

COMPUTER VISION

RESOURCE GUIDE



Taha Anwar



Bleed AI

Contents

- Introduction
- About The Author
- Python Libraries and Packages
- Websites & Blogs
- Youtube Channels & Social Sites
- Courses
- Books
- Computer Vision Resource Compilations.
- Tutorials
- Computer Vision Conferences
- Final Notes

Introduction

Who is this Road Map for?

So this guide is for you if you're a:

- A programmer
- Engineer
- Scientist
- Entrepreneur
- Hobbyist
- Hacker
- Or any other synonym of the word "programmer"

Haha, ignore the last one, but seriously this is a general-purpose guide aimed at anyone looking to get started in the computer vision field. In this guide, I will list out the best available learning resources (Books, courses, blogs, libraries, etc) in computer vision. Even if you already are a seasoned practitioner this will still be a really useful guide for you.

Note: I will not throw a compiled list of all the resources at you, but rather will only describe the most important resources and a few optional but good ones. Also since there are many great resources out there on Artificial intelligence and machine learning so the resources I mention would primarily be focused on Computer Vision.

About The Author



Before we get started, here's a little bit about me.

I'm Taha Anwar (Founder Bleed AI), a Computer Vision Engineer, an Applied Scientist, Entrepreneur, and a Mensan. In the past few years alone I have:


- Taught Computer Vision courses in Universities.
- Delivered Seminars in events like Pycon, Google Developer Groups, etc, and various other Prominent universities.
- Open Sourced a number of computer vision applications and published a convenience library in Vision.
- Published a number of interesting High-Quality blog posts in Vision here at Bleed AI and some at LearnOpenCV.com.
- Published a State of the Art course in **Computer Vision & Image processing** in Urdu/Hindi language which is unparalleled by any university course in vision.

I have also worked with the Official **OpenCV.org** Team for the past two years. For me, teaching and publishing Computer Vision tutorials and applications is a passion and I love it when I see people use vision to solve some interesting real-world problems.

Python Libraries & Packages

Alright, let's discuss some popular python packages that you would need to learn to use for computer vision-related tasks. This is not an exhaustive list and there are many other packages, I've just listed the ones which I use most often and which I think are most useful.

OpenCV




If you just wanted a taste of image processing and computer vision and you only had the time to learn one library then I would definitely recommend you to go for **OpenCV library**. OpenCV is a highly efficient computer vision library released in the year 2000, it's written in C++ but has bindings for other languages e.g. Python, java etc. It has more than 2500 optimized algorithms and has a user community of more than 47 thousand people. It's used to develop a variety of computer vision applications e.g. Object tracking, AR, 3d Reconstruction, image stitching, stereo images, Image recognition etc.

Website/Docs: [Link](#)

Quick Installation: [Link](#)

Tensorflow / Keras



OpenCV can handle a series of image processing and classical computer vision tasks but if you want to train Deep Learning models for vision related tasks then OpenCV is not the ideal option. For that you can use the **Tensorflow library**, a very popular option in the AI community. It's used for a wide range of tasks but it's primarily used to build, train and deploy state of the art deep learning models. With the release of it's version 2 the API has gotten even more user friendly. The library was first released in 2015 and is maintained by Google.

Website/Docs: [Link](#)

Quick Installation: [Link](#)

There's also this **Keras** library which is a high level wrapper over Tensorflow, you don't need to install keras separately as this wrapper now comes packaged with tensorflow and you can use it by using **tf.keras**.

Pytorch



A strong competitor to tensorflow is pytorch, before the release of TF 2, there used to be major differences between the way tensorflow and pytorch handled computation but now essentially both libraries do the same thing, pytorch's syntax is arguably more pythonic which is why it has seen a huge boost in recent years, especially in the research community.

Both tensorflow and pytorch are great choices, pytorch is more popular in the research community and tensorflow is more commonly used for model deployment.

Website/Docs: [Link](#)

Quick Installation: [Link](#)

Dlib



Dlib is a machine learning and numerical computation library developed in C++ by **Davis King**. It's a great library and has numerous use cases ranging from object tracking, facial recognition, training ML networks to Networking etc. But it's most popularly used for facial landmark detection tasks.

Website/Docs: [Link](#)

Quick Installation: [Link](#)

Scikit-Learn



Scikit-Learn (also known as Sklearn) which started in 2007 Google Summer of Code and saw its first public release in 2010 is another efficient and easy to use machine learning library.

Although this library is not ideal for training deep learning models but still it's widely used because with scikit-Learn you can train a variety of machine learning models with just a few lines of code, and most of the time you only need to change a single line of code to try out a different model.

Website/Docs: [Link](#)

Quick Installation: [Link](#)

PIL



PIL (Python Imaging Library, also called Pillow) is a simple image processing library, It's useful for image format conversions and other simpler tasks like reading/saving images.

Website/Docs: [Link](#)

Quick Installation: [Link](#)

Matplotlib

matplotlib

Matplotlib is not exactly a library that you can use for vision but rather it's a plotting library, it's tremendously useful for plotting multiple images, loss graphs and other similar graphical tasks.

Website/Docs: [Link](#)

Quick Installation: [Link](#)

Websites / Blogs:

[Bleed AI](#)

So **Bleed AI** is my blog, it's for all engineers, scientists, students, hobbyists, and practitioners who are interested in computer vision, machine learning, and deep learning. Here I publish hands-on tutorials with just the right mix of theory and math to lift you off the ground and make you start utilizing Computer Vision Skills efficiently and as fast as possible.

[LearnOpenCV](#)

LearnOpenCV is also a great computer vision blog by **Satya Mallick** and his team, where a lot of domain experts publish blog posts on a variety of vision topics. I have published a number of great posts on LearnOpenCV, which you can [check out here](#). Satya is also serving as CEO of OpenCV.org.

[Pyimagesearch](#)

Pyimagesearch is also a great computer vision website by **Adrian Rosebrock**. It was started in 2014 and now has 100s of high quality computer vision blog posts. Almost everyone learning vision would have come across this site.

[Machine Learning Mastery](#)

Machine Learning Mastery is a great website, it was founded by **Jason Brownlee**. This site is not primarily focused on computer vision but still it covers a variety of topics in computer vision, deep learning and machine learning. The site also has some great and affordable ebooks on vision and other areas.

Other Useful Blogs:

Here are 7 other really useful blogs/sites that I think are worth mentioning here:

- [Towards Data Science](#)
- [Analytics Vidhya](#)
- [Medium](#)
- [HACKERNOON](#)
- [CV-Tricks.com](#)
- [KDNuggets](#)
- [Neptune.AI](#)

Youtube Channels & Social Sites

Bleed AI

Although we do have an old channel with some tutorials [here](#), but we've just started a new channel with a free video course called **“Computer Vision For Everyone”**. This course is excellent for beginners, it's designed to cover a variety of topics in computer vision at a high level. Later on we also plan to cover other exciting topics in computer vision, so be sure to [Subscribe Here](#).

Murtaza Robotics

This is an excellent channel for you to not only learn computer vision but also combine vision techniques with drones and edge devices to build Cool robotic applications. Make sure to check this channel out.

First Principles Of Computer Vision

A relatively new channel that I just found, but to be honest I can't find a better channel than this that explains the maths, theoretical concepts behind different foundational Vision topics. Absolutely amazing.

Augmented Startups

Another versatile channel that covers a variety of computer vision topics and also combines it with edge devices to build interesting applications.

Pysource

Another channel that I found useful, if you like to see opencv and other vision based applications coded from scratch then this channel might be for you.

Other Youtube Channels:

Here are some other really useful channels that although are not directly focused on computer vision but I feel are really useful for building strong theoretical concepts in AI or learning about the latest trends.

- [*Yannic Kilcher YouTube Channel*](#)
- [*Arxiv Insights YouTube Channel*](#)
- [*Two Minute Papers YouTube Channel*](#)

Facebook & Reddit Groups

Here are a few facebook and reddit groups on vision and AI that I feel are really active and also have good quality engagement.

- [*AIDL \(Facebook Group\)*](#)
- [*Computer Vision \(Facebook Group\)*](#)
- [*Reddit Computer Vision*](#)
- [*Reddit Machine Learning*](#)

Tutorials & Applications

If you're just starting out in computer vision then here you will find some interesting beginner friendly computer vision tutorials and some very creative applications.

Computer Vision Tutorials:

- [Computer Vision Crash Course with OpenCV & Python](#)
- [A Crash Course with Dlib Library, 101 to Mastery](#)
- [Deep Learning with OpenCV DNN Module, A Comprehensive Guide](#)
- [Training a Custom Image Classifier with Tensorflow, Converting to ONNX and using it in OpenCV DNN module | Bleed AI](#)

Creative Vision Applications:

- [Building a Smart Intruder Detection System with OpenCV and your Phone](#)
- [Training a Custom Object Detector with DLIB & Making Gesture Controlled Applications](#)
- [Creating a Virtual Pen And Eraser with OpenCV](#)
- [Playing Rock, Paper, Scissors with AI](#)

Of Course there are many other creative applications and tutorials, these were just some of my favourite ones that I've written.

Courses

Alright, here I'm going to share some great computer vision courses that are out there. I'll cover both paid and free courses. In most cases the value provided in paid courses is far greater than what you'll get for free since the money is used to provide more support and build higher quality content.

Bleed AI Course

Our **Computer Vision & Image Processing** course (video lectures in urdu/hindi) is designed to help you master OpenCV with Python in just 3 months. This course is completely applied meaning you can get right into building powerful applications without pondering unnecessary math.

It doesn't matter if you're a developer, engineer, researcher, or even a university student. The only prerequisite you need is just some programming experience in any language.

OpenCV Courses

OpenCV also offers a series of online courses in computer vision to educate people around the globe. The courses are designed for students and domain professionals with the only prerequisite of Python or C++. Each course requires three to four months to finish.

PyimageSearch Courses

PyimageSearch also provides a number of computer vision courses, the courses offered here are on OpenCV, computer vision applications, raspberry PI/edge devices, Optical Character Recognition and Deep learning.

Udacity Computer Vision Nano Degree

Udacity, which is one of the leading MOOC platforms, also has a great course in computer vision, the full track is called “Computer Vision Nano Degree”, (it's not a college degree). Although the course outline seems good, the course itself is quite expensive and also requires some prerequisites including Python programming experience, statistics, machine learning basics, and Tensorflow or Pytorch experience.

Introduction to Computer Vision (FREE)

This is another vision course on Udacity, this course has a few pros and a few cons. On the positive end, it covers a lot of low level vision topics in depth and lasts 4 months, the overall content is great and videos are of high quality. Now the downside of this course is that first it's not in python, it uses matlab, it does not cover modern deep learning or even much of machine learning methods, it goes into some algorithms that are not much used today and thus can become boring at times. Even if you're not interested in low level vision Maths, there are still a number of videos in this course that can prove to be useful and intuitive.

Deep Learning For Computer Vision (University Of Michigan -FREE)

One of the most intuitive explanations of Modern deep learning techniques can be found in this deep learning course, by Justin Johnson. Justin formerly taught the same things in **CS231n course** by Stanford University. Since this course is taught on campus, it's not online, but you can find all the [video lectures here](#)..

[Intro to Deep Learning by MIT \(FREE\)](#)

This is a recent deep learning course by MIT University, similar to the above course. It has some great lectures on modern deep learning methods, you can watch the [video lectures here](#). Instead of just focusing on Computer vision, this course has lectures on other related deep learning topics too.


Beside these, you can find many other courses on different topics of AI on these Massive Open Online Course (MOOC) sites:

- [***Coursera***](#)
- [***Udacity***](#)
- [***Edx***](#)
- [***Udemy***](#)

Books For Computer Vision

I personally prefer blog sites, courses, and papers over books. This is because when working I want something hands-on, if the author explains some interesting concept then I want to see that algorithm in action, so in general, very few books are able to deliver hands-on coding experience with good concepts and intuition. I will try to list out some good books that do that well.

- Deep Learning with Python



This book is written by **Francois Chollet**, the guy who created Keras. This is by far the best book on Keras, it offers a practical, hands-on exploration of deep learning with python. It keeps away from mathematical notations and other jargon and goes directly into the code after theoretical explanation. The book is filled with best practices, tips, and tricks and is an excellent guide for anyone trying to learn deep learning for computer vision.

One issue is that Keras is no longer updated so the standard practice is that you first install TensorFlow and then use `tf.keras`. Because of that, you may have to change the code slightly to work. But Francois is working on the second edition of this book which will contain new chapters and refined code.

First Edition: [Link](#)

Second Edition: [Link](#)


- Pyimagesearch Books



There are some good books by **Adrian Rosebrock** which can help you on your computer vision and deep learning journey. According to Adrian's site, these books have great reviews and have helped a lot of people to learn Computer Vision.

List of Books: [Pyimagesearch](#).


- Programming Computer Vision with Python



The author **Jan Erik Solem** explains computer vision in broad terms in this book that provides enough understanding of the underlying theory and algorithms of computer vision to serve as a foundation for students, researchers, and enthusiasts. It provides complete code samples with explanations on how to build upon each example, along with exercises to help you apply what you've learned.

First Edition: [Link](#)


- Hands-On Machine Learning with Scikit-Learn and TensorFlow



This book by author **Aurélien Géron** shows you how to code a number of algorithms, starting with simple linear regression and progressing to deep neural networks in a very simple and intuitive manner. The two tools used in this book are **Scikit-Learn** and **TensorFlow**.

Second Edition: [Link](#)

- Computer Vision: Algorithms and Applications




This book is written by **Richard Szeliski** (a renowned CV Scientist), it covers a variety of methods that are being successfully used to analyze and interpret images. The author motivates students to push their creative boundaries by focusing on basic techniques that work in a real-world environment.

First Edition: [Link](#)

Second Edition: [Link](#)

- Computer Vision: A Modern Approach



The authors **David Forsyth** and **Jean Ponce** focus on modern computer vision methods in this book and describe ideas that are useful, or likely to be so in the future. The authors have tried to put emphasis on understanding the basic geometry and physics of imaging, and link this with actual applications.

Second Edition: [Link](#)

Computer Vision Resource Compilations

Here are some pages that contain a list of sorted resources for Computer vision.

- **Awesome Computer Vision**: provides a list of awesome computer vision resources from intermediate to advanced level. The resources include books, courses, papers, software, datasets, tutorials and talks, resources for students, blogs, links, and songs.
- **Awesome computer vision papers**: A list of Sorted papers on Computer Vision.
- **Awesome Deep Vision**: A list of Sorted papers on Deep learning base Computer Vision papers categorized according to subdomain.

Computer Vision Conferences

If you plan to go into research for computer vision then you should start reading papers from top tier conferences for computer vision and machine learning. The best sites to read papers are [Arxiv](#) or [Arxiv Sanity](#). Also do check if the paper you're planning to read got published in any of these reputable conferences.

Computer Vision and Pattern Recognition (CVPR)

Computer Vision and Pattern Recognition, previously named **Pattern Recognition and Image Processing** is one of the most important conferences in computer vision. It has a ~28% average acceptance rate for papers and considers a comprehensive list of topics related to computer vision and pattern recognition.

International Conference on Computer Vision (ICCV)

The **International Conference on Computer Vision** is also considered as a top tier conference with an average acceptance rate of ~26%. Every other year a large number of researchers compete to get their papers accepted by this conference.

European Conference on Computer Vision (ECCV)

The **European Conference on Computer Vision** is also on the list of top computer vision conferences and is held in those years in which ICCV is not. It is an exceptionally competitive conference with an acceptance rate of ~25%.

International Conference on Machine Learning (ICML)

International Conference on Machine Learning is one of the primary international academic conferences that have a high impact on machine learning research with an average acceptance rate of **~29%**.

Neural Information Processing Systems (NIPS)

Neural Information Processing Systems (abbreviated as **NeurIPS**) is a leading machine learning and computational neuroscience conference with an average acceptance rate of **~24%**.

Final Notes:

Alright so these were some of the most useful resources to learn and apply Computer Vision in practical life. You don't need to go over all of them in order to become proficient in the field, just start by picking a few of them and routinely go over them and practice/code each concept you learn.

Computer vision is a vast field, you can easily get lost, the biggest mistake you can make is not having a clear focus or plan regarding what you want. So at first you need to briefly explore different areas of vision and find what you love, after deciding upon the subdomain of vision you want to study, make sure to develop expertise in it.

Finally If you want to connect with me and my team at Bleed AI, then there are a number of sites that you can reach us on, starting from our website to our Facebook page. For learning purposes the best resources are the **Bleed AI** website and the Youtube Channel. For contacting us, the fastest way to get replies is by emailing support@bleedai.com

- [Bleed AI Website](#)
- [Youtube Channel](#)
- [Facebook Page](#)
- [Bleed AI LinkedIn](#)
- [Twitter Page](#)
- Email: Support@bleedai.com
- [Taha Anwar LinkedIn](#)

*If You Find This Guide Resourceful, Please Do
Share With Others That Might Also Find It Useful.
Thank You*