美国数学会经典影印系列

AMS
AMERICAN MATHEMATICAL SOCIETY

Poincaré's Legacies, Part II

pages from year two of a mathematical blog

庞加莱的遗产，第 II 部分

第二年的数学博客选文

Terence Tao

高等教育出版社

美国数学会经典影印系列
AMS
American Mathematical Society

Poincaré's Legacies, Part II

pages from y[illegible]
of a mathematical blog

庞加莱的遗产，第 II 部分

第二年的数学博客选文

Terence Tao

高等教育出版社·北京

图字：01-2016-2488 号

庞加莱的遗产，第 II 部分

Pangjialai de Yichan, Di II Bufen

图书在版编目 (CIP) 数据

庞加莱的遗产 . 第 II 部分 , 第二年的数学博客选文 = Poincare's Legacies, Part II: pages from year two of a mathematical blog : 英文 / (澳) 陶哲轩 (Terence Tao) 著 . — 影印本 . — 北京 : 高等教育出版社 , 2017.4
ISBN 978-7-04-046996-7
Ⅰ. ①庞… Ⅱ. ①陶… Ⅲ. ①数学—文集—英文
Ⅳ. ①O1-53
中国版本图书馆 CIP 数据核字 (2016) 第 326744 号

策划编辑 李 鹏　　责任编辑 李 鹏
封面设计 张申申　　责任印制 赵义民

出版发行 高等教育出版社
社址 北京市西城区德外大街 4 号
邮政编码 100120
购书热线 010-58581118
咨询电话 400-810-0598
网址 http://www.hep.edu.cn
http://www.hep.com.cn
网上订购 http://www.hepmall.com.cn
http://www.hepmall.com
http://www.hepmall.cn
印刷 北京中科印刷有限公司

开本 787mm × 1092mm 1/16
印张 19.5
字数 490 千字
版次 2017 年 4 月第 1 版
印次 2017 年 4 月第 1 次印刷
定价 135.00 元

[物 料 号 46996-00]

美国数学会经典影印系列

出版者的话

近年来，我国的科学技术取得了长足进步，特别是在数学等自然科学基础领域不断涌现出一流的研究成果。与此同时，国内的科研队伍与国外的交流合作也越来越密切，越来越多的科研工作者可以熟练地阅读英文文献，并在国际顶级期刊发表英文学术文章，在国外出版社出版英文学术著作。

然而，在国内阅读海外原版英文图书仍不是非常便捷。一方面，这些原版图书主要集中在科技、教育比较发达的大中城市的大型综合图书馆以及科研院所的资料室中，普通读者借阅不甚容易；另一方面，原版书价格昂贵，动辄上百美元，购买也很不方便。这极大地限制了科技工作者对于国外先进科学技术知识的获取，间接阻碍了我国科技的发展。

高等教育出版社本着植根教育、弘扬学术的宗旨服务我国广大科技和教育工作者，同美国数学会（American Mathematical Society）合作，在征求海内外众多专家学者意见的基础上，精选该学会近年出版的数十种专业著作，组织出版了“美国数学会经典影印系列”丛书。美国数学会创建于 1888 年，是国际上极具影响力的专业学术组织，目前拥有近 30000 会员和 580 余个机构成员，出版图书 3500 多种，冯·诺依曼、莱夫谢茨、陶哲轩等世界级数学大家都是其作者。本影印系列涵盖了代数、几何、分析、方程、拓扑、概率、动力系统等所有主要数学分支以及新近发展的数学主题。

我们希望这套书的出版，能够对国内的科研工作者、教育工作者以及青年学生起到重要的学术引领作用，也希望今后能有更多的海外优秀英文著作被介绍到中国。

高等教育出版社

2016 年 12 月

To Garth Gaudry, who set me on the road;
To my family, for their constant support;
And to the readers of my blog, for their feedback and contributions.

Preface

In February of 2007, I converted my "What's new" web page of research updates into a blog at `terrytao.wordpress.com`. This blog has since grown and evolved to cover a wide variety of mathematical topics, ranging from my own research updates, to lectures and guest posts by other mathematicians, to open problems, to class lecture notes, to expository articles at both basic and advanced levels.

With the encouragement of my blog readers, and also of the AMS, I published many of the mathematical articles from the first year (2007) of the blog as [**Ta2008b**], which will henceforth be referred to as *Structure and Randomness* throughout this book. This gave me the opportunity to improve and update these articles to a publishable (and citeable) standard, and also to record some of the substantive feedback I had received on these articles by the readers of the blog. Given the success of the blog experiment so far, I am now doing the same for the second year (2008) of articles from the blog. This year, the amount of material is large enough that the blog will be published in two volumes.

As with *Structure and Randomness*, each part begins with a collection of expository articles, ranging in level from completely elementary logic puzzles to remarks on recent research, which are only loosely related to each other and to the rest of the book. However, in contrast to the previous book, the bulk of these volumes is dominated by the lecture notes for two graduate courses I gave during the year. The two courses stemmed from two very different but fundamental contributions to mathematics by Henri Poincaré, which explains the title of the book.

This is the second of the two volumes, and it focuses on geometry, topology, and partial differential equations. In particular, Chapter 2 contains the lecture notes for my course on the famous *Poincaré conjecture* that every simply connected compact three-dimensional manifold is homeomorphic to a sphere, and its recent spectacular solution **[Pe2002]**, **[Pe2003]**, **[Pe2003b]** by Perelman. This conjecture is purely topological in nature, and yet Perelman's proof uses remarkably little topology, instead working almost entirely in the realm of Riemannian geometry and partial differential equations, and specifically in a detailed analysis of solutions to Ricci flows on three-dimensional manifolds, and the singularities formed by these flows. As such, the course will incorporate, along the way, a review of many of the basic concepts and results from Riemannian geometry (and to a lesser extent, from parabolic PDE), while being focused primarily on the single objective of proving the Poincaré conjecture. Due to the complexity and technical intricacy of the argument, we will not be providing a fully complete proof of this conjecture here (see **[MoTi2007]** for a careful and detailed treatment); but we will be able to cover the high-level features of the argument, as well as many of the specific components of that argument, in full detail, and the remaining components are sketched and motivated, with references to more complete arguments given. In principle, the course material is sufficiently self-contained that prior exposure to Riemannian geometry, PDE, or topology at the graduate level is not strictly necessary, but in practice, one would probably need some comfort with at least one of these three areas in order to not be totally overwhelmed by the material. (I ran this course as a topics course; in particular, I did not assign homework.)

A remark on notation

For reasons of space, we will not be able to define every single mathematical term that we use in this book. If a term is italicised for reasons other than emphasis or definition, then it denotes a standard mathematical object, result, or concept, which can be easily looked up in any number of references. (In the blog version of the book, many of these terms were linked to their Wikipedia pages, or other on-line reference pages.)

I will however mention a few notational conventions that I will use throughout. The cardinality of a finite set E will be denoted $|E|$. We will use the asymptotic notation $X = O(Y)$, $X \ll Y$, or $Y \gg X$ to denote the estimate $|X| \leq CY$ for some absolute constant $C > 0$. In some cases we will need this constant C to depend on a parameter (e.g. d), in which case we shall indicate this dependence by subscripts, e.g. $X = O_d(Y)$ or $X \ll_d Y$. We also sometimes use $X \sim Y$ as a synonym for $X \ll Y \ll X$.

In many situations there will be a large parameter n that goes off to infinity. When that occurs, we also use the notation $o_{n\to\infty}(X)$ or simply $o(X)$ to denote any quantity bounded in magnitude by $c(n)X$, where $c(n)$

is a function depending only on n that goes to zero as n goes to infinity. If we need $c(n)$ to depend on another parameter, e.g. d, we indicate this by further subscripts, e.g. $o_{n\to\infty;d}(X)$.

We will occasionally use the averaging notation

$$\mathbf{E}_{x\in X} f(x) := \frac{1}{|X|}\sum_{x\in X} f(x)$$

to denote the average value of a function $f : X \to \mathbf{C}$ on a non-empty finite set X.

Acknowledgments

The author is supported by a grant from the MacArthur Foundation, by NSF grant DMS-0649473, and by the NSF Waterman award. The author is also greatly indebted to Peter Petersen and William Wylie for many useful discussions on Ricci flow, which were instrumental in the preparation of the lecture notes on this topic.

Contents

Chapter 1

Expository Articles

1.1. Dvir's proof of the finite field Kakeya conjecture

One of my favourite unsolved problems in mathematics is the *Kakeya conjecture* in geometric measure theory. This conjecture is descended from the following question, posed by Soichi Kakeya in 1917:

Problem 1.1.1 (Kakeya needle problem). What is the least area in the plane required to continuously rotate a needle of unit length and zero thickness around completely (i.e., by 360°)?

For instance, one can rotate a unit needle inside a unit disk, which has area $\pi/4$. By using a *deltoid* one requires only $\pi/8$ area.

In [**Be1919**], [**Be1928**], Besicovitch showed that in fact one could rotate a unit needle using an *arbitrarily small* amount of positive area. This unintuitive fact was a corollary of two observations. The first, which is easy, is that one can *translate* a needle using arbitrarily small area, by sliding the needle along the direction it points in for a long distance (which costs zero area), turning it slightly (costing a small amount of area), sliding back, and then undoing the turn. The second fact, which is less obvious, can be phrased as follows. Define a *Kakeya set* in $\mathbf{R}^2$ to be any set which contains a unit line segment in each direction.

Theorem 1.1.2 ([**Be1919**]). *There exist Kakeya sets in $\mathbf{R}^2$ of arbitrarily small area (or more precisely, Lebesgue measure).*

In fact, one can construct such sets with zero Lebesgue measure. On the other hand, it was shown by Davies [**Da1971**] that even though these sets had zero area, they were still necessarily two-dimensional (in the sense of either Hausdorff or Minkowski dimension). This led to an analogous conjecture in higher dimensions:

Conjecture 1.1.3 (Kakeya conjecture). *A Besicovitch set in* $\mathbf{R}^n$ *(i.e., a subset of* $\mathbf{R}^n$ *that contains a unit line segment in every direction) has Minkowski and Hausdorff dimension equal to* n.

This conjecture remains open in dimensions three and higher (and gets more difficult as the dimension increases), although many partial results are known. For instance, if $n = 3$, it is known that Besicovitch sets have *Hausdorff dimension* at least $5/2$ (see [**Wo1995**]) and *upper Minkowski dimension* at least $5/2 + 10^{-10}$ (see [**KaLaTa2000**]). See also the surveys [**Ta2001**], [**KaTa2002**], [**Wo1999**].

In [**Wo1999**], Wolff proposed a simpler finite field analogue[1] of the Kakeya conjecture as a model problem that avoided all the technical issues involving Minkowski and Hausdorff dimension. If F^n is a vector space over a finite field F, define a *Kakeya set* to be a subset of F^n which contains a line in every direction.

Conjecture 1.1.4 (Finite field Kakeya conjecture). *Let* $E \subset F^n$ *be a Kakeya set. Then* E *has cardinality at least* $c_n|F|^n$, *where* $c_n > 0$ *depends only on* n.

This conjecture has had a significant influence in the subject, in particular inspiring work on the *sum-product phenomenon* in finite fields, which has since proven to have many applications in number theory and computer science. Modulo minor technicalities, the progress on the finite field Kakeya conjecture was, until very recently, essentially the same as that of the original "Euclidean" Kakeya conjecture.

Recently, the finite field Kakeya conjecture was proven using a beautifully simple argument by Dvir [**Dv2008**], based on the *polynomial method* in algebraic extremal combinatorics. The proof is so short that I can present it in full here.

The polynomial method is used to control the size of various sets E by looking at one or more polynomials P which vanish on that set E. This philosophy of course closely resembles that of algebraic geometry, and indeed one could classify the polynomial method as a kind of "combinatorial algebraic geometry". An important difference, though, is that in the combinatorial setting we work over fields that are definitely *not* algebraically

[1] Cf. Section 1.6 of *Structure and Randomness*.

closed; in particular, we are primarily interested in polynomials[2] and their zero sets over *finite* fields.

For instance, in high school we learn the following connection between one-dimensional sets E and polynomials $P(x)$ in one variable:

Theorem 1.1.5 (Factor theorem). *Let F be a field, and $d \geq 1$ an integer. Let $F[x]$ denote the polynomials in one variable with coefficients in F.*

(1) *If $P \in F[x]$ is a non-zero polynomial of degree at most d, then the set $\{x \in F : P(x) = 0\}$ has cardinality at most d.*
(2) *Conversely, given any set $E \subset F$ of cardinality at most d, there exists a non-zero polynomial $P \in F[x]$ of degree at most d that vanishes on E.*

Thus, to obtain an *upper bound* on the size of a one-dimensional set E, it would suffice to exhibit a non-zero low-degree polynomial that vanishes on E; conversely, to obtain a *lower bound* on the size of E, one would have to show that the only low-degree polynomial that vanishes on E is the zero polynomial. It is the latter type of observation which is of relevance to the finite field Kakeya problem.

There are analogues of both (1) and (2) in higher dimensions. For instance, the *Schwartz-Zippel lemma* [**Sc1980**] is a higher-dimensional analogue of (1), as is the combinatorial Nullstellensatz of Alon [**Al1999**] and *Bézout's theorem* from algebraic geometry, while Stepanov's method [**St1969**] exploits a higher-dimensional analogue of (2). These sorts of techniques and results are collectively referred to as the *polynomial method* in extremal algebraic combinatorics. For Dvir's argument, we will need a very simple higher-dimensional version of (2) that comes from basic linear algebra, namely

Lemma 1.1.6. *Let $E \subset F^n$ be a set of cardinality less than $\binom{n+d}{n}$ for some $d \geq 0$. Then there exists a non-zero polynomial $P \in F[x_1, \ldots, x_n]$ in n variables of degree at most d which vanishes on E.*

Proof. Let V be the vector space of polynomials in $F[x_1, \ldots, x_n]$ of degree at most d. Elementary combinatorics reveals that V has dimension $\binom{n+d}{n}$. On the other hand, the vector space F^E of F-valued functions on E has dimension $|E| < \binom{n+d}{d}$. Hence the evaluation map $P \mapsto (P(x))_{x \in E}$ from V to F^E is non-injective, and the claim follows. □

Dvir's argument combines this lemma with the following proposition.

[2] Also, whereas algebraic geometry is more concerned with *specific* (and often highly structured) polynomials, the polynomial method requires that one consider rather *generic* (and usually quite high degree) polynomials.

Proposition 1.1.7. *Let $P \in F[x_1, \dots, x_n]$ be a polynomial of degree at most $|F| - 1$ which vanishes on a Kakeya set E. Then P is identically zero.*

Proof. Suppose for contradiction that P is non-zero. We can write $P = \sum_{i=0}^{d} P_i$, where $0 \leq d \leq |F| - 1$ is the degree of P and P_i is the i^{th} homogeneous component, thus P_d is non-zero. Since P vanishes on E, d cannot be zero.

Let $v \in F^n \backslash \{0\}$ be an arbitrary direction. As E is a Kakeya set, E contains a line $\{x + tv : t \in F\}$ for some $x = x_v \in F^n$, thus $P(x + tv) = 0$ for all $t \in F$. The left-hand side is a polynomial in t of degree at most $|F| - 1$, and thus vanishes identically by the factor theorem. In particular, the t^d coefficient of this polynomial, which is $P_d(v)$, vanishes for any non-zero v. Since P_d is homogeneous of degree $d > 0$, P_d vanishes on all of F^n. Since P_d also has degree less than $|F|$, repeated application of the factor theorem for each variable in turn (or the Schwartz-Zippel lemma [**Sc1980**], which is much the same thing) shows that $P_d = 0$, a contradiction. □

Remark 1.1.8. The point here is that a low-degree polynomial which vanishes on a line must also vanish at the point at infinity where the line touches the hyperplane at infinity. Thus a polynomial which vanishes on a Kakeya set vanishes on the entire hyperplane at infinity. One can then divide out the defining polynomial for that hyperplane and repeat the process to conclude that the polynomial vanishes identically.

Combining the lemma and the proposition we obtain

Corollary 1.1.9. *Every Kakeya set in F^n has cardinality at least $\binom{|F|+n-1}{n}$.*

Since $\binom{|F|+n-1}{n} = \frac{1}{n!}|F|^n + O_n(|F|^{n-1})$, this establishes the finite field Kakeya conjecture.

This bound seems to be quite tight. For instance, it gives the lower bound of $\frac{|F|(|F|+1)}{2}$ for Kakeya sets in F^2 (which was already implicitly observed by Wolff); this is very close to the exact bound, which was recently established in [**Ba2008**], [**BlMa2008**] to be $\frac{|F|(|F|+1)}{2} + \frac{|F|-1}{2}|F|$ in the case when $|F|$ is odd. (Thanks to Simeon Ball and Francesco Mazzocca for these references.)

It now seems sensible to revisit other problems in extremal combinatorics over finite fields to see if the polynomial method can yield results there. Certainly close relatives of the Kakeya conjecture (e.g. the Nikodým set conjecture, or the Kakeya maximal function conjecture) should now be establishable by these methods. On the other hand, there are other problems (such as the sum-product problem, Szemerédi-Trotter type theorems, and distance set problems) which are sensitive to the choice of field F (and in particular, whether that field contains a subfield of index 2); see [**BoKaTa2004**]. It would be interesting to see if there are ways to adapt the polynomial method in order to detect the existence of subfields.

Very recently, the polynomial method has also been extended to yield some progress on the Euclidean case; see Section 1.7.

Notes. This article first appeared at

`terrytao.wordpress.com/2008/03/24`.

Thanks to ninguem for corrections.

Seva posed the question of determining the asymptotic best density for a Kakeya set in, say, F_3^n, as $n \to \infty$.

Some pictures of Kakeya sets can be found at

`www.math.ucla.edu/ tao/java/Besicovitch.html`

or `en.wikipedia.org/wiki/Kakeya_set`.

Further discussion of Dvir's result can be found online at

`ilaba.wordpress.com/2008/03/23`

and `quomodocumque.wordpress.com/2008/03/25`.

1.2. The Black-Scholes equation

In this article I would like to describe the mathematical derivation of the famous *Black-Scholes equation* in financial mathematics, at least in the simplified case in which time is discrete. This simplified model avoids many of the technicalities involving stochastic calculus, Itô's formula, etc., and brings the beautifully simple basic idea behind the derivation of this formula into focus.

The basic type of problem that the Black-Scholes equation solves (in particular models) is the following. One has an *underlying financial instrument* S, which represents some asset[3] which can be bought and sold at various times t, with the per-unit price S_t of the instrument varying with t. Given such an underlying instrument S, one can create *options* based on S and on some future time t_1, which give the buyer and seller of the options certain rights and obligations regarding S at an *expiration time* t_1. For instance,

(1) A *call option* for S at time t_1 and at a *strike price* P gives the buyer of the option the right (but not the obligation) to buy a unit of S from the seller of the option at price P at time t_1 (conversely, the seller of the option has the obligation but not the right to sell

[3]For the mathematical model, it is not relevant what type of asset S actually is, but one could imagine for instance that S is a stock, a commodity, a currency, or a bond.

a unit of S to the buyer of the option at time t_1, if the buyer so requests).

(2) A *put option* for S at time t_1 and at a strike price P gives the buyer of the option the right (but not the obligation) to sell a unit of S to the seller of the option at price P at time t_1 (and conversely, the seller of the option has the obligation but not the right to buy a unit of S from the buyer of the option at time t_1, if the buyer so requests).

(3) More complicated options, such as *straddles* and *collars*, can be formed by taking linear combinations of call and put options, e.g. simultaneously buying or selling a call and a put option. One can also consider "American options" which offer rights and obligations for an interval of time, rather than the "European options" described above which only apply at a fixed time t_1. The Black-Scholes formula applies only to European options, though extensions of this theory have been applied to American options.

The problem is this: what is the "correct" price, at time t_0, to assign to a European option (such as a put or call option) at a future expiration time t_1? Of course, due to the volatility of the underlying instrument S, the future price S_{t_1} of this instrument is not known at time t_0. Nevertheless—and this is really quite a remarkable fact—it is still possible to compute deterministically, at time t_0, the price of an option that depends on that unknown price S_{t_1}, under certain assumptions (one of which is that one knows exactly *how* volatile the underlying instrument is).

1.2.1. How to compute price. Before we do any mathematics, we must first settle a fundamental financial question—how can one compute the price of some asset A? In most economic situations, such a price would depend on many factors, such as the supply and demand of A, transaction costs in buying or selling A, legal regulations concerning A, or more intangible factors such as the current market sentiment regarding A. Any model that attempted to accurately describe all of these features would be hideously complicated and involve a large number of parameters that would be nearly impossible to measure directly. So, in general, one cannot hope to compute such prices mathematically.

But the situation is much simpler for purely financial products, such as options, at least when one has a highly deep and liquid market for the underlying instrument S. More precisely, we will make the following (unrealistic) assumptions.