美国数学会经典影印系列

AMS
American Mathematical Society

An Epsilon of Room, I: Real Analysis

pages from year three of a mathematical blog

ε 空间，I: 实分析 第三年的数学博客选文

Terence Tao

高等教育出版社

美国数学会经典影印系列

AMS
AMERICAN MATHEMATICAL SOCIETY

An Epsilon of Room, I: Real Analysis

pages from year three
of a mathematical blog

ε 空间, I: 实分析 第三年的数学博客选文

Terence Tao

高等教育出版社·北京

图字：01-2016-2518 号

ε 空间 . I, 实分析
ε Kongjian. I, Shifenxi

图书在版编目 (CIP) 数据

ε 空间 . I, 实分析 : 第三年的数学博客选文 = An Epsilon of Room, I: Real Analysis: pages from year three 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 , 2017.1
ISBN 978-7-04-046900-4
Ⅰ. ① ε··· Ⅱ. ①陶··· Ⅲ. ①实分析—文集—英文 Ⅳ. ①O174.1-53
中国版本图书馆 CIP 数据核字 (2016) 第 281730 号

策划编辑 赵天夫　　责任编辑 赵天夫
封面设计 张申申　　责任印制 毛斯璐

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 23
字数 490千字
版次 2017 年 1 月第 1 版
印次 2017 年 1 月第 1 次印刷
定价 135.00元

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 two years of the blog as [**Ta2008**] and [**Ta2009**], which will henceforth be referred to as *Structure and Randomness* and *Poincaré's Legacies Vols. I, II.* 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.

The current text contains many (though not all) of the posts for the third year (2009) of the blog, focusing primarily on those posts of a mathematical nature which were not contributed primarily by other authors, and which are not published elsewhere. It has been split into two volumes.

The current volume consists of lecture notes from my graduate real analysis courses that I taught at UCLA (Chapter 1), together with some related material in Chapter 2. These notes cover the second part of the graduate real analysis sequence here, and therefore assume some familiarity with general measure theory (in particular, the construction of Lebesgue measure and the Lebesgue integral, and more generally the material reviewed in Section 1.1), as well as undergraduate real analysis (e.g., various notions of limits and convergence). The notes then cover more advanced topics in measure theory (notably, the Lebesgue-Radon-Nikodym and Riesz representation theorems) as well as a number of topics in functional analysis, such as the theory of Hilbert and Banach spaces, and the study of key function spaces such as the Lebesgue and Sobolev spaces, or spaces of distributions. The general theory of the Fourier transform is also discussed. In addition, a number of auxiliary (but optional) topics, such as Zorn's lemma, are discussed in Chapter 2. In my own course, I covered the material in Chapter 1 only and also used Folland's text [**Fo2000**] as a secondary source. But I hope that the current text may be useful in other graduate real analysis courses, particularly in conjunction with a secondary text (in particular, one that covers the prerequisite material on measure theory).

The second volume in this series (referred to henceforth as *Volume II*) consists of sundry articles on a variety of mathematical topics, which is only occasionally related to the above course, and can be read independently.

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 for 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 online 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.

Thanks to Kestutis Cesnavicius, Wolfgang M., Daniel Mckenzie, Simion, Snegud, Blake Stacey, Konrad Swanepoel, and anonymous commenters for global corrections to the text, and to Edward Dunne at the American Mathematical Society for encouragement and editing.

Contents

Chapter 1

Real analysis

A quick review of measure and integration theory

In this section we quickly review the basics of abstract measure theory and integration theory, which was covered in the previous course but will of course be relied upon in the current course. This is only a brief summary of the material; certainly, one should consult a real analysis text for the full details of the theory.

1.1.1. Measurable spaces. Ideally, measure theory on a space X should be able to assign a measure (or *volume*, or *mass*, etc.) to every set in X. Unfortunately, due to paradoxes such as the *Banach-Tarski paradox*, many natural notions of measure (e.g., *Lebesgue measure*) cannot be applied to measure all subsets of X; instead, we must restrict our attention to certain measurable subsets of X. This turns out to suffice for most applications; for instance, just about any *non-pathological* subset of Euclidean space that one actually encounters will be Lebesgue measurable (as a general rule of thumb, any set which does not rely on the axiom of choice in its construction will be measurable).

To formalise this abstractly, we use

Definition 1.1.1 (Measurable spaces). A *measurable space* $(X, \mathcal{X})$ is a set X, together with a collection $\mathcal{X}$ of subsets of X which form a σ-algebra, thus $\mathcal{X}$ contains the empty set and X, and is closed under countable intersections, countable unions, and complements. A subset of X is said to be measurable with respect to the measurable space if it lies in $\mathcal{X}$.

A function $f : X \to Y$ from one measurable space $(X, \mathcal{X})$ to another $(Y, \mathcal{Y})$ is said to be *measurable* if $f^{-1}(E) \in \mathcal{X}$ for all $E \in \mathcal{Y}$.

Remark 1.1.2. The class of measurable spaces forms a *category*, with the measurable functions being the *morphisms*. The symbol σ stands for *countable union*; cf. σ-compact, σ-finite, F_σ set.

Remark 1.1.3. The notion of a measurable space $(X, \mathcal{X})$ (and of a measurable function) is superficially similar to that of a *topological space* $(X, \mathcal{F})$ (and of a *continuous function*); the topology $\mathcal{F}$ contains $\emptyset$ and X just as the σ-algebra $\mathcal{X}$ does, but is now closed under arbitrary unions and finite intersections, rather than countable unions, countable intersections, and complements. The two categories are linked to each other by the Borel algebra construction; see Example 1.1.5 below.

Example 1.1.4. We say that one σ-algebra $\mathcal{X}$ on a set X is *coarser* than another $\mathcal{X}'$ (or that $\mathcal{X}'$ is *finer* than $\mathcal{X}$) if $\mathcal{X} \subset \mathcal{X}'$ (or equivalently, if the identity map from $(X, \mathcal{X}')$ to $(X, \mathcal{X})$ is measurable); thus every set which is measurable in the coarse space is also measurable in the fine space. The coarsest σ-algebra on a set X is the trivial σ-algebra $\{\emptyset,\mathrm{X}\}$, while the finest is the discrete σ-algebra $2^X := \{E : E \subset \mathrm{X}\}$.

Example 1.1.5. The intersection $\bigwedge_{\alpha \in A} \mathcal{X}_\alpha := \bigcap_{\alpha \in A} \mathcal{X}_\alpha$ of an arbitrary family $(\mathcal{X}_\alpha)_{\alpha \in A}$ of σ-algebras on X is another σ-algebra on X. Because of this, given any collection $\mathcal{F}$ of sets on X we can define the σ-algebra $\mathcal{B}[\mathcal{F}]$ *generated by* $\mathcal{F}$, defined to be the intersection of all the σ-algebras containing $\mathcal{F}$, or equivalently the coarsest algebra for which all sets in $\mathcal{F}$ are measurable. (This intersection is non-vacuous, since it will always involve the discrete σ-algebra 2^X.) In particular, the open sets $\mathcal{F}$ of a topological space $(X, \mathcal{F})$ generate a σ-algebra, known as the *Borel σ-algebra* of that space.

We can also define the *join* $\bigvee_{\alpha \in A} \mathcal{X}_\alpha$ of any family $(\mathcal{X}_\alpha)_{\alpha \in A}$ of σ-algebras on X by the formula

$$\bigvee_{\alpha \in A} \mathcal{X}_\alpha := \mathcal{B}[\bigcup_{\alpha \in A} \mathcal{X}_\alpha]. \tag{1.1}$$

For instance, the *Lebesgue σ-algebra* $\mathcal{L}$ of Lebesgue measurable sets on a Euclidean space $\mathbf{R}^n$ is the join of the Borel σ-algebra $\mathcal{B}$ and of the algebra of null sets and their complements (also called *co-null* sets).

Exercise 1.1.1. A function $f : X \to Y$ from one topological space to another is said to be *Borel measurable* if it is measurable once X and Y are equipped with their respective Borel σ-algebras. Show that every continuous function is Borel measurable. (The converse statement, of course, is very far from being true; for instance, the pointwise limit of a sequence of measurable functions, if it exists, is also measurable, whereas the analogous claim for continuous functions is completely false.)

Remark 1.1.6. A function $f : \mathbf{R}^n \to \mathbf{C}$ is said to be *Lebesgue measurable* if it is measurable from $\mathbf{R}^n$ (with the Lebesgue σ-algebra) to $\mathbf{C}$ (with the Borel σ-algebra), or equivalently if $f^{-1}(B)$ is Lebesgue measurable for every open ball B in $\mathbf{C}$. Note the asymmetry between Lebesgue and Borel here; in particular, the composition of two Lebesgue measurable functions need not be Lebesgue measurable.

Example 1.1.7. Given a function $f : X \to Y$ from a set X to a measurable space $(Y, \mathcal{Y})$, we can define the *pullback* $f^{-1}(\mathcal{Y})$ of $\mathcal{Y}$ to be the σ-algebra $f^{-1}(\mathcal{Y}) := \{f^{-1}(E) : E \in \mathcal{Y}\}$; this is the coarsest structure on X that makes f measurable. For instance, the pullback of the Borel σ-algebra from $[0,1]$ to $[0,1]^2$ under the map $(x,y) \mapsto x$ consists of all sets of the form $E \times [0,1]$, where $E \subset [0,1]$ is Borel measurable.

More generally, given a family $(f_\alpha : X \to Y_\alpha)_{\alpha \in A}$ of functions into measurable spaces $(Y_\alpha, \mathcal{Y}_\alpha)$, we can form the σ-algebra $\bigvee_{\alpha \in A} f_\alpha^{-1}(\mathcal{Y}_\alpha)$ generated by the f_α; this is the coarsest structure on X that makes all the f_α simultaneously measurable.

Remark 1.1.8. In probability theory and information theory, the functions $f_\alpha : X \to Y_\alpha$ in Example 1.1.7 can be interpreted as *observables*, and the σ-algebra generated by these observables thus captures mathematically the concept of observable information. For instance, given a time parameter t, one might define the σ-algebra $\mathcal{F}_{\leq t}$ generated by all observables for some random process (e.g., *Brownian motion*) that can be made at time t or earlier; this endows the underlying event space X with an uncountable increasing family of σ-algebras.

Example 1.1.9. If E is a subset of a measurable space $(Y, \mathcal{Y})$, the pullback of $\mathcal{Y}$ under the inclusion map $\iota : E \to Y$ is called the *restriction* of $\mathcal{Y}$ to E and is denoted $\mathcal{Y}\lfloor_E$. Thus, for instance, we can restrict the Borel and Lebesgue σ-algebras on a Euclidean space $\mathbf{R}^n$ to any subset of such a space.

Exercise 1.1.2. Let M be an n-dimensional manifold, and let $(\pi_\alpha : U_\alpha \to V_\alpha)$ be an atlas of coordinate charts for M, where U_α is an open cover of M and V_α are open subsets of $\mathbf{R}^n$. Show that the Borel σ-algebra on M is the unique σ-algebra whose restriction to each U_α is the pullback via π_α of the restriction of the Borel σ-algebra of $\mathbf{R}^n$ to V_α.

Example 1.1.10. A function $f : X \to A$ into some index set A will partition X into level sets $f^{-1}(\{\alpha\})$ for $\alpha \in A$; conversely, every partition $X = \bigcup_{\alpha \in A} E_\alpha$ of X arises from at least one function f in this manner (one can just take f to be the map from points in X to the partition cell in which that point lies). Given such an f, we call the σ-algebra $f^{-1}(2^A)$ the σ-algebra *generated by* the partition; a set is measurable with respect to this structure if and only if it is the union of some subcollection $\bigcup_{\alpha\in B} E_\alpha$ of cells of the partition.

Exercise 1.1.3. Show that a σ-algebra on a finite set X necessarily arises from a partition $X = \bigcup_{\alpha\in A} E_\alpha$ as in Example 1.1.10, and furthermore the partition is unique (up to relabeling). Thus, in the finitary world, σ-algebras are essentially the same concept as partitions.

Example 1.1.11. Let $(X_\alpha, \mathcal{X}_\alpha)_{\alpha\in A}$ be a family of measurable spaces, then the Cartesian product $\prod_{\alpha\in A} X_\alpha$ has canonical projection maps $\pi_\beta : \prod_{\alpha\in A} X_\alpha \to X_\beta$ for each $\beta \in A$. The product σ-algebra $\prod_{\alpha\in A} \mathcal{X}_\alpha$ is defined as the σ-algebra on $\prod_{\alpha\in A} X_\alpha$ generated by the π_α, as in Example 1.1.7.

Exercise 1.1.4. Let $(X_\alpha)_{\alpha\in A}$ be an at most countable family of second countable topological spaces. Show that the Borel σ-algebra of the product space (with the product topology) is equal to the product of the Borel σ-algebras of the factor spaces. In particular, the Borel σ-algebra on $\mathbf{R}^n$ is the product of n copies of the Borel σ-algebra on $\mathbf{R}$. (The claim can fail when the countability hypotheses are dropped, though in most applications in analysis, these hypotheses are satisfied.) We caution however that the Lebesgue σ-algebra on $\mathbf{R}^n$ is not the product of n copies of the one-dimensional Lebesgue σ-algebra, as it contains some additional null sets; however, it is the completion of that product.

Exercise 1.1.5. Let $(X, \mathcal{X})$ and $(Y, \mathcal{Y})$ be measurable spaces. Show that if E is measurable with respect to $\mathcal{X} \times \mathcal{Y}$, then for every $x \in X$, the set $\{y \in Y : (x, y) \in E\}$ is measurable in $\mathcal{Y}$, and similarly for every $y \in Y$, the set $\{x \in X : (x, y) \in E\}$ is measurable in $\mathcal{X}$. Thus, sections of Borel measurable sets are again Borel measurable. (The same is not true for Lebesgue measurable sets.)

1.1.2. Measure spaces. Now we endow measurable spaces with a measure, turning them into measure spaces.

Definition 1.1.12 (Measures). A (non-negative) *measure* μ on a measurable space $(X, \mathcal{X})$ is a function $\mu : \mathcal{X} \to [0, +\infty]$ such that $\mu(\emptyset) = 0$, and such that we have the countable additivity property $\mu(\bigcup_{n=1}^\infty E_n) = \sum_{n=1}^\infty \mu(E_n)$ whenever $E_1, E_2, \ldots$ are disjoint measurable sets. We refer to the triplet $(X, \mathcal{X}, \mu)$ as a *measure space*.

A measure space $(X, \mathcal{X}, \mu)$ is *finite* if $\mu(X) < \infty$; it is a *probability space* if $\mu(X) = 1$ (and then we call μ a *probability measure*). It is *σ-finite* if X can be covered by countably many sets of finite measure.

A measurable set E is a *null set* if $\mu(E) = 0$. A property on points x in X is said to hold for *almost every* $x \in X$ (or *almost surely*, for probability spaces) if it holds outside of a null set. We abbreviate "almost every" and "almost surely" as a.e. and a.s., respectively. The complement of a null set is said to be a *co-null set* or to have *full measure*.

Example 1.1.13 (Dirac measures). Given any measurable space $(X, \mathcal{X})$ and a point $x \in X$, we can define the *Dirac measure* (or *Dirac mass*) δ_x to be the measure such that $\delta_x(E) = 1$ when $x \in E$ and $\delta_x(E) = 0$, otherwise. This is a probability measure.

Example 1.1.14 (Counting measure). Given any measurable space $(X, \mathcal{X})$, we define *counting measure* $\#$ by defining $\#(E)$ to be the cardinality $|E|$ of E when E is finite, or $+\infty$ otherwise. This measure is finite when X is finite, and σ-finite when X is at most countable. If X is also finite, we can define *normalised counting measure* $\frac{1}{|E|}\#$; this is a probability measure, also known as the *uniform probability measure* on X (especially if we give X the discrete σ-algebra).

Example 1.1.15. Any finite non-negative linear combination of measures is again a measure; any finite convex combination of probability measures is again a probability measure.

Example 1.1.16. If $f : X \to Y$ is a measurable map from one measurable space $(X, \mathcal{X})$ to another $(Y, \mathcal{Y})$, and μ is a measure on $\mathcal{X}$, we can define the *push-forward* $f_*\mu : \mathcal{Y} \to [0, +\infty]$ by the formula $f_*\mu(E) := \mu(f^{-1}(E))$; this is a measure on $(Y, \mathcal{Y})$. Thus, for instance, $f_*\delta_x = \delta_{f(x)}$ for all $x \in X$.

We record some basic properties of measures of sets:

Exercise 1.1.6. Let $(X, \mathcal{X}, \mu)$ be a measure space. Show the following statements:

(i) *Monotonicity.* If $E \subset F$ are measurable sets, then $\mu(E) \leq \mu(F)$. (In particular, any measurable subset of a null set is again a null set.)

(ii) *Countable subadditivity.* If $E_1, E_2, \ldots$ are a countable sequence of measurable sets, then $\mu(\bigcup_{n=1}^\infty E_n) \leq \sum_{n=1}^\infty \mu(E_n)$. (Of course, one also has subadditivity for finite sequences.) In particular, any countable union of null sets is again a null set.

(iii) *Monotone convergence for sets.* If $E_1 \subset E_2 \subset \cdots$ are measurable, then $\mu(\bigcup_{n=1}^\infty E_n) = \lim_{n\to\infty} \mu(E_n)$.

(iv) *Dominated convergence for sets.* If $E_1 \supset E_2 \supset \cdots$ are measurable, and $\mu(E_1)$ is finite, then $\mu(\bigcap_{n=1}^\infty E_n) = \lim_{n\to\infty} \mu(E_n)$. Show that the claim can fail if $\mu(E_1)$ is infinite.