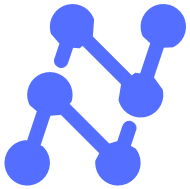

Nanonets

Building inbound GTM at Nanonets

Prathamesh Juvatkar

Frictionless workflow automation platform.

Making unstructured data interoperable

$11.5M Raised

ELEVATION

$10M Series A - Oct 2021

Y Combinator

$1.5M Seed - Apr 2017

400+ Customers worldwide

Inbound Traffic in US

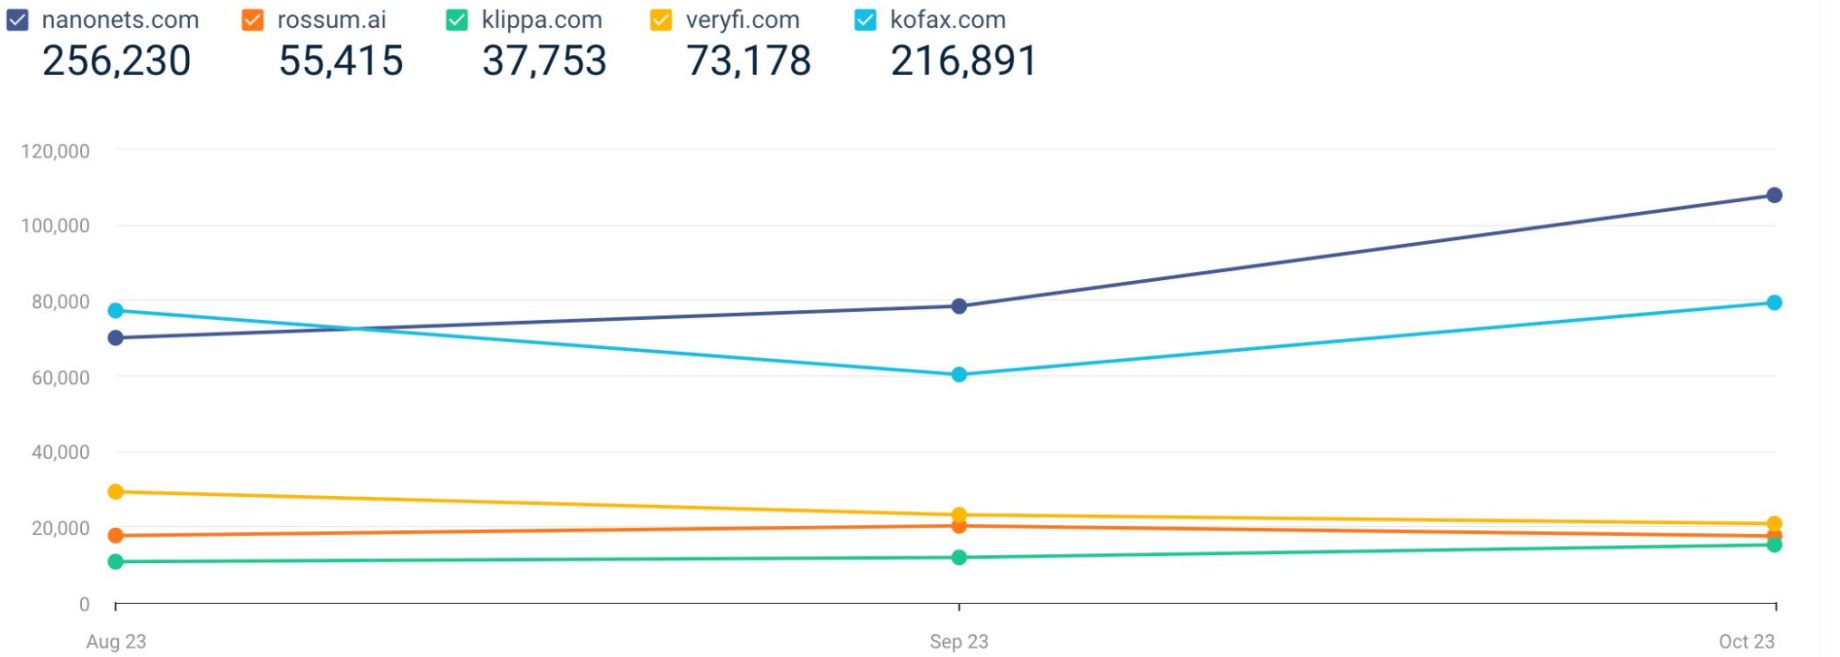

Nanonets

Traffic grew 300% in last 12 Months

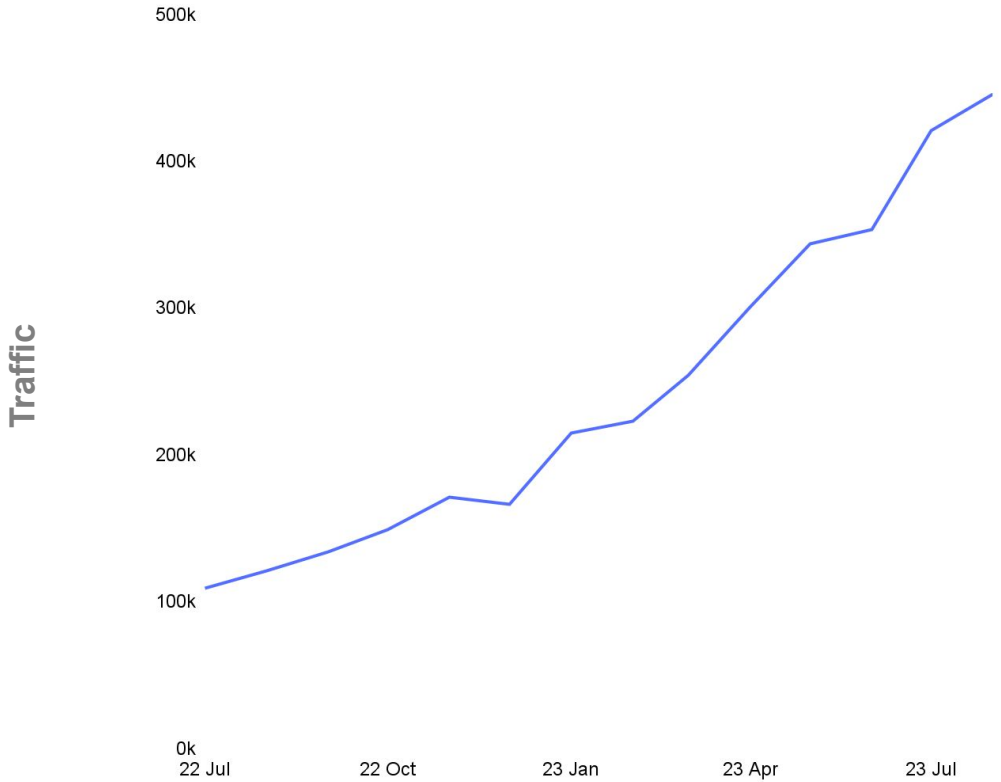

Highly Efficient - PLG Motion

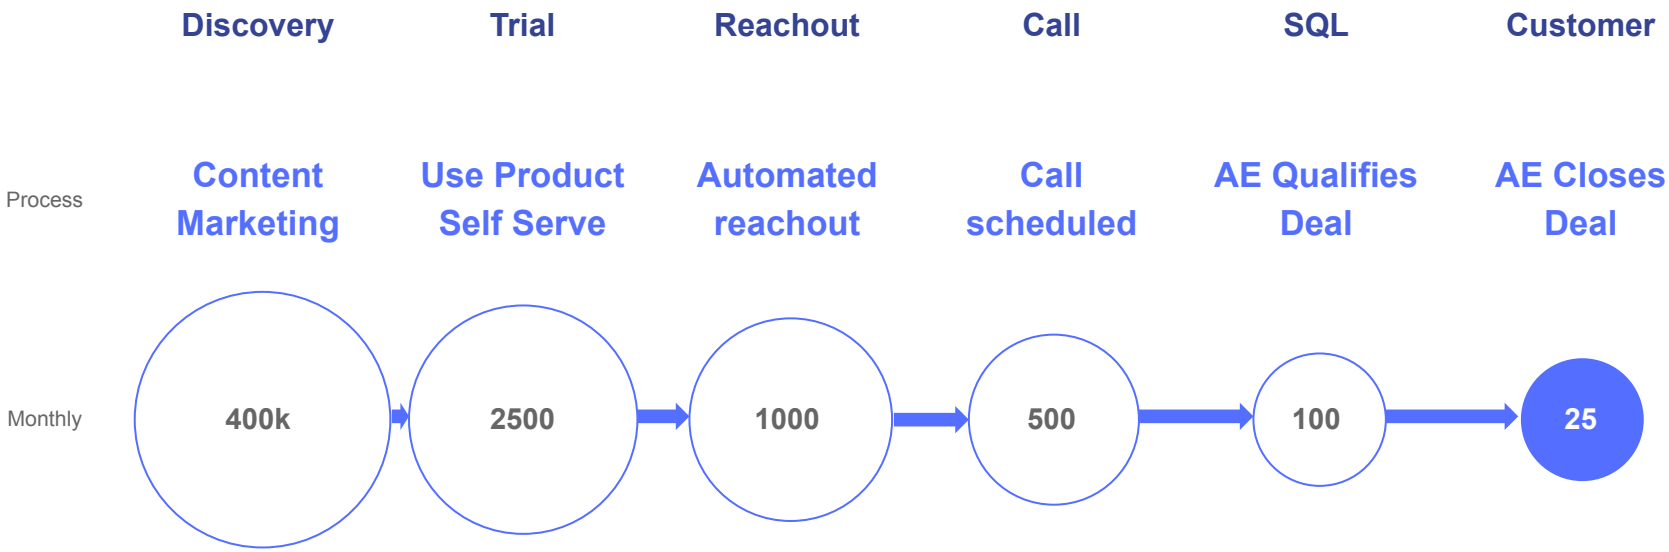

34% of Fortune 500 have used the Product

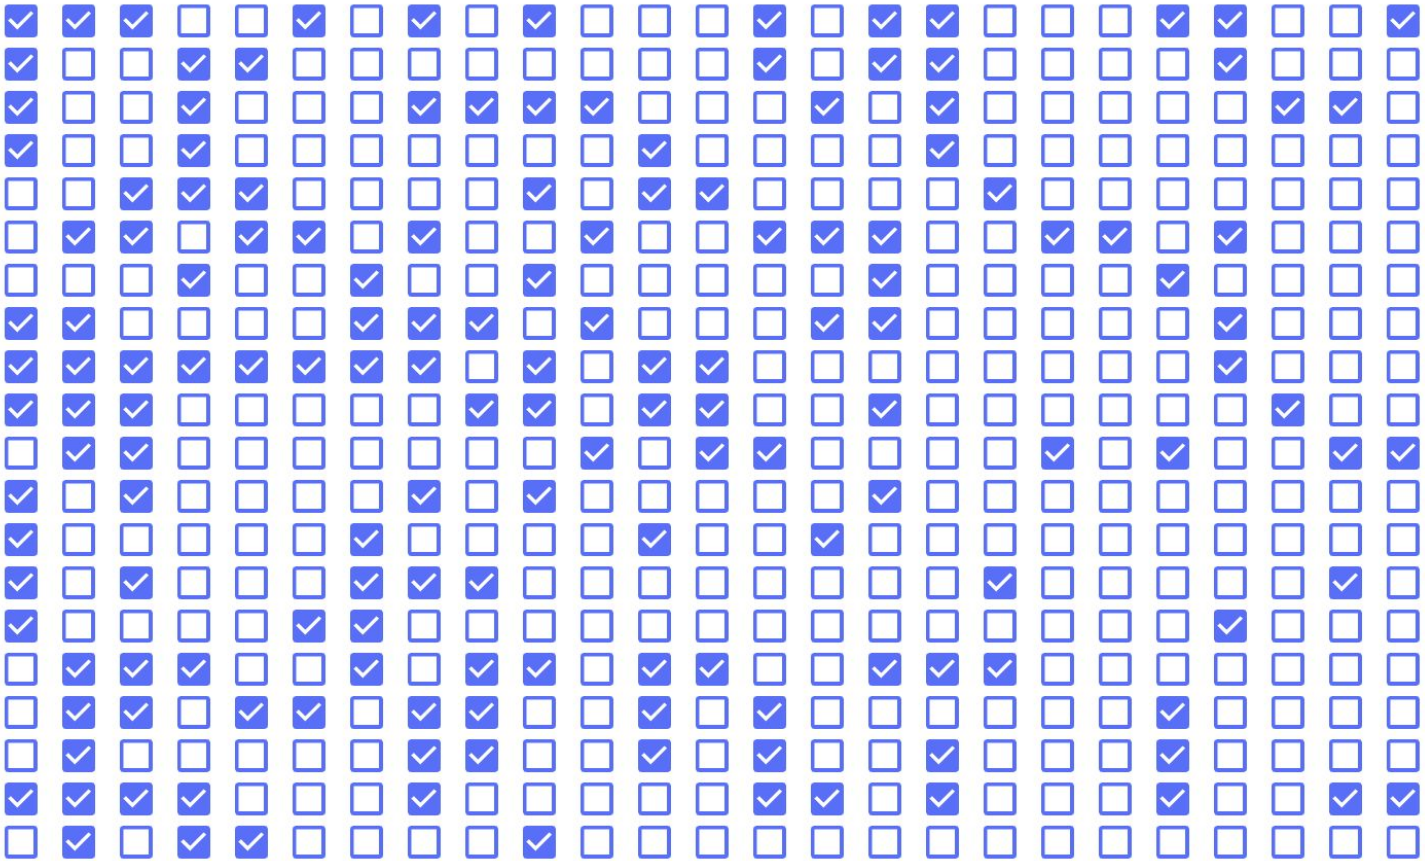

Our YC interview - 2017

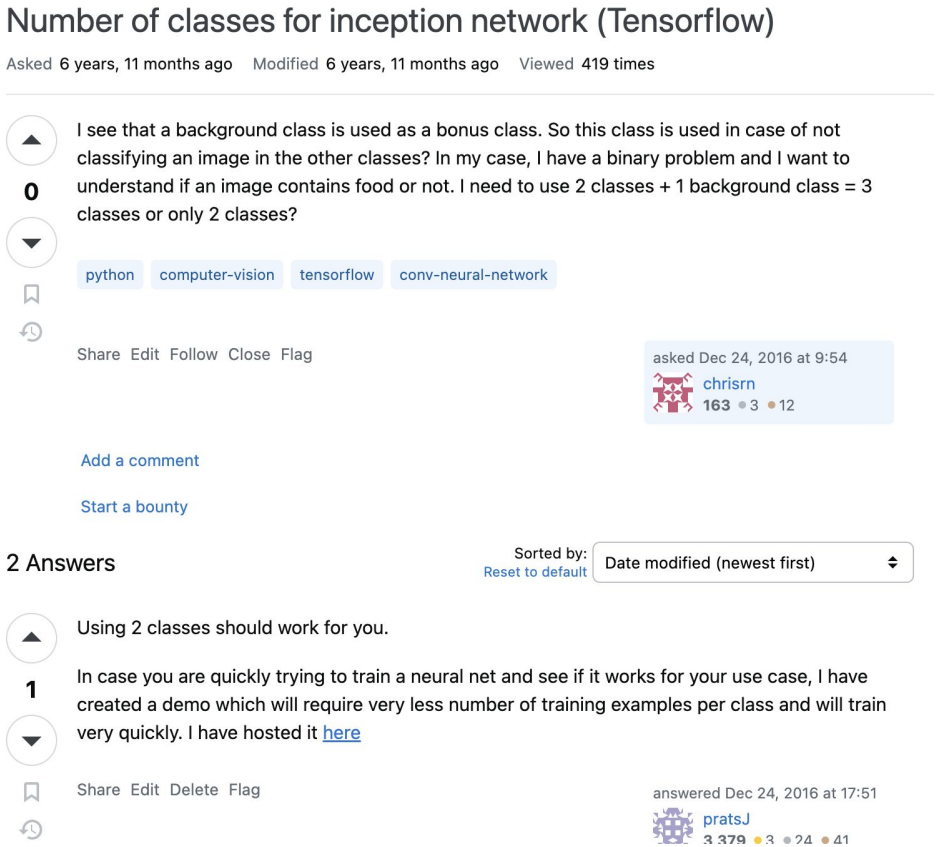

Number of classes for inception network (Tensorflow)

Asked 6 years, 11 months ago Modified 6 years, 11 months ago Viewed 419 times

0

I see that a background class is used as a bonus class. So this class is used in case of not classifying an image in the other classes? In my case, I have a binary problem and I want to understand if an image contains food or not. I need to use 2 classes + 1 background class = 3 classes or only 2 classes?

python computer-vision tensorflow conv-neural-network

Share Edit Follow Close Flag

asked Dec 24, 2016 at 9:54
chrisrn
163 3 12

Add a comment

Start a bounty

2 Answers

Sorted by: Reset to default Date modified (newest first)

1

Using 2 classes should work for you.

In case you are quickly trying to train a neural net and see if it works for your use case, I have created a demo which will require very less number of training examples per class and will train very quickly. I have hosted it here

Share Edit Delete Flag

answered Dec 24, 2016 at 17:51
pratsJ
3,379 3 24 41

- Developer trying to solve a problem with ML
- Doesn’t have in-depth ML knowledge
- Trying to build solution with OSS frameworks
- Asking questions on SO

Rank the blog post on Hacker News

Rank the blog post on Hacker News

- Place where devs with lot of clout spend time
- Lot of publications will pickup articles from Hacker News
- Will get lot of high quality backlinks

Rank the blog post on Hacker News

- Informational content works best on Hacker News
- Will get lot of feedback and discussions on the post
- Ideas for next content as well as product feedback

Rank the blog post on Hacker News

- Ranking on Hacker News needs high quality content
- Initial traffic and low bounce rate helps with ranking
- More likely to get ranked on search engines as well

Hacker News posts

▲ How to use Deep Learning when you have Limited Data (medium.com/nanonets)

252 points by sarthakjain on Feb 2, 2017 | hide | past | favorite | 50 comments

sarthakjain on Feb 2, 2017 | next [–]

Machine Learning and AI seem to be in vogue but become tough to implement unless you have boatloads of data. We've personally had multiple frustrating experiences over the last ~7 years of trying to solve problems using ML. In almost all the cases we failed to ship due to lack of data. Transfer Learning is a major breakthrough in ML where companies with little data can also build state of the art models. Unfortunately not enough people know about it. We are trying to do our part to make it easier to use Transfer Learning as well as increase awareness about it.

Using Transfer Learning we can build a model to identify cats and dogs in images with a few (<100) images as compared to the few thousands it would take before.

To make Transfer Learning easy we are building https://nanonets.ai that has multiple pretrained models that can be augmented with your data to create state of the art models. We are currently in the process of building our first few models. Image Labeling and Object Detection (in Images) work with a few Text based models coming up in the next few weeks.

▲ chriskanan on Feb 2, 2017 | parent | next [–]

Is transfer learning really not widely known by people doing AI? In my field, computer vision, it is used by most of the papers in the past three years in CVPR, etc. All of the students that take either my deep learning or my computer vision courses have to do assignments on transfer learning with deep neural networks.

▲ sarthakjain on Feb 2, 2017 | root | parent | next [–]

Totally agree, everybody in the industry knows about it. However if you look at https://www.google.com/trends/explore?date=2014-01-01%202017... nobody outside seems to know. I might be wrong but a lot of people outside the ML community seem to be hesitant to using ml because they don't have enough data, trying to remove the misconception if it exists

▲ qxcv on Feb 3, 2017 | root | parent | next [–]

Google Trends can tell you a lot of things :)

https://www.google.com/trends/explore?date=2014-01-01%202017...

More seriously though: others have pointed out that finetuning is pretty popular in some subfields, but it's just one hammer in a of a whole toolbox of techniques which are necessary to make neural nets train (even when you have a tonne of data). Standardisation, choice of initialisation, and choice of learning rate schedule all come to mind as other factors which *seem* simple, but which can have a huge impact in practice.

Of course, each tool has its limitations. The most obvious limitation of finetuning is that you need a network that's already been trained on vaguely similar data. Pretraining on ImageNet is probably not going to help you solve problems where the size of objects matters, for example, because most ImageNet performance tends to benefit from scale invariance.

I wish you luck with nanonets.ai, but I think it's irresponsible to market this as the "1 weird trick" to bring data efficiency to neural nets.

▲ averagewall on Feb 3, 2017 | root | parent | prev | next [–]

That graph might be more because it's not as exciting. Google searches are probably dominated by hobbyists and casually interested people. If they're not trying to achieve a specific goal, then they might prefer to work out the basics and make it themselves instead of just taking part of someone else's work and reusing it. If you were going to do that for fun, why not go the whole hog and reuse an entire pretrained network?

Personally, I'm a hobbyist and I don't want to know about these shortcuts until I start to need them - which is a stage I might never reach. People who've progressed far enough to need them are probably far fewer than those who are just curious what these words mean.

Another possibility is the words "transfer learning" might be more generally meaningful outside the ML field than the other search terms on the graph, so most of the searches for it are really schoolteachers or something else.

Hacker News posts

How to Easily Detect Objects with Deep Learning on Raspberry Pi (medium.com/nanonets)

308 points by sarthakjain on April 2, 2018 | hide | past | favorite | 50 comments

* 9 points by prats226 on April 2, 2018 | next [–]

I am Prathamesh, co-founder of https://nanonets.com

We were working on a project to detect objects using deep learning with raspberry pi and we have benchmarked various deep learning architectures on pi. With ~100-200 images, you can create a detector of your own with this method.

In this post, we have detected vehicles in Indian traffic using pi and also added github links to code to train the model on your own dataset and then script to get inference on pi. Hope this helps!

thedirt0115 on April 2, 2018 | parent | next [–]

Hi, I have a question about the performance benchmarks section. The best performaing model from your benchmarks, ssd_mobilenet_v1, has a prediction time in seconds of 0.72 -- Is that the total runtime of the script? I'm wondering if I could achieve ~1 FPS running in realtime (basically looping as fast as possible against the camera input), or is there more overhead? (edit -- made question more specific)

* 4 points by prats226 on April 2, 2018 | root | parent | next [–]

Yeah its total runtime of script. However you can get upto 3-4 FPS with more optimizations. We are going to try more quantization options soon with release of tensorflow 1.7 and will report our findings (Will post updates in blog). Also pi camera code needs to be optimized which will further increase FPS. One big advantage of this method is just by collecting ~200 images for any use case, you can have detector ready in couple of hours.

One advantage of using API based approach is you can get much higher FPS without compromising accuracy and is also independent of pi CPU power and heating etc.

▲ jd20 on April 2, 2018 | root | parent | next [–]

Curious, have you looked at any other embedded platforms besides the Pi? The Tegra might be an interesting comparison point, I wonder what kind of FPS the onboard GPU would buy you.

* 3 points by prats226 on April 2, 2018 | root | parent | next [–]

Hey, good suggestion to benchmark against other SOC's as well. I heard raspberry pi recently added support for external graphics card. Haven't tried yet.

▲ Aspos on April 2, 2018 | prev | next [–]

How to draw an owl, yeah http://i0.kym-cdn.com/photos/images/original/000/572/078/d6d...

▲ sarthakjain on April 2, 2018 | parent | next [–]

Thanks for feedback. What could we possibly to do make it easier to follow the intermediate steps?

▲ Aspos on April 2, 2018 | root | parent | next [–]

Say I have a bunch of photos of my cat. Want to be able to use Raspberry PI to recognize my cat from the others. How can I create the dataset to feed to your ML engine? Would love to see end-to-end how-to, seriously.

▲ c54 on April 3, 2018 | root | parent | next [–]

Quality over quantity

How to easily Detect Objects with Deep Learning on Raspberry Pi

The real world poses challenges like having limited data and having tiny hardware like Mobile Phones and Raspberry Pis which can't run complex Deep Learning models. This post demonstrates how you can do object detection using a Raspberry Pi. Like cars on a road, oranges in a fridge, signatures in a document and teslas in space.

Sarthak Jain · Following
Published in NanoNets · 10 min read · Mar 20, 2018

5.8K 19

Disclaimer: I'm building nanonets.com to help build ML with less data and no hardware

> If you're impatient scroll to the bottom of the post for the Github Repos

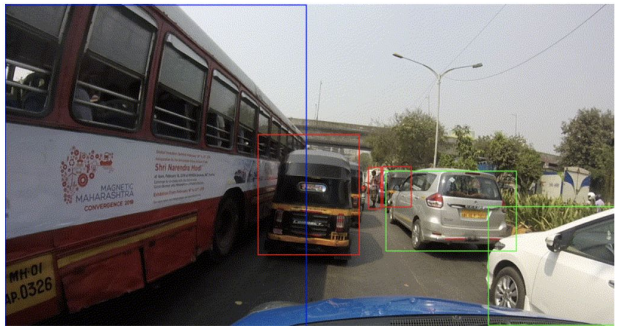

- Bought a raspberry pi, mounted on our car
- Drove through Mumbai traffic
- Quantized models to work on it real time
- Even gave away all the code we had written for it

0 → 1 playbook

- Start from day 1, takes time to build skillset. Founder led.
- Find one channel where your ICP spends time
- Write something they would love to read, focus on quality a lot
- Distribute heavily on other channels

0 → 1 playbook

- Don’t worry about keywords, technical SEO etc too much to begin with
- Keep iterating content from feedback
- Ok to measure leading indicators like views, # of shares, upvotes

1 → 10 playbook

- Instead of writing blogs yourself, start creating outlines
- Figure out topics that people like to read about
- Find out part-time writers who get views on medium in your domain
- Setup analytics for attribution
- Start measuring signups, MQL’s etc rather than views

1 → 10 playbook

page	impressions	signup
https://nanonets.com/blog/convert-handwriting-to-text/	11,964	1,047
https://nanonets.com/blog/ocr-software-best-ocr-software/	2,761	194
https://nanonets.com/blog/convert-image-to-text/	8,041	510
https://nanonets.com/blog/extract-data-from-pdf/	3,509	140
https://nanonets.com/blog/best-reconciliation-software/	261	5

10 → 100 playbook

- Invest in marketing team
- Start focusing on technical elements of inbound
  - Technical SEO
  - Internal Backlinking
  - Keyword research and prioritization
  - Long tail keywords, pSEO (if applicable)

10 → 100 playbook

- Have a content funnel and calendar to maintain pace
- TOFU/MOFU/BOFU content segregation
- Landing pages (once you have domain authority)
- Start measuring more down the funnel metrics
    - MQL’s based on ICP
    - Qualified calls setup

Content Marketing Strategy

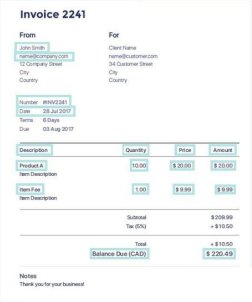

1. Task oriented

Extract Data from Invoices

Demo of invoice extraction with signup

2. State of industry

State of AP Automation in 2023

Highlighting what is going on in AP Automation with a plug

3. Goal oriented

Write Invoices to Quickbooks

How can they accomplish this goal and a CTA to try

4. Technology/Tool

Convert PDF to CSV

Tool to convert with a CTA to bulk convert using Nanonets

Mistakes to avoid

- Never lose focus on quality of the content
- Hire team who understands the product and content quality more than technical aspects of SEO
- Keep expanding scope of articles to write, never slow down pace of pushing out content

Any Questions?