

Gaussian Processes For Machine Learning

Gaussian Processes For Machine Learning

Downloaded from blog.amf.com by guest

GAUSSIAN PROCESSES FOR MACHINE LEARNING SUMMARY COLLECTION: OPEN THE SIGNIFICANCE IN BITE-SIZED CHUNKS

Welcome to our captivating book summary collection. We are excited to introduce you to the globe of Gaussian Processes For Machine Learning recaps and just how they can improve your reading experience. As enthusiastic readers ourselves, we comprehend the worth of diving into the heart of every story and uncovering its essence in bite-sized chunks.

Gaussian Processes For Machine Learning book recap collection offers simply that - a concise and interesting recap of the key points and themes of a publication. In today's fast-paced world, we understand that time is valuable, and our summaries are developed to conserve you time by offering a quick overview of Gaussian Processes For Machine Learning's content and insights.

Our team of expert writers carefully curates our publication recap of Gaussian Processes For Machine Learning collection to make certain that we give you with top quality summaries that catch the significance of each book. Whether you are looking to explore brand-new categories, find brand-new authors, or merely obtain deeper insights right into your favorite publications, our collection has something for every person.

Join us today and unlock the world of Gaussian Processes For Machine Learning summaries. Discover the benefits of condensing intricate concepts right into basic and easy-to-understand language. Our publication summaries are a wonderful way to increase your expertise and widen your perspectives without needing to invest hours of your time.

Remain tuned as we explore the idea of Gaussian Processes For Machine Learning, discuss their advantages, and give ideas on just how to create effective recaps. With our aid, you'll find the ideal book for your passions and unlock a globe of expertise.

DISCOVERING PUBLICATION RECAPS OF GAUSSIAN PROCESSES FOR MACHINE LEARNING

Machine Learning and Knowledge Discovery in Databases Cambridge University Press

From an engineering standpoint, the increasing complexity of robotic systems and the increasing demand for more autonomously learning robots, has become essential. This book is largely based on the successful workshop "From motor to interaction learning in robots" held at the IEEE/RSJ International Conference on Intelligent Robot Systems. The major aim of the book is to give students interested the topics described above a chance to get started faster and researchers a helpful compandium.

OECD Studies on Water Groundwater Allocation Managing Growing Pressures on Quantity and Quality Springer Nature

The three volume set LNAI 9284, 9285, and 9286 constitutes the refereed proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases, ECML PKDD 2015, held in Porto, Portugal, in September 2015. The 131 papers presented in these proceedings were carefully reviewed and selected from a total of 483 submissions. These include 89 research papers, 11 industrial papers, 14 nectar papers, 17 demo papers. They were organized in topical sections named: classification, regression and supervised learning; clustering and unsupervised learning; data preprocessing; data streams and online learning; deep learning; distance and metric learning; large scale learning and big data; matrix and tensor analysis; pattern and sequence mining; preference learning and label ranking; probabilistic, statistical, and graphical approaches; rich data; and social and graphs. Part III is structured in industrial track, nectar track, and demo track.

Artificial Neural Networks and Machine Learning - ICANN 2011 Springer

This volume contains selected tutorial and young scientist school papers of the 5th RAAI Summer School on Artificial Intelligence, held in July 2019 at Institute of Physics and Technology (MIPT) campus in Dolgoprudny, a suburb of Moscow, Russia. The 11 chapters in this volume present papers focusing on various important aspects of Multiagent systems; Behavior planning; Natural language processing; Modeling of reasoning; and Machine learning and data analysis.

Bayesian Learning for Neural Networks Courier Corporation

This New York Times bestseller "elegantly weaves evidence and insights . . . into a single, accessible historical narrative" (Bill Gates) and presents a captivating history of the universe -- from the Big Bang to dinosaurs to mass globalization and beyond. Most historians study the smallest slivers of time, emphasizing specific dates, individuals, and documents. But what would it look like to study the whole of history, from the big bang through the present day -- and even into the remote future? How would looking at the full span of time change the way we perceive the universe, the earth, and our very existence? These were the questions David Christian set out to answer when he created the field of "Big History," the most exciting new approach to understanding where we have been, where we are, and where we are going. In *Origin Story*, Christian takes readers on a wild ride through the entire 13.8 billion years we've come to know as "history." By focusing on defining events (thresholds), major trends, and profound questions about our origins, Christian exposes the hidden threads that tie everything together -- from the creation of the planet to the advent of agriculture, nuclear war, and beyond. With stunning insights into the origin of the universe, the beginning of life, the emergence of humans, and what the future might bring, *Origin Story* boldly reframes our place

in the cosmos.

Advanced Lectures on Machine Learning Springer Nature

Content Description. #Includes bibliographical references and index.

Machine Learning and Knowledge Discovery in Databases Cambridge University Press

This book introduces Bayesian reasoning and Gaussian processes into machine learning applications. Bayesian methods are applied in many areas, such as game development, decision making, and drug discovery. It is very effective for machine learning algorithms in handling missing data and extracting information from small datasets. Bayesian Reasoning and Gaussian Processes for Machine Learning Applications uses a statistical background to understand continuous distributions and how learning can be viewed from a probabilistic framework. The chapters progress into such machine learning topics as belief network and Bayesian reinforcement learning, which is followed by Gaussian process introduction, classification, regression, covariance, and performance analysis of Gaussian processes with other models. FEATURES Contains recent advancements in machine learning Highlights applications of machine learning algorithms Offers both quantitative and qualitative research Includes numerous case studies This book is aimed at graduates, researchers, and professionals in the field of data science and machine learning.

At our publication recap collection, we firmly count on the power of checking out Gaussian Processes For Machine Learning. Not only can this open new expertise and insights, but it can additionally save readers time and aid them determine which books to spend their time in. Allow's study the idea of Gaussian Processes For Machine Learning recaps and their advantages.

WHAT ARE PUBLICATION SUMMARIES?

Schedule recaps are compressed versions of a book's key points and themes. They provide a quick overview of Gaussian Processes For Machine Learning's significance in bite-sized pieces. They can range from a couple of paragraphs to a couple of pages.

WHY ARE THEY USEFUL?

Gaussian Processes For Machine Learning summaries are important due to the fact that they enable visitors to gain a deeper understanding of a publication's bottom lines and styles without having to review the complete publication. They are specifically valuable for hectic individuals that wish to remain enlightened but might not have the time to check out an entire publication of Gaussian Processes For Machine Learning.

HOW CAN THEY PROFIT GAUSSIAN PROCESSES FOR MACHINE LEARNING VISITORS?

Book recaps can profit viewers by conserving time, giving a hassle-free summary of Gaussian Processes For Machine Learning's essence, and assisting readers identify which books deserve investing even more time in. They permit viewers to rapidly and quickly obtain understandings and knowledge without having to devote to reading the full book of Gaussian Processes For Machine Learning.

- Saves time
- Offers a fast summary
- Aids Gaussian Processes For Machine Learning readers choose which publications to invest even more time in

Remain tuned for our following area where we will certainly dive deeper into the benefits of Gaussian Processes For Machine Learning.

Bayesian Reasoning and Machine Learning Springer Nature

Implement TensorFlow's offerings such as TensorBoard, TensorFlow.js, TensorFlow Probability, and TensorFlow Lite to build smart automation projects Key FeaturesUse machine learning and deep learning principles to build real-world projectsGet to grips with TensorFlow's impressive range of module offeringsImplement projects on GANs, reinforcement learning, and capsule networkBook Description TensorFlow has transformed the way machine learning is perceived. TensorFlow Machine Learning Projects teaches you how to exploit the benefits—simplicity, efficiency, and flexibility—of using TensorFlow in various real-world projects. With the help of this book, you'll not only learn how to build advanced projects using different datasets but also be able to tackle common challenges using a range of libraries from the TensorFlow ecosystem. To start with, you'll get to grips with using TensorFlow for machine learning projects; you'll explore a wide range of projects using TensorForest and TensorBoard for detecting exoplanets, TensorFlow.js for sentiment analysis, and TensorFlow Lite for digit classification. As you make your way through the book, you'll build projects in various real-world domains, incorporating natural language processing (NLP), the Gaussian process, autoencoders, recommender systems, and Bayesian neural networks, along with trending areas such as Generative Adversarial Networks (GANs), capsule networks, and reinforcement learning. You'll learn how to use the TensorFlow on Spark API and GPU-accelerated computing with TensorFlow to detect objects, followed by how to train and develop a recurrent neural network (RNN) model to generate book scripts. By the end of this book, you'll have gained the required expertise to build full-fledged machine learning projects at work. What you will learnUnderstand the TensorFlow ecosystem using various datasets and techniquesCreate recommendation systems for quality product recommendationsBuild projects using CNNs, NLP, and Bayesian neural networksPlay Pac-Man using deep reinforcement learningDeploy scalable TensorFlow-based machine learning systemsGenerate your own book script using RNNsWho this book is for TensorFlow Machine Learning Projects is for you if you are a data analyst, data scientist, machine learning professional, or deep learning enthusiast with basic knowledge of TensorFlow. This book is also for you if you want to build end-to-end projects in the machine learning domain using supervised, unsupervised, and reinforcement learning techniques

Advances in Knowledge Discovery and Data Mining Springer

This book constitutes the refereed proceedings of the First International Workshop on Machine Learning held in Sheffield, UK, in September 2004. The 19 revised full papers presented were carefully reviewed and selected for inclusion in the book. They address all current issues in the rapidly maturing field of machine learning that aims to provide practical methods for data discovery,

categorisation and modelling. The particular focus of the workshop was advanced research methods in machine learning and statistical signal processing.

MIT Press

A comprehensive introduction to machine learning that uses probabilistic models and inference as a unifying approach. Today's Web-enabled deluge of electronic data calls for automated methods of data analysis. Machine learning provides these, developing methods that can automatically detect patterns in data and then use the uncovered patterns to predict future data. This textbook offers a comprehensive and self-contained introduction to the field of machine learning, based on a unified, probabilistic approach. The coverage combines breadth and depth, offering necessary background material on such topics as probability, optimization, and linear algebra as well as discussion of recent developments in the field, including conditional random fields, L1 regularization, and deep learning. The book is written in an informal, accessible style, complete with pseudo-code for the most important algorithms. All topics are copiously illustrated with color images and worked examples drawn from such application domains as biology, text processing, computer vision, and robotics. Rather than providing a cookbook of different heuristic methods, the book stresses a principled model-based approach, often using the language of graphical models to specify models in a concise and intuitive way. Almost all the models described have been implemented in a MATLAB software package—PMTK (probabilistic modeling toolkit)—that is freely available online. The book is suitable for upper-level undergraduates with an introductory-level college math background and beginning graduate students.

Kernel Mean Embedding of Distributions American Mathematical Soc.

The three volume proceedings LNAI 11906 - 11908 constitutes the refereed proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases, ECML PKDD 2019, held in Würzburg, Germany, in September 2019. The total of 130 regular papers presented in these volumes was carefully reviewed and selected from 733 submissions; there are 10 papers in the demo track. The contributions were organized in topical sections named as follows: Part I: pattern mining; clustering, anomaly and outlier detection, and autoencoders; dimensionality reduction and feature selection; social networks and graphs; decision trees, interpretability, and causality; strings and streams; privacy and security; optimization. Part II: supervised learning; multi-label learning; large-scale learning; deep learning; probabilistic models; natural language processing. Part III: reinforcement learning and bandits; ranking; applied data science: computer vision and explanation; applied data science: healthcare; applied data science: e-commerce, finance, and advertising; applied data science: rich data; applied data science: applications; demo track. Chapter "Incorporating Dependencies in Spectral Kernels for Gaussian Processes" is available open access under a Creative Commons Attribution 4.0 International License via link.springer.com.

[Machine Learning and Knowledge Discovery in Databases](#) Springer

An overview of the theory and application of kernel classification methods. Linear classifiers in kernel spaces have emerged as a major topic within the field of machine learning. The kernel technique takes the linear classifier—a limited, but well-established and comprehensively studied

model—and extends its applicability to a wide range of nonlinear pattern-recognition tasks such as natural language processing, machine vision, and biological sequence analysis. This book provides the first comprehensive overview of both the theory and algorithms of kernel classifiers, including the most recent developments. It begins by describing the major algorithmic advances: kernel perceptron learning, kernel Fisher discriminants, support vector machines, relevance vector machines, Gaussian processes, and Bayes point machines. Then follows a detailed introduction to learning theory, including VC and PAC-Bayesian theory, data-dependent structural risk minimization, and compression bounds. Throughout, the book emphasizes the interaction between theory and algorithms: how learning algorithms work and why. The book includes many examples, complete pseudo code of the algorithms presented, and an extensive source code library.

Learning Kernel Classifiers Gaussian Processes for Machine Learning

Bayesian nonparametrics comes of age with this landmark text synthesizing theory, methodology and computation.

BENEFITS OF GAUSSIAN PROCESSES FOR MACHINE LEARNING PUBLICATION SUMMARIES

At our publication recap collection, our team believe in the countless advantages of reading Gaussian Processes For Machine Learning recaps. Here are a couple of crucial advantages:

- **Time-saving:** With our active routines, it can be testing to discover time to read every publication we desire. Our book recaps use a quick overview of one of the most crucial points without needing to spend a number of hours in checking out Gaussian Processes For Machine Learning entire book.
- **Quick overview of Gaussian Processes For Machine Learning:** If there is a book you're interested in, yet you're uncertain if it's ideal for you, our publication summaries supply a glimpse right into the writer's essences and composing style before purchasing the complete book.
- **Enhanced understanding in Gaussian Processes For Machine Learning:** For those who have reviewed the whole book, our publication recaps offer an opportunity to revitalize your memory and rediscover the key points and styles.

In general, book recaps of Gaussian Processes For Machine Learning offer an useful device to boost your reading experience and maximize your time and effort.

EXACTLY HOW TO COMPOSE A PUBLICATION SUMMARY OF GAUSSIAN PROCESSES FOR MACHINE LEARNING

Writing a book summary might seem like a daunting job, yet it can actually be an enjoyable and satisfying experience. Right here are some crucial elements to remember when writing your book recap:

1. **Concentrate on the essence:** The goal of a book summary is to catch the significance of

Gaussian Processes For Machine Learning in a succinct and engaging means. Avoid getting caught up in the information and rather focus on the bottom lines and motifs that the writer is trying to communicate.

2. **Maintain it quick:** Gaussian Processes For Machine Learning summary is indicated to be a quick summary, so maintain it concise. Stay with one of the most vital details and prevent entering into excessive deepness.
3. **Include the main personalities:** See to it to include a brief description of the primary characters, including their names and any type of defining qualities or qualities.
4. **Highlight the main motifs:** Recognize the main themes of Gaussian Processes For Machine Learning and highlight them in your summary. This will offer visitors a far better concept of what guide has to do with and what they can anticipate to gain from it.

By maintaining these crucial elements in mind, you can write an effective and appealing book recap that records the essence of Gaussian Processes For Machine Learning book and leaves visitors wanting more.

LOCATING THE RIGHT GAUSSIAN PROCESSES FOR MACHINE LEARNING BOOK RECAPS

Are you having a hard time to find the best Gaussian Processes For Machine Learning summaries for your passions? Don't fret, we've obtained you covered. Here are some ideas on discovering high-quality book summaries:

1. ONLINE PLATFORMS

Among the simplest methods to find Gaussian Processes For Machine Learning summaries is via online systems. Websites like Blinkist, getAbstract, and Sumizeit supply a range of recaps for various categories and styles. You can likewise take a look at Amazon Kindle's "Short Reads" area for quick, easy-to-digest recaps.

2. RESERVE REVIEW WEBSITES

Schedule review websites like Goodreads and BookPage frequently include recaps alongside their testimonials. They can give a deeper understanding of Gaussian Processes For Machine Learning plot and themes while also supplying understanding right into the viewers's experience. You can additionally take a look at their "advised" web page to uncover new summaries.

3. CURATED COLLECTIONS

Statistical Machine Learning MIT Press

The recent rapid growth in the variety and complexity of new machine learning architectures requires the development of improved methods for designing, analyzing, evaluating, and communicating machine learning technologies. Statistical Machine Learning: A Unified Framework provides students, engineers, and scientists with tools from mathematical statistics and nonlinear

optimization theory to become experts in the field of machine learning. In particular, the material in this text directly supports the mathematical analysis and design of old, new, and not-yet-invented nonlinear high-dimensional machine learning algorithms. Features: Unified empirical risk minimization framework supports rigorous mathematical analyses of widely used supervised, unsupervised, and reinforcement machine learning algorithms Matrix calculus methods for supporting machine learning analysis and design applications Explicit conditions for ensuring convergence of adaptive, batch, minibatch, MCEM, and MCMC learning algorithms that minimize both unimodal and multimodal objective functions Explicit conditions for characterizing asymptotic properties of M-estimators and model selection criteria such as AIC and BIC in the presence of possible model misspecification This advanced text is suitable for graduate students or highly motivated undergraduate students in statistics, computer science, electrical engineering, and applied mathematics. The text is self-contained and only assumes knowledge of lower-division linear algebra and upper-division probability theory. Students, professional engineers, and multidisciplinary scientists possessing these minimal prerequisites will find this text challenging yet accessible. About the Author: Richard M. Golden (Ph.D., M.S.E.E., B.S.E.E.) is Professor of Cognitive Science and Participating Faculty Member in Electrical Engineering at the University of Texas at Dallas. Dr. Golden has published articles and given talks at scientific conferences on a wide range of topics in the fields of both statistics and machine learning over the past three decades. His long-term research interests include identifying conditions for the convergence of deterministic and stochastic machine learning algorithms and investigating estimation and inference in the presence of possibly misspecified probability models.

Deterministic and Statistical Methods in Machine Learning CRC Press

Provides a comprehensive review of kernel mean embeddings of distributions and, in the course of doing so, discusses some challenging issues that could potentially lead to new research directions. The targeted audience includes graduate students and researchers in machine learning and statistics.

Efficient Reinforcement Learning Using Gaussian Processes OECD Publishing

Gaussian Process Regression Analysis for Functional Data presents nonparametric statistical methods for functional regression analysis, specifically the methods based on a Gaussian process prior in a functional space. The authors focus on problems involving functional response variables and mixed covariates of functional and scalar variables. Covering the basics of Gaussian process regression, the first several chapters discuss functional data analysis, theoretical aspects based on the asymptotic properties of Gaussian process regression models, and new methodological developments for high dimensional data and variable selection. The remainder of the text explores advanced topics of functional regression analysis, including novel nonparametric statistical methods for curve prediction, curve clustering, functional ANOVA, and functional regression analysis of batch data, repeated curves, and non-Gaussian data. Many flexible models based on Gaussian processes provide efficient ways of model learning, interpreting model structure, and carrying out inference, particularly when dealing with large dimensional functional data. This book shows how to use these Gaussian process regression models in the analysis of functional data. Some MATLAB® and C codes

are available on the first author's website.

Machine Learning Cambridge University Press

This book introduces machine learning methods in finance. It presents a unified treatment of machine learning and various statistical and computational disciplines in quantitative finance, such as financial econometrics and discrete time stochastic control, with an emphasis on how theory and hypothesis tests inform the choice of algorithm for financial data modeling and decision making. With the trend towards increasing computational resources and larger datasets, machine learning has grown into an important skillset for the finance industry. This book is written for advanced graduate students and academics in financial econometrics, mathematical finance and applied statistics, in addition to quants and data scientists in the field of quantitative finance. *Machine Learning in Finance: From Theory to Practice* is divided into three parts, each part covering theory and applications. The first presents supervised learning for cross-sectional data from both a Bayesian and frequentist perspective. The more advanced material places a firm emphasis on neural networks, including deep learning, as well as Gaussian processes, with examples in investment management and derivative modeling. The second part presents supervised learning for time series data, arguably the most common data type used in finance with examples in trading, stochastic volatility and fixed income modeling. Finally, the third part presents reinforcement learning and its applications in trading, investment and wealth management. Python code examples are provided to support the readers' understanding of the methodologies and applications. The book also includes more than 80 mathematical and programming exercises, with worked solutions available to instructors. As a bridge to research in this emergent field, the final chapter presents the frontiers of machine learning in finance from a researcher's perspective, highlighting how many well-known concepts in statistical physics are likely to emerge as important methodologies for machine learning in finance.

Gaussian Processes for Machine Learning CRC Press

Computer simulation experiments are essential to modern scientific discovery, whether that be in physics, chemistry, biology, epidemiology, ecology, engineering, etc. Surrogates are meta-models of computer simulations, used to solve mathematical models that are too intricate to be worked by hand. Gaussian process (GP) regression is a supremely flexible tool for the analysis of computer simulation experiments. This book presents an applied introduction to GP regression for modelling and optimization of computer simulation experiments. Features:

- Emphasis on methods, applications, and reproducibility.
- R code is integrated throughout for application of the methods.
- Includes more than 200 full colour figures.
- Includes many exercises to supplement understanding, with separate solutions available from the author.
- Supported by a website with full code available to reproduce all methods and examples.

The book is primarily designed as a textbook for postgraduate students studying GP regression from mathematics, statistics, computer science, and engineering. Given the breadth of examples, it could also be used by researchers from these fields, as well as from economics, life science, social science, etc.

Bayesian Reasoning and Gaussian Processes for Machine Learning Applications Springer

Nature

Machine Learning: A Bayesian and Optimization Perspective, 2nd edition, gives a unified perspective on machine learning by covering both pillars of supervised learning, namely regression and classification. The book starts with the basics, including mean square, least squares and maximum likelihood methods, ridge regression, Bayesian decision theory classification, logistic regression, and decision trees. It then progresses to more recent techniques, covering sparse modelling methods, learning in reproducing kernel Hilbert spaces and support vector machines, Bayesian inference with a focus on the EM algorithm and its approximate inference variational versions, Monte Carlo methods, probabilistic graphical models focusing on Bayesian networks, hidden Markov models and particle filtering. Dimensionality reduction and latent variables modelling are also considered in depth. This palette of techniques concludes with an extended chapter on neural networks and deep learning architectures. The book also covers the fundamentals of statistical parameter estimation, Wiener and Kalman filtering, convexity and convex optimization, including a chapter on stochastic approximation and the gradient descent family of algorithms, presenting related online learning techniques as well as concepts and algorithmic versions for distributed optimization. Focusing on the physical reasoning behind the mathematics, without sacrificing rigor, all the various methods and techniques are explained in depth, supported by examples and problems, giving an invaluable resource to the student and researcher for understanding and applying machine learning concepts. Most of the chapters include typical case studies and computer exercises, both in MATLAB and Python. The chapters are written to be as self-contained as possible, making the text suitable for different courses: pattern recognition, statistical/adaptive signal processing, statistical/Bayesian learning, as well as courses on sparse modeling, deep learning, and probabilistic graphical models. New to this edition: Complete re-write of the chapter on Neural Networks and Deep Learning to reflect the latest advances since the 1st edition. The chapter, starting from the basic perceptron and feed-forward neural networks concepts, now presents an in depth treatment of deep networks, including recent optimization algorithms, batch normalization, regularization techniques such as the dropout method, convolutional neural networks, recurrent neural networks, attention mechanisms, adversarial examples and training, capsule networks and generative architectures, such as restricted Boltzman machines (RBMs), variational autoencoders and generative adversarial networks (GANs). Expanded treatment of Bayesian learning to include nonparametric Bayesian methods, with a focus on the Chinese restaurant and the Indian buffet processes. Presents the physical reasoning, mathematical modeling and algorithmic implementation of each method. Updates on the latest trends, including sparsity, convex analysis and optimization, online distributed algorithms, learning in RKH spaces, Bayesian inference, graphical and hidden Markov models, particle filtering, deep learning, dictionary learning and latent variables modeling. Provides case studies on a variety of topics, including protein folding prediction, optical character recognition, text authorship identification, fMRI data analysis, change point detection, hyperspectral image unmixing, target localization, and more.

For readers that like a more personalized touch, curated collections are a fantastic option. These collections are usually produced by sector experts or fanatics and supply a list of must-read recaps

for various genres. You can find them on blog sites, podcasts, and even social media groups.

With these pointers, you can locate the ideal Gaussian Processes For Machine Learning book recaps for your interests and choices. Delighted reading!

REVIEW OF GAUSSIAN PROCESSES FOR MACHINE LEARNING

- A Tree Grows in Brooklyn is an amazing novel, a kind of New York To Kill a Mockingbird. Beginning in 1912 but then flashing back to the early 1900s, it is the story of Francie Nolan, the daughter in a poor family. Her mother, Kate, is determined to get the best for her son, Neely, who is a year younger than Francie. Their father, Johnny, is handsome and funny and flamboyant, and also drinks way too much. And Francie loves to go the library. She's working her way through the books, from A to Z. A Tree Grows in Brooklyn is the story of the Rommely and the Nolan families' lives, centering on Francie. The title comes from a tree that grows in the yard. It's hard to pin-point what's so great about this novel. I read it 2 years ago, and loved it. A Tree Grows in Brooklyn is less compelling than To Kill a Mockingbird, less deep, I guess you would say, but it's still a classic in its own right, and written much earlier. And it is pretty deep, whatever that actually means. Betty Smith brings to life the squalor and poverty of Brooklyn, New York, with her elegant prose. She creates wonderful, memorable characters and scenes. I loved reading about the various things that Francie and Neely did to earn a bit of candy money (though obviously, it was kind of sad), and reading about the various eccentricities of everyone in the neighborhood. You can certainly tell that Betty Smith was knowledgeable about Brooklyn; I'm pretty sure this book is based upon her own childhood at the

turn of the century, and it shows. It feels realistic. I also loved the way that school was portrayed. A lot of the kids are awful to one another; even after they go through something, they tease other kids for having the same thing happen to them. The teachers ignore when the children need to use the bathroom, and bullying is rampant. And yet, there are moments of brightness: when the music teacher comes, when Sissy (Francie's aunt) intercedes for her. Francie actually really does like school. A thing I found interesting was that even though Manhattan isn't that far away, it seems so foreign to Francie, a whole other world that she can't even imagine. The characters were amazing. Katie, Johnny, Francie, Neely, Sissy...all of them and more felt realistic and were really fun to read about. Francie, particularly. She's an avid reader much like myself, "on that day when she first knew she could read, she made a vow to read one book a day as long as she lived." I can identify with that, and I would highly recommend A Tree Grows in Brooklyn. All of my reviews can be read at my blog, novareviews.blogspot.com.

- It's certainly easy to see why this book is a classic. The haunting story of Francie Nolan, a young girl living in poverty in Brooklyn, strikes a chord deep within all of us. Beautifully told, this story draws us into Francie's early-20th century life—her father's drunkenness, her mother's disillusionment with life, her deep love of books— and makes us respond with empathy. This is certainly one of those books that transcends descriptions—it is simply a book about life that has become a classic because of its enduring truths. Betty Smith's strong main character, keen and startlingly truthful observations about life, and stark descriptions make this an unforgettable read—deserving many more than five "stars."