PUBLICATIONS MATHÉMATIQUES DE BESANÇON

Algèbre et Théorie des nombres

Calum Spicer

Recent progress on the Minimal Model Program for foliations

2026, p. 109-128.

https://doi.org/10.5802/pmb.73

PRESSES UNIVERSITAIRES DE FRANCHE-COMTÉ

Les Publications mathématiques de Besançon *sont membres du Centre Mersenne pour l'édition scientifique ouverte*
http://www.centre-mersenne.org/
e-ISSN : 2592-6616

Publications mathématiques de Besançon – 2026, 109-128

RECENT PROGRESS ON THE MINIMAL MODEL PROGRAM FOR FOLIATIONS

by

Calum Spicer

Abstract. — We recall the work of Brunella, McQuillan and Mendes on the classification of foliations on projective surfaces. We then survey some recent work extending this classification to foliations on higher dimensional projective varieties.

Résumé. — Nous rappelons les travaux de Brunella, McQuillan et Mendes sur la classification des feuilletages sur les surfaces projectives. Nous présentons ensuite un aperçu de travaux récents qui étendent cette classification aux feuilletages sur les variétés projectives de dimension ≥ 3.

1. Introduction

These are the notes for a lecture series given at Luminy in February 2025. The aim of the lectures was to explain an ongoing program of research into the classification of foliations on projective varieties. The rough goal of this project is the following.

Group all foliations on projective varieties into families depending on natural invariants.

Our approach to this problem will be guided by the ideas and insights of the Minimal Model Program (MMP). There are many wonderful surveys on the MMP, for instance [23], but in these notes we will assume that the reader has no prior familiarity with the MMP.
We will work over $\mathbb{C}$ throughout.

Acknowledgments. — I thank the organisers for the invitation to speak at this conference and for the suggestion to write up these notes. I thank my collaborators, especially Paolo Cascini, Jihao Liu, Michael McQuillan, Jorge Vitório Pereira and Roberto Svaldi for many stimulating discussions over the years on the material covered in these notes. Finally, my thanks to the referee for many thoughtful and constructive comments. The work described here was partially supported by EPSRC.

2020 ***Mathematics Subject Classification.*** — 37F75.

2. The birational classification of surface foliations

2.1. Some preliminary definitions. — A morphism $f\colon X \to Z$ between normal varieties is said to be a *contraction* if it is surjective and $f_*\mathcal{O}_X = \mathcal{O}_Z$.
Let $\varphi\colon X \dashrightarrow Y$ be a birational map between normal algebraic varieties. The *exceptional locus* of φ is the closed subset $\operatorname{Exc}\varphi$ of X where φ is not an isomorphism. We say that φ is a *birational contraction* if $\operatorname{Exc}\varphi^{-1}$ does not contain any divisor. Let $\varphi\colon X \dashrightarrow Y$ be a birational contraction between normal varieties and let D be a $\mathbb{Q}$-Cartier $\mathbb{Q}$-divisor on X such that $D_Y := \varphi_* D$ is also $\mathbb{Q}$-Cartier. Then φ is called *D-negative* (resp. *D-trivial*) if there exist a normal variety W and birational morphisms $p\colon W \to X$ and $q\colon W \to Y$ resolving the indeterminacy locus of φ such that $p^*D = q^*D_Y + E$ where $E \geq 0$ is a $\mathbb{Q}$-divisor whose support coincides with the exceptional locus of q (resp. $p^*D = q^*D_Y$).
A foliation, denoted $\mathcal{F}$, of rank r on a normal algebraic variety X is a rank r coherent subsheaf $T_{\mathcal{F}} \subset T_X$ such that

(1) $T_{\mathcal{F}}$ is saturated in T_X, and

(2) $T_{\mathcal{F}}$ is closed under Lie bracket.

Note that if $r = 1$ then (2) is automatically satisfied. We call $T_{\mathcal{F}}$ the tangent sheaf of $\mathcal{F}$.
If X is a smooth algebraic variety and $\mathcal{F}$ is a foliation on X we define $\operatorname{Sing}\mathcal{F}$ to be the set of points where $T_{\mathcal{F}}$ is not a sub-bundle of T_X (we refer to [19, Definition 5.4] for a more refined notion of singular locus which behaves better on singular varieties).

2.2. Kodaira dimension and numerical dimension. — To make sense of our main goal, let us first recall the work of Brunella, McQuillan and Mendes on the classification of surface foliations, [7, 25, 28]. The insight of Brunella, McQuillan and Mendes was that there should be a Kodaira–Enriques type classification for surface foliations. We will explain the framework for this classification now.
To any normal projective surface X equipped with a foliation $\mathcal{F}$ there is a canonical invariant we can associate to $\mathcal{F}$, namely, the canonical class $K_{\mathcal{F}}$. We define $K_{\mathcal{F}}$ to be a Weil divisor such that $\mathcal{O}(-K_{\mathcal{F}}) = \det T_{\mathcal{F}}$. From $K_{\mathcal{F}}$ we can extract some numerical invariants:

(1) the plurigenera $h^0(X, \mathcal{O}(mK_{\mathcal{F}}))$ and the *Kodaira dimension*
$$\kappa(X,\mathcal{F}) = \max\{\dim \phi_{|mK_{\mathcal{F}}|}(X) : m = 1, 2, 3, \dots\};$$
and

(2) the *numerical dimension*
$$\nu(X,\mathcal{F}) = \begin{cases} \max\{k : P^k \not\equiv 0\} & \text{if } K_{\mathcal{F}} \text{ is pseudo-effective and} \\ & \qquad \text{has Zariski decomposition } K_{\mathcal{F}} = P + N \\ -\infty & \text{if } K_{\mathcal{F}} \text{ is not pseudo-effective.} \end{cases}$$

The idea behind the classification of Brunella, McQuillan and Mendes is to classify foliations according to $\nu(X,\mathcal{F})$ and $\kappa(X,\mathcal{F})$.
We first observe that these invariants are birational invariants (at least for foliations with mild singularities, see Exercise 3.3 for a precise statement) and therefore our classification should

itself be birational in nature. Let us take a moment to explain why κ and ν are birational invariants in a special case.
Let X be a smooth projective surface, let $\mathcal{F}$ be a rank one foliation on X and let $x \in X\backslash \mathrm{Sing}\,\mathcal{F}$. Let $b\colon X' \to X$ be the blow up of X at x. Let $E = b^{-1}(x)$ be the exceptional divisor. Note that there is a uniquely determined transform of $\mathcal{F}$ as a foliation on X', which we will denote $b^{-1}\mathcal{F}$. The construction of $\mathcal{F}'$ is as follows: by considering $\mathcal{F}$ as a foliation on $X \setminus \{x\}$ we get a foliation on $X' \setminus E$. Observe that a coherent saturated subsheaf of $T_{X'}|_{X'\setminus E}$ admits a unique extension to a coherent saturated subsheaf of $T_{X'}$. This unique extension gives us $b^{-1}\mathcal{F}$.
Having defined $b^{-1}\mathcal{F}$, we can then ask how the canonical divisor transforms under blow ups.

Exercise 2.1. — Use notation as above.

(1) Show that $K_{b^{-1}\mathcal{F}} = b^*K_{\mathcal{F}} + E$.

(2) Show that $h^0(X, \mathcal{O}(mK_{\mathcal{F}})) = h^0(X', \mathcal{O}(mK_{b^{-1}\mathcal{F}}))$ for $m \geq 0$.

(3) Show that $\nu(X, \mathcal{F}) = \nu(X', b^{-1}\mathcal{F})$.

Exercise 2.2. — More generally, given a normal variety X, a foliation $\mathcal{F}$ on X and any birational map $X \dashrightarrow X'$ define the transform of $\mathcal{F}$ on X'.

While $(X, \mathcal{F})$ and $(X', b^{-1}\mathcal{F})$ have the same set of canonically determined invariants, it is intuitively clear that $(X, \mathcal{F})$ is "simpler" from the perspective of classification. This suggests a natural approach to our classification.

- First, we should try to find a member $(X_{\min}, \mathcal{F}_{\min})$ in each birational equivalence class of foliated varieties which is minimal in the sense that every other member of the birational class arises as a blow up of $(X_{\min}, \mathcal{F}_{\min})$. (This is not always possible, as we will see, but we will describe a process which gets us as close as possible to this situation.)

- Then, we should try to describe the possible geometries of $(X_{\min}, \mathcal{F}_{\min})$ in terms of the Kodaira and numerical dimension.

The majority of what will be discussed in this survey relates to the first point in the above scheme, namely, finding a simplest model in each birational equivalence class. Indeed, as we will see, this part of our strategy will have many implications for the study of foliations outside of its role in the classification scheme. Unfortunately, we will not say much about the second aspect of the classification (but see Section 6.2 for a brief summary).
It must also be emphasised that the role of singularities is very important in this classification scheme, and we will discuss it in greater detail in the coming sections. For the purposes of understanding the birational theory of surface foliations the key class of singularities is that of simple singularities, which we will now recall (in Section 3 we will expand this class of singularities to the class of log canonical singularities).
Let $x \in X$ be a germ of a smooth surface and let $\mathcal{F}$ be a rank one foliation on X defined by a vector field ∂. We say that $x \in X$ is a simple singularity for $\mathcal{F}$ provided one of the following holds:

- ∂ is non-singular at x; or

– ∂ is singular at x and if λ, μ are the eigenvalues of the linear part of ∂ at x, then at least one (say μ) is non-zero and $\frac{\lambda}{\mu} \notin \mathbb{Q}_{>0}$.

We note that in some places in the literature simple singularities are called reduced singularities.

2.3. Alternative classification schemes. — It is worth taking a moment to consider how the classification proposed above relates to another very natural way to categorise foliations on surfaces. Let X be a smooth projective surface together with a foliation $\mathcal{F}$. Let $x \in X \setminus \operatorname{Sing} \mathcal{F}$ be a point and let L_x be the leaf of $\mathcal{F}$ through x. Since L_x is a Riemann surface its universal cover $\widetilde{L_x}$ is either $\mathbb{P}^1, \mathbb{C}$ or $\mathbb{D} \subset \mathbb{C}$. For general $x \in X$, $\widetilde{L_x}$ is independent of the choice of x (this follows from the continuity of the leafwise Poincaré metric away from an algebraic subset of $Z \subset X$, see for instance [25, Corollary V.1.4]) and so we can group surface foliations into three classes based on the universal cover of a general leaf.
To compare these two groupings consider the following table. We assume that X is a smooth projective surface and $\mathcal{F}$ is a foliation on X with simple singularities.

$\nu(K_\mathcal{F})$	**Universal cover of L_x**
$-\infty$	$\mathbb{P}^1$
0	$\mathbb{C}$
1	$\mathbb{C}$ or $\mathbb{D}$
2	$\mathbb{D}$

I should emphasise that the above table is not obvious at all, but can be deduced from the Kodaira–Enriques style classification of foliations which is summarised in Section 6.2 below. An important takeaway from this table is that grouping foliations in terms of canonically determined invariants $\nu(X, \mathcal{F})$ and $\kappa(X, \mathcal{F})$ also groups foliations into natural families from a geometric perspective.

2.4. Finding a minimal representative. — We now return to the problem of finding a "minimal" representative in any birational class.
The fundamental result in this direction is the following.

Theorem 2.3. — *Let X be a smooth projective surface and let $\mathcal{F}$ be a rank one foliation on X with simple singularities. Then there is a $K_\mathcal{F}$-negative birational contraction $f\colon X \to Y$ such that one of two (mutually exclusive) possibilities occur where $\mathcal{F}_Y$ is the transform of $\mathcal{F}$ on Y.*

(1) *$K_{\mathcal{F}_Y}$ is nef, i.e., $K_{\mathcal{F}_Y} \cdot C \geq 0$ for any curve $C \subset Y$; or*

(2) *there is a $\mathbb{P}^1$-fibration $Y \to Z$ such that $T_{\mathcal{F}_Y} = T_{Y/Z}$ and $\rho(Y/Z) = 1$.*

Before explaining how Theorem 2.3 is proven, let us make a few comments.
Intuitively, the morphism has the effect of "undoing" blow ups. Note that the blow up $b\colon X' \to X$ considered in Exercise 2.1 is $K_{b^{-1}\mathcal{F}}$-negative. Indeed, more generally, if one starts with a smooth projective surface X and a foliation $\mathcal{F}$ where $K_\mathcal{F}$ is nef and then performs a sequence of blow ups $p\colon \widetilde{X} \to X$ centered outside of the singular locus of the foliation, then the map

guaranteed to exist by applying Theorem 2.3 to $\widetilde{X}$ and $p^{-1}\mathcal{F}$ is precisely the morphism $\widetilde{X} \to X$.
Since the morphism f is $K_{\mathcal{F}}$-negative, it follows that the canonically determined invariants of $\mathcal{F}$ and $\mathcal{F}_Y$ are the same.

Exercise 2.4. — Notation as above. Suppose that $K_{\mathcal{F}}$ is pseudo-effective and let $K_{\mathcal{F}} = P + N$ be the Zariski decomposition. Show that $\operatorname{Exc} f = \operatorname{Supp} N$.

Exercise 2.5. — Notation as above. Show that $\kappa(X, \mathcal{F}) = \kappa(Y, \mathcal{F}_Y)$ and $\nu(X, \mathcal{F}) = \nu(Y, \mathcal{F}_Y)$.

One last remark before turning to a sketch of the proof of Theorem 2.3. In general Y will not be smooth. Consider the following example.

Example 2.6. — Let $Y = \{xy - z^2 = 0\} \subset \mathbb{C}^3$. Equivalently, we can think of Y as being the quotient singularity $\mathbb{C}^2/(x, y) \sim (-x, -y)$. The foliation on $\mathbb{C}^2$ defined by the vector field $\frac{\partial}{\partial x}$ descends to a foliation on Y, denoted $\mathcal{F}_Y$.
The singular surface Y can be resolved by a single blow up at the origin $b\colon X \to Y$ and $\mathcal{F}$, the transform of $\mathcal{F}_Y$ has the following properties.

- $E = b^{-1}(0)$ is invariant by $\mathcal{F}$;
- $E \cong \mathbb{P}^1$;
- $\mathcal{F}$ has simple singularities;
- there is a unique singular point of $\mathcal{F}$ on E; and
- $K_{\mathcal{F}} \cdot E = -1$.

Exercise 2.7. — Verify the claims in Example 2.6.

Example 2.6 is a local example, but can easily be globalised to give an example of a surface X and a foliation $\mathcal{F}$ such that the birational model where $K_{\mathcal{F}}$ becomes nef is singular.

2.5. Proof sketch of Theorem 2.3 and the Minimal Model Program. — The idea to proving Theorem 2.3 is to run an (almost) algorithmic process called the Minimal Model Program. We sketch this process now.
Our starting data is a smooth projective surface X together with a foliation $\mathcal{F}$ with simple singularities.

Step 1. — We first ask if $K_{\mathcal{F}}$ is nef. If it is, then we stop: we have already found a birational model where $K_{\mathcal{F}}$ is nef and we are in Case (1) of Theorem 2.3. Otherwise, we move onto Step 2.

Step 2. — If $K_{\mathcal{F}}$ is not nef, then (by definition) there exists a curve C such that $K_{\mathcal{F}} \cdot C < 0$. The *Cone theorem* (more on this later) guarantees that there exists a morphism $c\colon \mathbb{P}^1 \to X$ such that if $\Sigma = c(\mathbb{P}^1)$,

- $K_{\mathcal{F}} \cdot \Sigma < 0$;
- Σ is invariant by $\mathcal{F}$; and

– $\Sigma^2 \leq 0$ (strictly speaking we should take Σ to span an extremal ray in $\overline{NE}(X)$, but more on this later).

The *Contraction theorem* (again, more on this later) then guarantees the existence of a contraction $f\colon X \to Y$ such that $f(C)$ is a point if and only if $[C] \in \mathbb{R}_+[\Sigma]$.
Having produced a contraction we move onto Step 3.

Step 3. — We now analyse the morphism $f\colon X \to Y$ more carefully in terms of $\dim Y$. There are three possibilities.

– Case A (Birational Contraction): $\dim Y = 2$. In this case f is birational and contracts only Σ.

– Case B (Fibre type contraction): $\dim Y = 1$. In this case $X \to Y$ is fibration whose general fibre is $\mathbb{P}^1$, and the morphism $X \to Y$ induces the foliation $\mathcal{F}$.

– Case C: $\dim Y = 0$. In fact, since we assumed that $\mathcal{F}$ has simple singularities, this case does not occur, see [3, Proposition 5.3].

In Case B we are in Case (2) of Theorem 2.3, and so we conclude the process. In Case A, we replace X by Y and return to Step 1 and repeat the process. (Strictly speaking Y will usually be singular and so $\mathcal{F}_Y$ will not have simple singularities. This issue can be handled and is one of the motivations for our discussion of singularities in Section 3 below.)
Each time we cycle through Steps 1–3 we either produce a $\mathbb{P}^1$ fibration or $\rho(X)$ drops by 1. It follows that we can go through this process only finitely many times and therefore this process must terminate in one of the two alternatives of Theorem 2.3. □

Remark 2.8. — We observe that if $K_{\mathcal{F}}$ is pseudo-effective then we always terminate in a model where $K_{\mathcal{F}_Y}$ is nef and if $K_{\mathcal{F}}$ is not pseudo-effective we will always terminate in a model where the foliation is given by a $\mathbb{P}^1$-fibration.

Remark 2.9. — The extremal ray we find in Step 2 is not necessarily unique, and so a choice must be made at this step in the process. As a result, the output of the MMP is not necessarily unique either. Nevertheless, we expect there to be a precise way in which the different outputs of the MMP are related.

3. Singularities of the MMP

While simple singularities are a very natural class of foliation singularities on smooth varieties, as we saw in the previous section, when running the MMP we typically make our underlying variety singular. Thus, from the perspective of the MMP the class of simple singularities is insufficiently flexible. In this section we will introduce the notion of canonical and log canonical singularities. These will be turn out to be the right classes of singularities to work with for the MMP.

3.1. Canonical and log canonical singularities. — We start off by defining the discrepancy. Let X be a normal variety and let $\mathcal{F}$ be a foliation on X such that $K_{\mathcal{F}}$ is $\mathbb{Q}$-Cartier. Let E be any divisor on some birational model $p\colon X' \to X$. Let $\mathcal{F}' = p^{-1}\mathcal{F}$. We define the *discrepancy* of $\mathcal{F}$ with respect to E to be

$$a(E, \mathcal{F}) \coloneqq \text{ coefficient of } E \text{ in } K_{\mathcal{F}'} - p^* K_{\mathcal{F}}$$

where $K_{\mathcal{F}}$ and $K_{\mathcal{F}'}$ are chosen so that $p_* K_{\mathcal{F}'} = K_{\mathcal{F}}$. We define the *log discrepancy* to be

$$A(E, \mathcal{F}) \coloneqq a(E, \mathcal{F}) + \varepsilon(E)$$

where

$$\varepsilon(E) = \begin{cases} 1 & \text{if } E \text{ is not } \mathcal{F}' \text{ invariant} \\ 0 & \text{if } E \text{ is } \mathcal{F}' \text{ invariant.} \end{cases}$$

Note that in [27] what is called "discrepancy" is what we call "log discrepancy".

Exercise 3.1. — Show that the discrepancy is well defined, that is, show that it is independent of the choice of model $p\colon X' \to X$ and the choice of canonical divisor $K_{\mathcal{F}}$.

Let X be a normal variety and let $\mathcal{F}$ be a foliation on X such that $K_{\mathcal{F}}$ is $\mathbb{Q}$-Cartier. We say that $\mathcal{F}$ has *canonical* (resp. *log canonical*) singularities provided $a(E, \mathcal{F}) \geq 0$ (resp. $A(E, \mathcal{F}) \geq 0$) for all divisors E appearing on any birational model of X.
One advantage of the classes of canonical and log canonical singularities is that they make sense on singular varieties. The other benefit is that they are preserved (almost by definition) by the kind of operations we do in the MMP.

Exercise 3.2. — Let X and Y be normal varieties, let $\mathcal{F}$ be a foliation on X such that $K_{\mathcal{F}}$ is $\mathbb{Q}$-Cartier and let $f\colon X \dashrightarrow Y$ be a $K_{\mathcal{F}}$-negative birational contraction. Let $\mathcal{F}_Y$ be the transform of $\mathcal{F}$ on Y and suppose that $K_{\mathcal{F}_Y}$ is $\mathbb{Q}$-Cartier. If $\mathcal{F}$ has canonical (resp. log canonical) singularities, then $\mathcal{F}_Y$ has canonical (resp. log canonical) singularities.

In fact, these notions of singularity work in the more general setting of pairs $(\mathcal{F}, \Delta)$. By a pair we mean the data of a foliation $\mathcal{F}$ and $\mathbb{Q}$-Weil divisor $\Delta \geq 0$ such that $K_{\mathcal{F}} + \Delta$ is $\mathbb{Q}$-Cartier. Working in this more general setting is absolutely critical (as we will see), but for the reader encountering the subject for the first time you won't miss much of the bigger picture by simply taking $\Delta = 0$ wherever you see it later in these notes.

Exercise 3.3. — Let X and Y be normal varieties and suppose that there exists a birational map $f\colon X \dashrightarrow Y$. Let $\mathcal{F}$ be a foliation on X and let $\mathcal{F}_Y$ be the transform of $\mathcal{F}$ on Y. Suppose that $\mathcal{F}$ and $\mathcal{F}_Y$ have canonical singularities. Then, $\kappa(X, \mathcal{F}) = \kappa(Y, \mathcal{F}_Y)$.

Remark 3.4. — We remark that there are many other important classes of singularities from the perspective of the MMP, such as terminal, dlt, klt and F-dlt. To not overburden the first time reader with a large zoo of singularities we have chosen to focus only on canonical and log canonical singularities and we refer the reader to [12] and [22] for the definitions of these other classes of singularities.

3.2. Examples. — In general, it seems difficult to tell if a given foliation is log canonical. Part of the issue is that there is no reduction of singularities statement for foliations in complete generality. Nevertheless, there are many important special cases where it is feasible to determine whether or not a foliation is log canonical.

Let ∂ be a vector field on a smooth variety X and let $x \in \operatorname{Sing} \partial$ be a closed point. Notice that ∂ defines a linear map

$$L(\partial, x) \colon \mathfrak{m}_x/\mathfrak{m}_x^2 \to \mathfrak{m}_x/\mathfrak{m}_x^2$$

where $\mathfrak{m}_x$ is the ideal of $x \in X$. We refer to $L(\partial, x)$ as the linear part of ∂ at x. If $\mathcal{F}$ is a rank one foliation and $x \in \operatorname{Sing} \mathcal{F}$ is a closed point we define the linear part of $\mathcal{F}$ at x to be the linear part of any local generator of $\mathcal{F}$ at x. Of course, this is only well defined up to scalar multiples, but this seems unlikely to cause any confusion.

The following characterisation of log canonical singularities is very useful.

Theorem 3.5. — *Let X be a smooth variety and let $\mathcal{F}$ be a rank one foliation on X. The following are equivalent.*

- *$\mathcal{F}$ is log canonical.*
- *For any closed point $x \in \operatorname{Sing} \mathcal{F}$, the linear part of $\mathcal{F}$ at x is a non-nilpotent linear map.*

Proof. — See for instance [27, Fact I.ii.4]. □

Many natural classes of singularities are reasonable from the perspective of the MMP.

Theorem 3.6. — *Let X be a smooth variety and let $\mathcal{F}$ be a foliation on X. Then the following hold.*

(1) *If $\mathcal{F}$ is smooth, then $\mathcal{F}$ has canonical singularities.*

(2) *If $\mathcal{F}$ is a co-rank one foliation and $\mathcal{F}$ has simple singularities, then $\mathcal{F}$ has canonical singularities.*

(3) *If $\mathcal{F}$ is induced by a toroidal morphism $f \colon X \to Y$, then $\mathcal{F}$ has log canonical singularities.*

Proof. — Item (1) follows from [19, Lemma 5.9]. Item (2) follows from [16, Lemma 2.9]. Item (3) follows from [2, Lemma 3.1]. □

Exercise 3.7. — Let X and Y be normal varieties and let $q \colon X \to Y$ be a quasi-étale morphism (e.g., $X \cong \mathbb{C}^n$ and $Y = \mathbb{C}^n/G$ where G is a group acting without pseudo-reflections). Let $\mathcal{G}$ be a foliation on Y and let $\mathcal{F} = q^{-1}\mathcal{G}$. Then, $\mathcal{F}$ is log canonical if and only if $\mathcal{G}$ is log canonical.

4. MMP in higher dimensions

We now turn the problem of finding the "simplest" representative of any given birational class of foliated varieties in arbitrary dimension. By taking Theorem 2.3 as our guide we hope that given a smooth projective variety X equipped with a foliation $\mathcal{F}$ we can find a $K_{\mathcal{F}}$-negative birational contraction $f \colon X \dashrightarrow Y$ such that either

- $K_{\mathcal{F}_Y}$ is nef (*minimal model*); or

– there exists a fibration $g\colon Y \to Z$ such that $-K_{\mathcal{F}_Y}$ is g-ample (*Mori fibre space*).

As with Theorem 2.3 we should try to realise this goal as an output of an algorithmic process. We also remark that in the second alternative described above we do not expect that the fibration $g\colon Y \to Z$ induces the foliation (in contrast to the surface case). However, in this second case we do have that through a general point of Y there is a rational curve tangent to the foliation, see [8].

4.1. Cone and Contraction conjectures. — In this section we will explain the Cone and Contraction results mentioned above. Neither of these results are known in complete generality, although it now seems clear what the correct statements should be. These conjectures predict a connection between the birational geometry of X and $\mathcal{F}$ and the convex geometry of the Kleiman–Mori cone of X. Before precisely stating the Cone and Contraction conjectures we briefly recall the definition of the Kleiman–Mori cone, cf. [22, Definition 1.17]. Let X be a proper variety. We denote by $N_1(X)$ the real vector space generated by 1-cycles modulo numerical equivalence. The Kleiman–Mori cone, denoted $\overline{NE}(X) \subset N_1(X)$, is the closed convex cone generated by effective 1-cycles.

***Conjecture 4.1* (Cone Conjecture).** — *Let X be a normal $\mathbb{Q}$-factorial projective variety, let $\mathcal{F}$ be a foliation on X and let $\Delta \geq 0$ be a $\mathbb{Q}$-Weil divisor such that $K_{\mathcal{F}}+\Delta$ is $\mathbb{Q}$-Cartier and $(\mathcal{F}, \Delta)$ is log canonical. Then there exists a countable collection of rational curves $\xi_1, \xi_2, \dots$ tangent to $\mathcal{F}$ such that*

$$\overline{NE}(X) = \overline{NE}(X)_{K_{\mathcal{F}}+\Delta \geq 0} + \sum_i \mathbb{R}_+[\xi_i]$$

and $0 < -K_{\mathcal{F}} \cdot \xi \leq 2 \dim X$.

We say a ray $R \subset \overline{NE}(X)$ is *extremal* provided that whenever we have $\alpha, \beta \in \overline{NE}(X)$ such that $\alpha + \beta \in R$ then $\alpha, \beta \in R$. A direct consequence of the Cone Conjecture is that every $K_{\mathcal{F}}$-negative extremal ray is spanned by the class of a curve.

***Conjecture 4.2* (Contraction Conjecture).** — *Let X be a normal $\mathbb{Q}$-factorial projective variety, let $\mathcal{F}$ be a foliation on X and let $\Delta \geq 0$ be a $\mathbb{Q}$-Weil divisor such that $K_{\mathcal{F}} + \Delta$ is $\mathbb{Q}$-Cartier and $(\mathcal{F}, \Delta)$ is log canonical. Let R be a $K_{\mathcal{F}}$-negative extremal ray. Then there exists a contraction $c_R\colon X \to Y$ such that for any proper curve $C \subset X$, $\dim f(C) = 0$ if and only if $[C] \in R$.*

For the reader seeing these statements for the first time, you will not lose much in the way of the bigger picture by assuming that X is smooth.
Let us observe that if $K_{\mathcal{F}}$ is not nef, then by definition there exists a curve C such that $K_{\mathcal{F}} \cdot C < 0$. This implies that there is some $K_{\mathcal{F}}$-negative extremal ray. As we will see, the natural thing to study is not the curves on which $K_{\mathcal{F}}$ is negative, but rather the extremal rays. Indeed, for the Contraction Conjecture it is critical that we work with an extremal ray: the Contraction Conjecture *does not* guarantee that any $K_{\mathcal{F}}$-negative curve can be contracted.

Exercise 4.3. — Let X be a smooth projective surface.

(1) Let $C \subset X$ be a curve with $C^2 < 0$. Show that the class of C spans an extremal ray in $\overline{NE}(X)$.

(2) Show by example that an extremal ray in $\overline{NE}(X)$ is not necessarily spanned by the class of a curve $C \subset X$.

I will now summarise what is known about these two conjectures.

- When $\mathcal{F}$ is the foliation with one leaf, i.e., $T_{\mathcal{F}} = T_X$, then both of these conjectures are well known, see [22, Theorem 3.7] for the case where X is klt and [1] for the case where X is log canonical.
- When $\dim X \leq 3$ and $\mathcal{F}$ is co-rank one, Conjecture 4.1 is known by [16, 31, 32].
- When $\dim X \leq 3$, X has klt singularities, $\mathcal{F}$ is co-rank one and $\mathcal{F}$ has canonical singularities (or F-dlt singularities), Conjecture 4.2 is known by [12]. When $\mathcal{F}$ has log canonical singularities, Conjecture 4.2 is known by [17].
- When $\mathcal{F}$ is rank one, Conjecture 4.1 is known by [13].
- When $\dim X \leq 3$, X has klt singularities and $\mathcal{F}$ is rank one, Conjecture 4.2 is known by [16].
- When $\mathcal{F}$ is algebraically integrable, Conjecture 4.1 is known by [2].
- When $\mathcal{F}$ is algebraically integrable and X has klt singularities Conjecture 4.2 is known by [14, 24].

Remark 4.4. —

- The $\mathbb{Q}$-factorial hypothesis is a natural one from the perspective of MMP, but there are many situations where it would be preferable to know these Conjectures without any $\mathbb{Q}$-factoriality hypothesis. It is usually easy to drop the $\mathbb{Q}$-factoriality hypothesis from the Cone theorem, but removing it from the Contraction theorem is often subtle.
- In most cases where we have proven the Contraction theorem, we also require that the underlying variety have klt singularities. This is a natural condition and is frequently satisfied in situations of interest, but it would be interesting to know if it could be removed.

4.2. Flips. — If we attempt to repeat the proof of Theorem 2.3 in dimension ≥ 3 a new feature of higher dimensional geometry gets in the way. Let R be a $K_{\mathcal{F}}$-negative extremal ray and let $c_R \colon X \to Y$ be the associated contraction morphism. If $\operatorname{Exc} c_R$ is a divisor, we can proceed much in the same way as in the proof of Theorem 2.3, but if c_R is a small contraction, i.e., $\dim \operatorname{Exc} c_R \leq \dim X - 2$, we run into a serious issue: $K_{\mathcal{F}_Y}$ is not $\mathbb{Q}$-Cartier. Indeed, suppose for sake of contradiction that $K_{\mathcal{F}_Y}$ were $\mathbb{Q}$-Cartier. Since c_R is small and $K_{\mathcal{F}_Y}$ is $\mathbb{Q}$-Cartier we have $K_{\mathcal{F}} \sim_{\mathbb{Q}} c_R^* K_{\mathcal{F}_Y}$. However, if C is any curve such that $[C] \in R$ then we get the immediate contradiction

$$0 > K_{\mathcal{F}} \cdot C = (c_R^* K_{\mathcal{F}_Y}) \cdot C = 0$$

where the last equality holds by the projection formula.
Since $K_{\mathcal{F}_Y}$ is not $\mathbb{Q}$-Cartier, we can no longer continue the MMP process, indeed, it no longer makes sense to ask if $K_{\mathcal{F}_Y}$ is nef or not. Something else must be done at this point and this

is where the flip surgery comes in to play. The next example shows that small contractions really do occur in the course of the MMP, and also hopefully gives a feeling for how the flip surgery works.

Example 4.5. — Let S be a smooth surface, let $P \in S$ be a point and let $b\colon X \to S \times \mathbb{P}^1$ be the blow up of $S \times \mathbb{P}^1$ at $(P, [1:0])$. Let $\mathcal{F}$ be the foliation on X induced by the fibration $f\colon X \to S$. Let $C \subset X$ denote the strict transform of $\{P\} \times \mathbb{P}^1$ and let E denote the b-exceptional divisor. We make the following observations:

- C and E are both $\mathcal{F}$-invariant;
- the restriction of $\mathcal{F}$ to E is the pencil of lines passing through $C \cap E$;
- if L is one of these lines, then L spans a $K_{\mathcal{F}}$-negative extremal ray, denoted R_1; and
- C spans a $K_{\mathcal{F}}$-negative extremal ray, denoted R_2.

Observe that the contraction $c_{R_2}\colon X \to Y$ associated to R_2 is a small contraction (it contracts exactly the curve C), so in fact small contractions do occur in the course of the MMP. (The contraction associated to R_1 is simply the blow down of E.)
To get around this issue we "flip" the curve C by blowing it up and then blowing it down. To be precise, let us notice that $N_{C/X} = \mathcal{O}(-1) \oplus \mathcal{O}(-1)$. Thus, if $p\colon Y \to X$ is the blow up in C, then $D := \operatorname{Exc} p \cong \mathbb{P}^1 \times \mathbb{P}^1$. We can contract $D \subset Y$ along either of its two rulings, so let $q\colon Y \to X^+$ be the contraction of the other ruling and let $C^+ := q(D)$.
Note that $q \circ p^{-1}\colon X \dashrightarrow X^+$ gives an isomorphism between $X \setminus C$ and $X^+ \setminus C^+$. Moreover, letting $\mathcal{F}^+$ denote the transform of $\mathcal{F}$ on X, we can compute that $K_{\mathcal{F}^+} \cdot C^+ > 0$. Observe that $\mathcal{F}^+$ is induced by an equidimensional contraction $f\colon X^+ \to \widetilde{S}$ where $\widetilde{S} \to S$ is the blow up of S at P and there is a unique $K_{\mathcal{F}^+}$-negative extremal ray whose associated contraction is precisely the fibration $X^+ \to \widetilde{S}$. Thus, by flipping C we have side stepped the issue of the small contraction. We are able to continue running the MMP by continuing from X^+ rather than continuing from Y.

Exercise 4.6. — Modify Example 4.5 to produce an example where there is a unique $K_{\mathcal{F}}$-negative extremal ray, and the associated contraction is small. Show that in this example, the flip of C produces a model where $K_{\mathcal{F}^+}$ is nef.

In general, given a small contraction $c_R\colon X \to Y$, the *flip* of the contraction is the following diagram

$$\begin{tikzcd} X \arrow[rr, dashed] \arrow[dr, "c_R"'] & & X^+ \arrow[dl, "c_R^+"] \\ & Y & \end{tikzcd}$$

where

- $X \dashrightarrow X^+$ is an isomorphism in codimension $= 1$;
- $\rho(X^+/Y) = 1$; and
- $K_{\mathcal{F}^+}$ is c_R^+-ample.

From this description it is unclear if the flip of a given contraction in fact exists. However, we conjecture that it is always possible to construct the flip.

***Conjecture 4.7* (Existence of flips).** — *Let X be a normal $\mathbb{Q}$-factorial projective variety, let $\mathcal{F}$ be a foliation on X and let $\Delta \geq 0$ be a $\mathbb{Q}$-Weil divisor such that $K_{\mathcal{F}} + \Delta$ is $\mathbb{Q}$-Cartier and $(\mathcal{F}, \Delta)$ is log canonical. Let R be a $K_{\mathcal{F}}$-negative extremal ray and suppose that c_R is small contraction. Then the flip exists.*

Remark 4.8. — The flip, if it exists, is unique. Indeed, one can show that for some integer $l \gg 0$,

$$X^+ \cong \operatorname{Proj}_Y \bigoplus_{m\geq 0} c_{R*}\mathcal{O}(ml(K_{\mathcal{F}} + \Delta)).$$

Exercise 4.9. — Toric varieties and toric maps provide a wealth of concrete examples of the operations of the MMP. The flip described in Example 4.5 is in fact a toric flip. Describe this flip using operations on the associated fan. Try to write down a few more examples of flips between torus invariant foliations on toric varieties.

I will summarise what is known about this conjecture.

- When $\mathcal{F}$ is the foliation with one leaf, i.e., $T_{\mathcal{F}} = T_X$, Conjecture 4.7 follows from [5] in the klt case and [4, 20] in the log canonical case.
- When $\dim X \leq 3$, X has klt singularities, $\mathcal{F}$ is co-rank one and $\mathcal{F}$ has canonical singularities (or F-dlt singularities), Conjecture 4.7 is known by [12] and [17].
- When $\dim X \leq 3$, X has klt singularities and $\mathcal{F}$ is rank one, Conjecture 4.7 is known by [16] (cf. also [26]).
- When X has klt singularities and $\mathcal{F}$ is algebraically integrable, Conjecture 4.7 is known by [14, 24].

4.3. Running the algorithm in higher dimensions and termination. — Let's explain now how to run the MMP assuming Conjectures 4.1, 4.2 and 4.7.
Our starting data is a $\mathbb{Q}$-factorial projective variety X and $\mathcal{F}$ is a foliation on X with log canonical singularities.

Step 1. — We ask if $K_{\mathcal{F}}$ is nef. If $K_{\mathcal{F}}$ is nef, then we terminate, we have reached our minimal model. Otherwise we go to Step 2.

Step 2. — If $K_{\mathcal{F}}$ is not nef there exists $K_{\mathcal{F}}$-negative extremal ray R. By the Cone Conjecture R is spanned by the class of a curve ξ. By the Contraction Conjecture there exists a contraction $c_R \colon X \to Y$ associated to R.

Step 3. — We argue in cases based on the morphism $c_R \colon X \to Y$.

- (Fibre type contraction) If $c_R \colon X \to Y$ is not birational, i.e., $\dim X > \dim Y$, then we are in the Mori fibre space case and we terminate.
- (Divisorial contraction) If $c_R \colon X \to Y$ is birational and $\operatorname{Exc} c_R$ is a divisor, then Y is again a $\mathbb{Q}$-factorial projective variety, and $\mathcal{F}_Y$ has log canonical singularities. We may then replace X by Y and return to Step 1.

- (Flipping contraction) If $c_R\colon X \to Y$ is birational and $\operatorname{Exc} c_R$ is not a divisor, then by Conjecture 4.7 the flip $X \dashrightarrow X^+$ exists. X^+ is a $\mathbb{Q}$-factorial projective variety and $\mathcal{F}^+$ has log canonical singularities. We may then replace X by X^+ and return to Step 1.

Remark 4.10. — As noted above, in most of the cases where Conjectures 4.2 and 4.7 are known it is under the additional hypothesis that X is klt. In these cases it possible to show that the contraction/flips preserves klt singularities. This is not obvious, but follows from the way the contraction/flip is constructed. It would be desirable to have a "conceptual" reason that the klt condition is preserved by contractions/flips with respect to $K_{\mathcal{F}}$.

We now turn to the question of whether or not this process terminates. In this process we can only have finitely many divisorial contractions since each divisorial contraction lowers $\rho(X)$ by 1. However, a flip leaves the Picard number unchanged, i.e., $\rho(X) = \rho(X^+)$, so we cannot use the Picard number to deduce that we only encounter finitely many flips. Indeed, showing that there is no infinite sequence of flips is a real mystery (even in the case of varieties!).

***Conjecture 4.11* (Termination of Flips).** — *There is no infinite sequence of flips. In particular, the MMP terminates.*

This is a subtle problem, and comparatively speaking not much is known.

- When $\dim X = 3$ any sequence of flips terminates. In the case where $T_{\mathcal{F}} = T_X$ proof goes by showing that an invariant called the difficulty drops after every flip, see e.g. [22, Theorem 6.17]. This strategy needs to be modified to work in general, but this can be done, cf. [32, Section 2]. Unfortunately, the difficulty is not a sufficiently refined invariant to prove termination in dimension ≥ 4.

- When the rank of $\mathcal{F} = 1$ any sequence of flips terminates, cf. [16].

4.4. Termination with Scaling. — While Conjecture 4.11 seems to be a challenging conjecture, there is a weaker version of termination, namely Termination of Flips with Scaling, which suffices to prove the termination of the MMP algorithm in many cases of interest. The idea behind Termination of Scaling is to run the MMP with an auxiliary divisor which guides our choice of extremal ray. To be precise, consider the following set up. Let X be a normal $\mathbb{Q}$-factorial projective variety, let $\mathcal{F}$ be a foliation on X and let $\Delta \geq 0$ be a $\mathbb{Q}$-Weil divisor such that $(\mathcal{F}, \Delta)$ is log canonical. Let $C \geq 0$ be a $\mathbb{Q}$-divisor such that $(\mathcal{F}, \Delta + C)$ is log canonical and $K_{\mathcal{F}} + \Delta + C$ is nef.
Set $X_0 := X$, $\mathcal{F}_0 := \mathcal{F}$, $\Delta_0 := \Delta$ and $C_0 := C$. We define a sequence of divisorial contractions/flips with scaling of C

$$X_0 \dashrightarrow X_1 \dashrightarrow X_2 \dashrightarrow \dots$$

as follows. Set $\lambda_0 := \inf\{t : K_{\mathcal{F}} + \Delta + tC \text{ is nef.}\}$. The Cone Conjecture implies that there exists a $K_{\mathcal{F}}+\Delta$-negative extremal ray $\mathbb{R}_+[\xi]$ such that $(K_{\mathcal{F}}+\Delta+\lambda_0 C)\cdot\xi = 0$. Let $\phi_0\colon X_0 \dashrightarrow X_1$ be the divisorial contraction/flip associated to $\mathbb{R}_+[\xi]$, let $\mathcal{F}_1$ be the transform of $\mathcal{F}_0$, let $\Delta_1 := \phi_{0*}\Delta_0$ and let $C_1 := \phi_{0*}C_0$. We can then continue this process and produce a sequence of varieties $X_0 \dashrightarrow X_1 \dashrightarrow X_2 \dashrightarrow \dots$ and scaling constants $\lambda_0 \geq \lambda_1 \geq \dots$ such that at each step we contract/flip a $K_{\mathcal{F}_i} + \Delta_i + \lambda_i C_i$-trivial extremal ray.

***Conjecture 4.12* (Termination of Flips with Scaling).** — *Any sequence of flips with scaling terminates.*

Remark 4.13. — In practice we usually take C_0 to be a sufficiently ample divisor, but we remark that we usually lose ampleness after the first step in the MMP, i.e., typically C_1 is not an ample divisor.

Much more is known about Conjecture 4.12.

- If (X, Δ) is klt, C is ample, and either $K_X + \Delta$ is big or Δ is big, then any sequence of flips with scaling of C divisor terminates, cf. [5].

- Suppose that X is klt, $\mathcal{F}$ is an algebraically integrable foliation and $\Delta \geq 0$ is a $\mathbb{Q}$-divisor such that $(\mathcal{F}, \Delta)$ is log canonical. Let A, H be two ample divisors. Then any sequence of $K_{\mathcal{F}} + \Delta + A$-flips with scaling of H terminates, [10]

It must be emphasised that proving Termination of Flips with Scaling of an ample divisor is quite subtle. Part of the difficulty is the failure of Bertini's theorem for foliations, cf. [14, Section 3.5.2]. The way to address this technical complication is the systematic use of generalised pairs with foliations, see for instance, [24].

4.5. Summary. — While we now know a lot about the Minimal Model Program for foliations, it is still the case that a complete statement is only known in dimension ≤ 3. In addition to the technical challenges mentioned above, part of the issue is that a complete reduction of singularities statement is only known in dimension ≤ 3, see [9, 27]. Nevertheless, the statement in dimension ≤ 3 is quite satisfactory and for completeness's sake we state it here.

Theorem 4.14. — *Let X be a smooth projective variety of dimension ≤ 3, let $\mathcal{F}$ be a foliation on X and let $\Delta \geq 0$ be a $\mathbb{Q}$-divisor such that $(\mathcal{F}, \Delta)$ has log canonical singularities. Then, there exists a $K_{\mathcal{F}} + \Delta$-negative birational contraction $\phi\colon X \dashrightarrow Y$ such that if $\mathcal{F}_Y$ denotes the transform of $\mathcal{F}$ and $\Delta = \phi_* \Delta_Y$ then either*

(1) *$K_{\mathcal{F}_Y} + \Delta_Y$ is nef; or*

(2) *there exists a fibration $f\colon Y \to Z$ such that $-(K_{\mathcal{F}_Y} + \Delta_Y)$ is f-ample.*

5. A proof sketch of the Cone Theorem in dimension three

In this section, I will give a brief sketch of the proof of the Cone Theorem for foliations where the underlying variety is of dimension three. This is partly to give a flavour of some of standard proof techniques in the MMP, but also to introduce Bend and Break.

5.1. Bend and Break. — Mori's Bend and Break is a beautiful technique for producing rational curves on algebraic varieties which proceeds by finding rational curves on the variety after a reduction to characteristic $= p$, and then lifting these curves back to characteristic $= 0$. As observed by Miyaoka, the Bend and Break technique can be adapted to produce rational curves tangent to a foliation. We refer to [30] for a complete proof of this technique, but in practice it seems like the following statement of the Bend and Break technique is particularly useful.

***Theorem 5.1* (Bend and Break).** — *Let X be a normal projective variety of dimension $= n$, let $\mathcal{F}$ be a foliation on X and let $D_1, \dots, D_n$ be $\mathbb{R}$-Cartier nef divisors such that*

(1) $D_1 \cdot \ldots \cdot D_n = 0$; *and*

(2) $-K_{\mathcal{F}} \cdot D_2 \cdot \ldots \cdot D_n > 0$.

Then, through a general point of X there is a rational curve Σ which is tangent to $\mathcal{F}$ and $D_1 \cdot \Sigma = 0$. Moreover, if M is any nef divisor then we have

$$M \cdot \Sigma \leq 2 \dim X \frac{M \cdot D_2 \cdot \ldots \cdot D_n}{-K_{\mathcal{F}} \cdot D_2 \cdot \ldots \cdot D_n}.$$

Proof. — See [31, Corollary 2.28], cf. also [21]. □

5.2. Proof sketch of the Cone Theorem when $\dim X = 3$. — We argue by induction on the dimension of X, so we may assume that the Cone Theorem holds for foliations on surfaces.

Let R be a $K_{\mathcal{F}}$-negative extremal ray (strictly speaking we should take an exposed extremal ray). Let H_R be a supporting hyperplane to R, i.e. H_R is a nef $\mathbb{R}$-Cartier divisor which is characterised (up to scaling) by the property that $H_R \cdot \alpha = 0$ if and only if $\alpha \in R$. Note that by Kleiman's criterion (up to replacing H_R by a multiple) we have $A := H_R - K_{\mathcal{F}}$ is an ample divisor.

We now argue in cases based on $\nu(H_R)$.

Case 1: $\nu(H_R) < 3$. — In this case we define

$$D_i = \begin{cases} H_R & \text{for } 1 \leq i \leq \nu(H_R) + 1 \\ A & \text{for } \nu(H_R) + 1 < i \leq 3. \end{cases}$$

It is easy to verify that $D_1 \cdot D_2 \cdot D_3 = 0$ and that $-K_{\mathcal{F}} \cdot D_2 \cdot D_3 = A \cdot D_2 \cdot D_3 > 0$. We may then apply Bend and Break, Theorem 5.1, to produce a rational curve Σ such that $0 = D_1 \cdot \Sigma = H_R \cdot \Sigma$, in particular Σ spans R as required.

Remark 5.2. — In fact, the argument in Case 1 works in arbitrary dimension, mutatis mutandis.

Case 2: $\nu(H_R) = 3$. — Since H_R is nef and $H_R^3 \neq 0$ we see that H_R is a big divisor. By Kodaira's Lemma we see that $H_R \sim_{\mathbb{R}} A + E$ where A is ample and $E \geq 0$. There exists an irreducible component S of E such that R lies in the image of the natural map $\overline{NE}(T) \to \overline{NE}(X)$ where $n \colon T \to S \subset X$ is the normalisation of S. We have a restricted foliation $\mathcal{F}_T$ on T (which might be the entire tangent sheaf of T) and by foliation adjunction, cf. [14], we may write $n^*(K_{\mathcal{F}} + \varepsilon(S)S) = K_{\mathcal{F}_T} + \Delta_T$ where $\Delta_T \geq 0$. (N.b. this shows why working with pairs is important in the MMP: pairs will naturally arise when proceeding by induction on dimension by using adjunction)

There exists a $K_{\mathcal{F}_T} + \Delta_T$-negative extremal ray $R' \subset \overline{NE}(T)$ whose image in $\overline{NE}(X)$ is R. Since $\dim T \leq 2$ we would like apply the Cone theorem for $(\mathcal{F}_T, \Delta_T)$ to produce a rational curve spanning R', and hence a rational curve spanning R. The issue is that we do not have any control on the singularities of T or of $(\mathcal{F}_T, \Delta_T)$. The way around this issue is to replace $(\mathcal{F}, S)$ by a partial resolution of singularities called an F-dlt modification. On the F-dlt modification we can guarantee that $(\mathcal{F}_T, \Delta_T)$ has log canonical singularities, and so our inductive argument can proceed. Proving the existence of F-dlt modifications is subtle and while we expect in can be done in full generality, we are only able to prove it in some special cases, e.g., in dimension ≤ 3.

6. Important and Interesting things that I didn't get to talk about

Here is a quick list of some important and interesting results and techniques that I didn't get a chance to discuss, but I wish I could have.

6.1. Failure of the base point free theorem and abundance. — The base point free theorem says the following (see for instance [22, Theorem 3.3]).

Theorem 6.1. — *Let X be a normal projective variety and let $\Delta \geq 0$ be a $\mathbb{Q}$-divisor such that (X, Δ) is klt. Let H be a $\mathbb{Q}$-Cartier and nef divisor such that $H - (K_X + \Delta)$ is big and nef. Then H is semi-ample.*

This theorem plays a central role in the proofs of the Cone and Contraction theorem for varieties, but this type of theorem does not hold for foliations more generally:

Example 6.2. — Let X_0 be the Bailey–Borel compactification of a bi-disc quotient by an irreducible lattice and let $\mathcal{F}_0$ be one of the tautological foliations on X_0. Let $H \subset X_0$ be a general sufficiently ample divisor and let $Y_0 \to X_0$ be a ramified cover which is ramified along H to order r. Note that Y_0 has some number of singularities which are (locally) isomorphic to the cusps of X_0. Let $p \colon Y \to Y_0$ be the resolution of these cusps and let $\mathcal{G}$ be the pull-back of $\mathcal{F}_0$ along the natural map $Y \to X_0$.
By the Riemann–Hurwitz formula for foliations (see [19, Lemma 3.4] and [16, Proposition 2.2]) $K_{\mathcal{G}}$ is big and nef, however, by [25, Corollary IV.2.3] $K_{\mathcal{G}}$ is not semi-ample.

Because of this failure the proofs of the Cone and Contraction theorem for foliations need to take a very different strategy than the one taken for varieties (we remark, however, that a weak version of the base point free theorem is still expected to hold, see [15] for one such statement).
This example also demonstrates a very intriguing difference between the birational geometry of varieties and the birational geometry of foliations. The Abundance Conjecture is the following central conjecture in the birational geometry of algebraic varieties.

***Conjecture 6.3* (Abundance Conjecture).** — *Let X be a normal projective variety with log canonical singularities. If K_X is nef, then K_X is semi-ample.*

The foliations described above show that if $K_{\mathcal{F}}$ is nef, then it is not necessarily the case that $K_{\mathcal{F}}$ is semi-ample, in particular, the Abundance Conjecture does not hold for foliations. Nevertheless, as we will see in Section 6.2, at least for surfaces, the Abundance Conjecture only fails to hold for very specific foliations. This perhaps suggests the possibility of a refined version of the Abundance Conjecture which holds for foliations as well.

6.2. Kodaira–Enriques Classification. — While the existence of minimal models or Mori fibre spaces is important and interesting in its own right, we should not forget its role in the classification problem we started with. Indeed, in the case of foliations on surfaces a rather complete classification has been worked out which we summarise here.
In the following table we assume that $(X, \mathcal{F})$ is the output of the $K_{\mathcal{F}}$-MMP starting from a smooth surface and a foliation with canonical singularities.

$\nu(X, \mathcal{F})$	$\kappa(X, \mathcal{F})$	**Description**
$-\infty$	$-\infty$	$\mathcal{F}$ is induced by a $\mathbb{P}^1$-fibration.
0	0	Up to a quasi-étale cover, one of the following holds – $X \cong E \times C$ where E is an elliptic curve and $\mathcal{F}$ is induced by the projection to C; or – X is an equivariant compactification of a Lie group and $\mathcal{F}$ is generated by an element in the Lie algebra.
1	$-\infty$	X is a resolution of the cusps of the Bailey–Borel compactification of a bi-disc quotient by an irreducible lattice and $\mathcal{F}$ is induced by one of the two tautological foliations on the bi-disc.
1	1	Up to replacing X by cover ramified only on $\mathcal{F}$-invariant curves, – X is a $\mathbb{P}^1$-bundle over a curve and $\mathcal{F}$ Ricatti with repsect to this structure; – X admits an isotrivial elliptic fibration and $\mathcal{F}$ is turbulent with respect to this fibration; – $\mathcal{F}$ is induced by a non-isotrivial elliptic fibration; or – $\mathcal{F}$ is induced by an isotrivial fibration in curves of genus ≥ 2.
2	2	General type

We point out a very interesting feature of the above classification. We always have $\nu(X, \mathcal{F}) = \kappa(X, \mathcal{F})$ *except* in the case of bi-disc quotients. As noted earlier these foliations give examples where abundance fails. A remarkable feature of the classification is that these are the only such failures in the case of surface foliations. It would be very interesting to understand this phenomenon in higher dimensions.

Question 6.4. — *Can we classify all foliations with log canonical singularities for which $\nu(X, \mathcal{F}) \neq \kappa(X, \mathcal{F})$?*

We also remark that the classification does not say anything in particular about the general type case. In this case, in analogy with the moduli space of genus ≥ 2 curves, the expectation is that one should produce a moduli space which parametrises foliations of general type. This will be explored in forthcoming work, [34].

6.3. Adjoint foliated structures. — Starting with [29] it was observed that many of the "pathologies" of the birational geometry of foliations can be handled by working with divisors of the form $(1-t)K_{\mathcal{F}}+tK_X$ where $0 < t \ll 1$. This line of thought was pursued further in [33] where the boundedness of certain classes of foliations was proven. In [10, 11] this idea was pushed further in the case of algebraically integrable foliations where we consider all values of $t \in [0,1]$.

6.4. Generalised pairs. — As alluded to previously, it has become clear recently that the framework of generalised pairs is the correct one for studying many questions in the birational geometry of foliations. Generalised pairs were introduced in [6] and have proven to be a powerful tool in birational geometry more generally. We refer to [10, 11, 18, 24] for some examples of the use of generalised pairs in the context of the study of foliations.

6.5. Singularities. — The MMP is an effective tool for the study of singularities of algebraic varieties, and the foliated MMP seems to be a similarly useful tool for the study of foliation singularities. We refer to [32] for some illustrative examples.

6.6. More on the MMP for algebraically integrable foliations. — A contraction $f\colon X \to Y$ between normal projective varieties induces a foliation $\mathcal{F}$. If X has log canonical singularities, then we can run a K_X-MMP relatively over Y, i.e., we only contract/flip curves which are contained in fibres of f. Should this MMP terminate, the output will either be a relative Mori fibre space over Y, or model where K_X is nef over Y. If $\mathcal{F}$ has log canonical singularities, then we can also run a $K_{\mathcal{F}}$-MMP. In general, the $K_{\mathcal{F}}$-MMP differs from the relative K_X-MMP, that is to say, a sequence of steps of the relative K_X-MMP is not necessarily a sequence of steps in the $K_{\mathcal{F}}$-MMP, and vice versa. It must also be remarked that if $K_{\mathcal{F}}$ is pseudo-effective, then the $K_{\mathcal{F}}$-MMP terminates in a model where $K_{\mathcal{F}}$ is globally nef (not just relatively nef over Y).
We also remark that even if X has log canonical singularities, this is not enough to guarantee that $\mathcal{F}$ log canonical singularities. It would be interesting to understand more precisely the relationship between the singularities of the fibration $f\colon X \to Y$ and the singularities of the induced foliation. We refer to [2, Proposition 3.7] for one such result in this direction.

References

[1] F. Ambro, "Quasi-log varieties", *Tr. Mat. Inst. Steklova* **240** (2003), p. 220-239.

[2] F. Ambro, P. Cascini, V. Shokurov & C. Spicer, "Positivity of the moduli part", 2022, `https://arxiv.org/abs/2111.00423v2`.

[3] C. Araujo & S. Druel, "On Fano foliations", *Adv. Math.* **238** (2013), p. 70-118.

[4] C. Birkar, "Existence of log canonical flips and a special LMMP", *Publ. Math., Inst. Hautes Étud. Sci.* **115** (2012), p. 325-368.

[5] C. Birkar, P. Cascini, C. D. Hacon & J. McKernan, "Existence of minimal models for varieties of log general type", *J. Am. Math. Soc.* **23** (2010), no. 2, p. 405-468.

[6] C. Birkar & D.-Q. Zhang, "Effectivity of Iitaka fibrations and pluricanonical systems of polarized pairs", *Publ. Math., Inst. Hautes Étud. Sci.* **123** (2016), p. 283-331.

[7] M. Brunella, *Birational geometry of foliations*, IMPA Monographs, vol. 1, Springer, 2015, reprint of the 2000 edition with new results.

[8] F. Campana & M. Păun, "Foliations with positive slopes and birational stability of orbifold cotangent bundles", *Publ. Math., Inst. Hautes Étud. Sci.* **129** (2019), p. 1-49.

[9] F. Cano, "Reduction of the singularities of codimension one singular foliations in dimension three", *Ann. Math. (2)* **160** (2004), no. 3, p. 907-1011.

[10] P. Cascini, J. Han, J. Liu, F. Meng, C. Spicer, R. Svaldi & L. Xie, "Minimal model program for algebraically integrable adjoint foliated structures", 2024, https://arxiv.org/abs/2408.14258.

[11] ———, "On finite generation and boundedness of adjoint foliated structures", 2025, https://arxiv.org/abs/2504.10737.

[12] P. Cascini & C. Spicer, "MMP for co-rank one foliations on threefolds", *Invent. Math.* **225** (2021), no. 2, p. 603-690.

[13] ———, "Foliation adjunction", *Math. Ann.* **391** (2025), no. 4, p. 5695-5727.

[14] ———, "MMP for algebraically integrable foliations", in *Higher dimensional algebraic geometry—a volume in honor of V. V. Shokurov*, London Mathematical Society Lecture Note Series, vol. 489, Cambridge University Press, 2025, p. 69-84.

[15] ———, "On base point freeness for rank one foliations", 2025, https://arxiv.org/abs/2509.03109.

[16] ———, "On the MMP for rank one foliations on threefolds", *Forum Math. Pi* **13** (2025), article no. e20 (38 pages).

[17] P. Chaudhuri & R. Mascharak, "Log canonical minimal model program for corank one foliations on threefolds", 2024, https://arxiv.org/abs/2410.05178.

[18] G. Chen, J. Han, J. Liu & L. Xie, "Minimal model program for algebraically integrable foliations and generalized pairs", 2023, https://arxiv.org/abs/2309.15823.

[19] S. Druel, "Codimension 1 foliations with numerically trivial canonical class on singular spaces", *Duke Math. J.* **170** (2021), no. 1, p. 95-203.

[20] C. D. Hacon & C. Xu, "Existence of log canonical closures", *Invent. Math.* **192** (2013), no. 1, p. 161-195.

[21] S. Keel, K. Matsuki & J. McKernan, "Log abundance theorem for threefolds", *Duke Math. J.* **75** (1994), no. 1, p. 99-119.

[22] J. Kollár & S. Mori, *Birational geometry of algebraic varieties. With the collaboration of C. H. Clemens and A. Corti.*, Cambridge Tracts in Mathematics, vol. 134, Cambridge University Press, 1998, viii+254 pages.

[23] B. Lehmann, "A snapshot of the minimal model program", in *Surveys on recent developments in algebraic geometry*, Proceedings of Symposia in Pure Mathematics, vol. 95, American Mathematical Society, 2017, p. 1-32.

[24] J. Liu, F. Meng & L. Xie, "Minimal model program for algebraically integrable foliations on klt varieties", 2025, https://arxiv.org/abs/2404.01559v3.

[25] M. McQuillan, "Canonical models of foliations", *Pure Appl. Math. Q.* **4** (2008), no. 3, p. 877-1012.

[26] ———, "Semi-stable reduction of foliations", in *Arithmetic and algebraic geometry. A mathematical tribute to Yuri Manin*, Simons Symposia, Springer, 2024, p. 239-351.

[27] M. McQuillan & D. Panazzolo, "Almost étale resolution of foliations", *J. Differ. Geom.* **95** (2013), no. 2, p. 279-319.

[28] L. G. Mendes, "Kodaira dimension of holomorphic singular foliations", *Bol. Soc. Bras. Mat., Nova Sér.* **31** (2000), no. 2, p. 127-143.

[29] J. V. Pereira & R. Svaldi, "Effective algebraic integration in bounded genus", *Algebr. Geom.* **6** (2019), no. 4, p. 454-485.

[30] N. I. Shepherd-Barron, "Miyaoka's theorems on the generic seminegativity of T_X and on the Kodaira dimension of minimal regular threefolds", in *Flips and abundance for algebraic threefolds* (J. Kollár, ed.), Astérisque, vol. 211, Société Mathématique de France, 1992, p. 103-114.

[31] C. Spicer, "Higher-dimensional foliated Mori theory", *Compos. Math.* **156** (2020), no. 1, p. 1-38.

[32] C. Spicer & R. Svaldi, "Local and global applications of the Minimal Model Program for co-rank 1 foliations on threefolds", *J. Eur. Math. Soc.* **24** (2022), no. 11, p. 3969-4025.

[33] ———, "Effective generation for foliated surfaces: results and applications", *J. Reine Angew. Math.* **795** (2023), p. 45-84.

[34] C. Spicer, R. Svaldi & S. Velazquez, "On moduli of foliated surfaces", in preparation.

Calum Spicer, Department of Mathematics, King's College London, Strand, London WC2R 2LS, UK
E-mail : calum.spicer@kcl.ac.uk