subseteq \mathbb{R}^2$ (recall HW7) has a non-affine invertible collineation, even though neither $\mathbb{R}^2$ nor $\mathbb{Q}^2$ does.