Taming Service Variability, Building Worldwide Systems, and Scaling Deep Learning

Jeff Dean
Google

Joint work with many collaborators

A Talk in Two Parts...

- Part 1: Building and deploying large-scale services
  - achieving low latency in shared environments
  - dealing with worldwide deployment and operations
- Part 2: Learning high-level representations via deep learning
  - automatically, from raw data
  - useful in wide variety of domains

Faster Is Better

Google

Faster Is Better

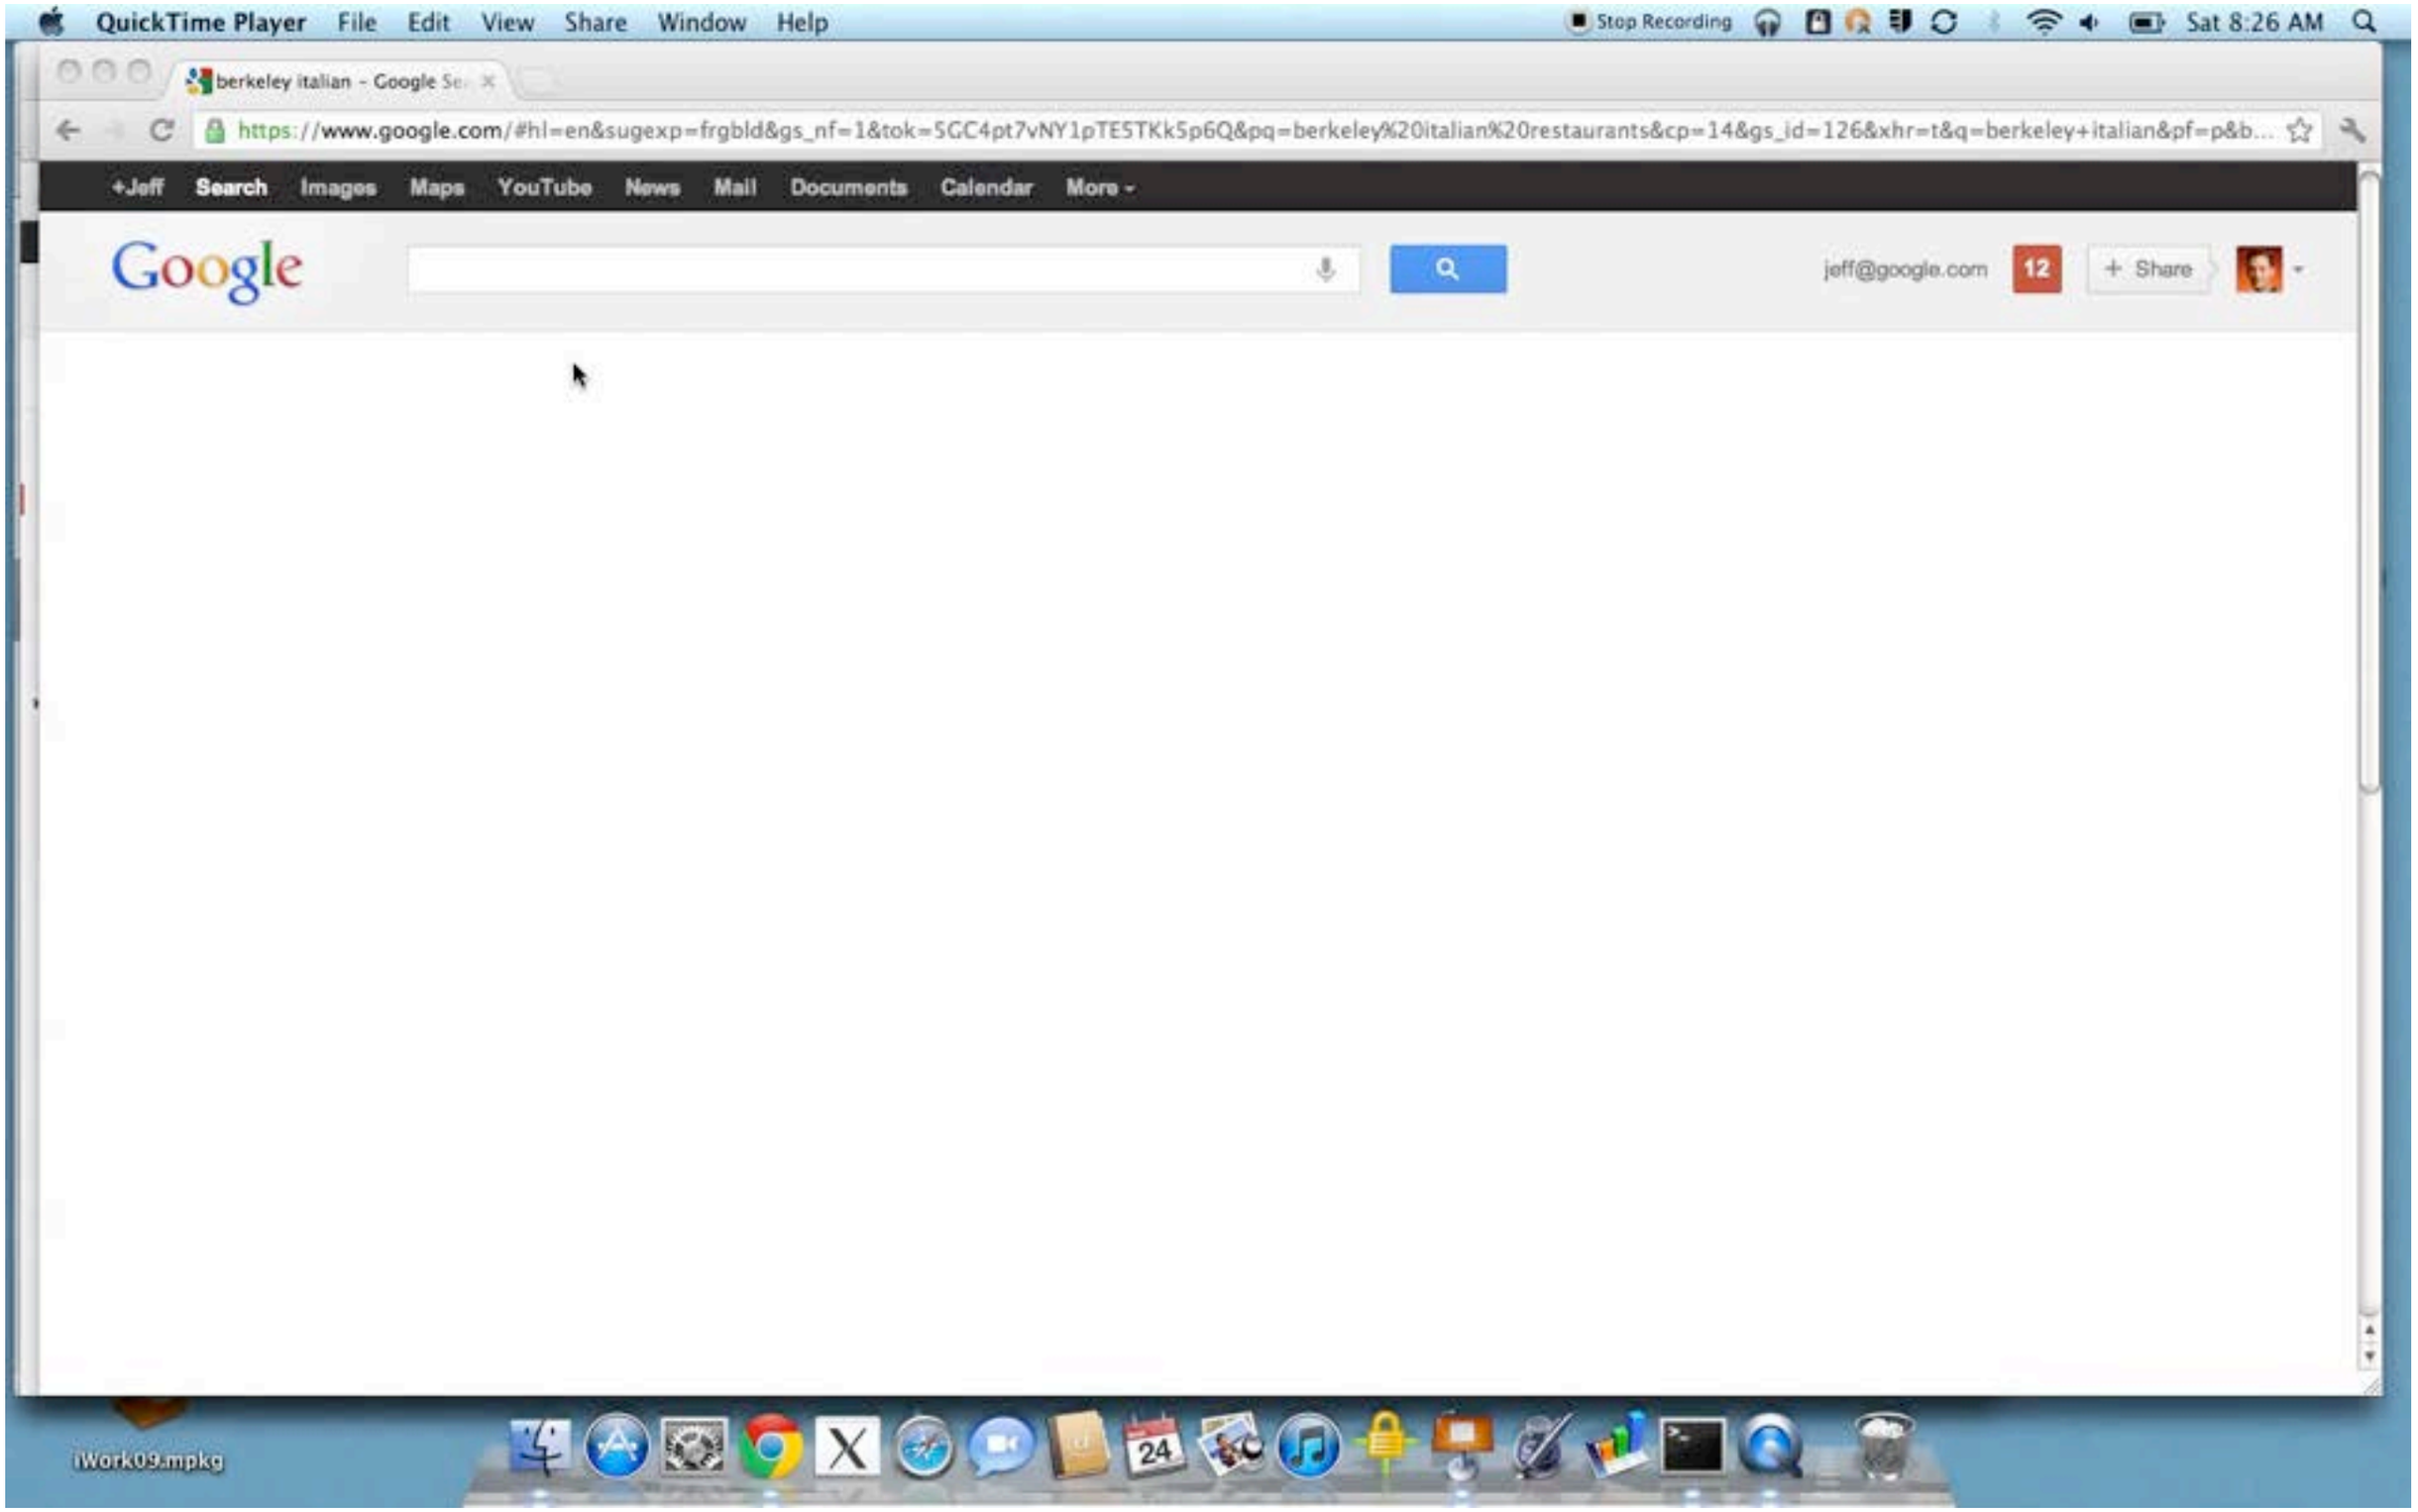

Large Fanout Services

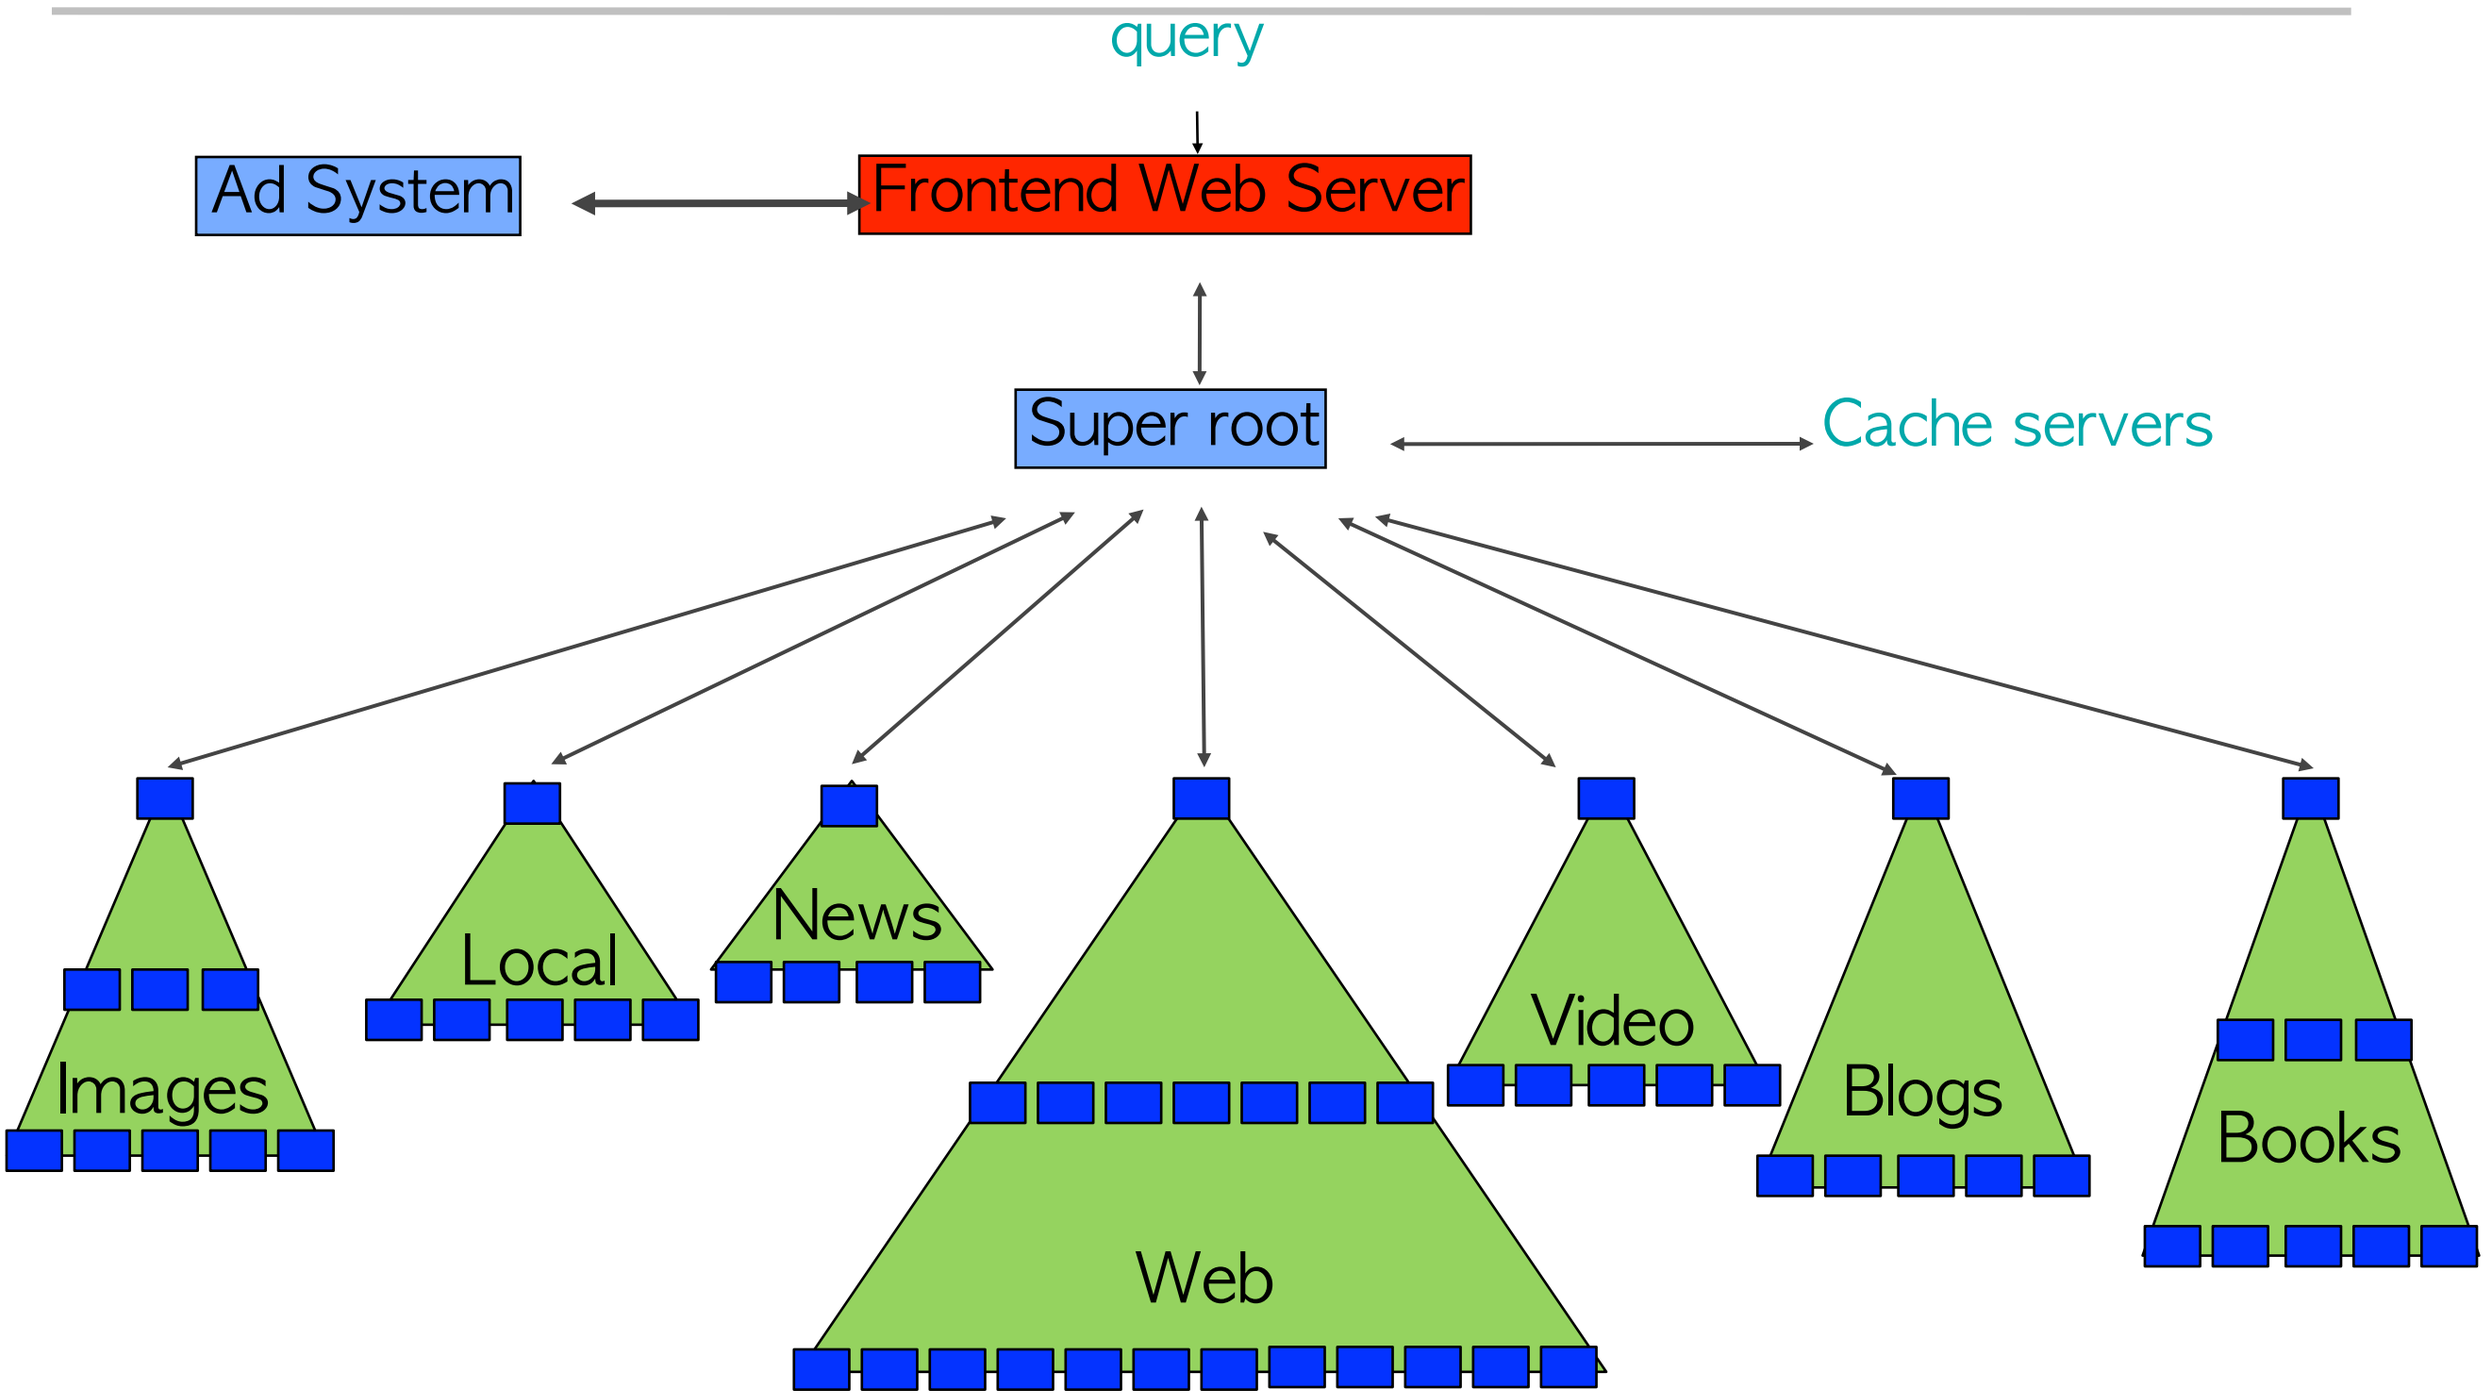

Google

Shared Environment

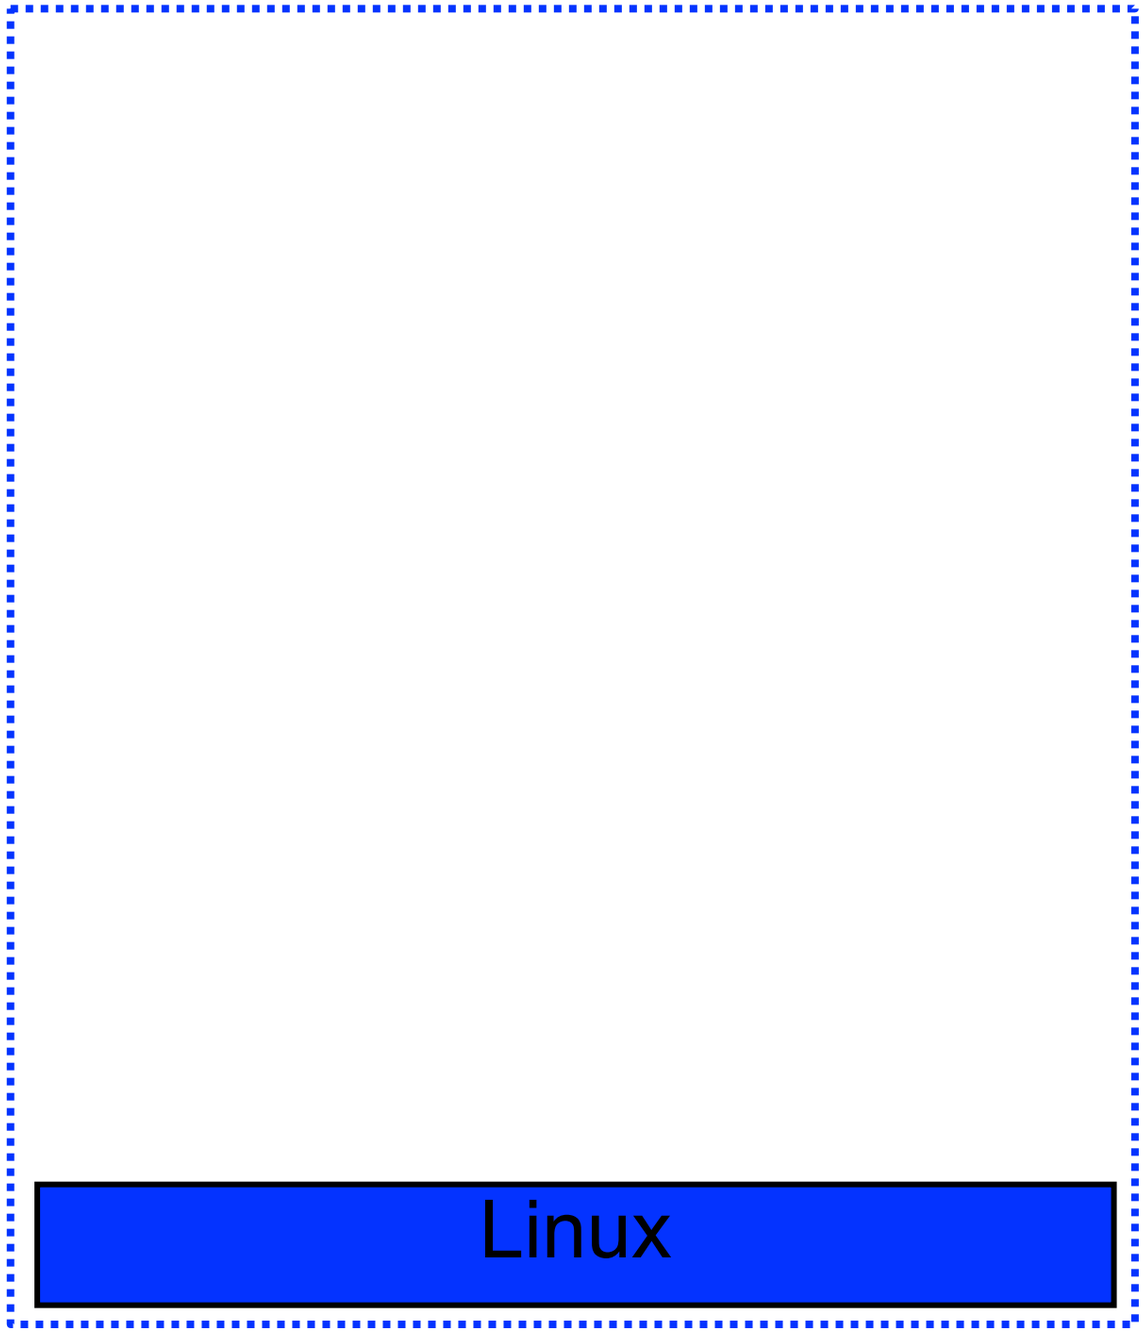

Google

Shared Environment

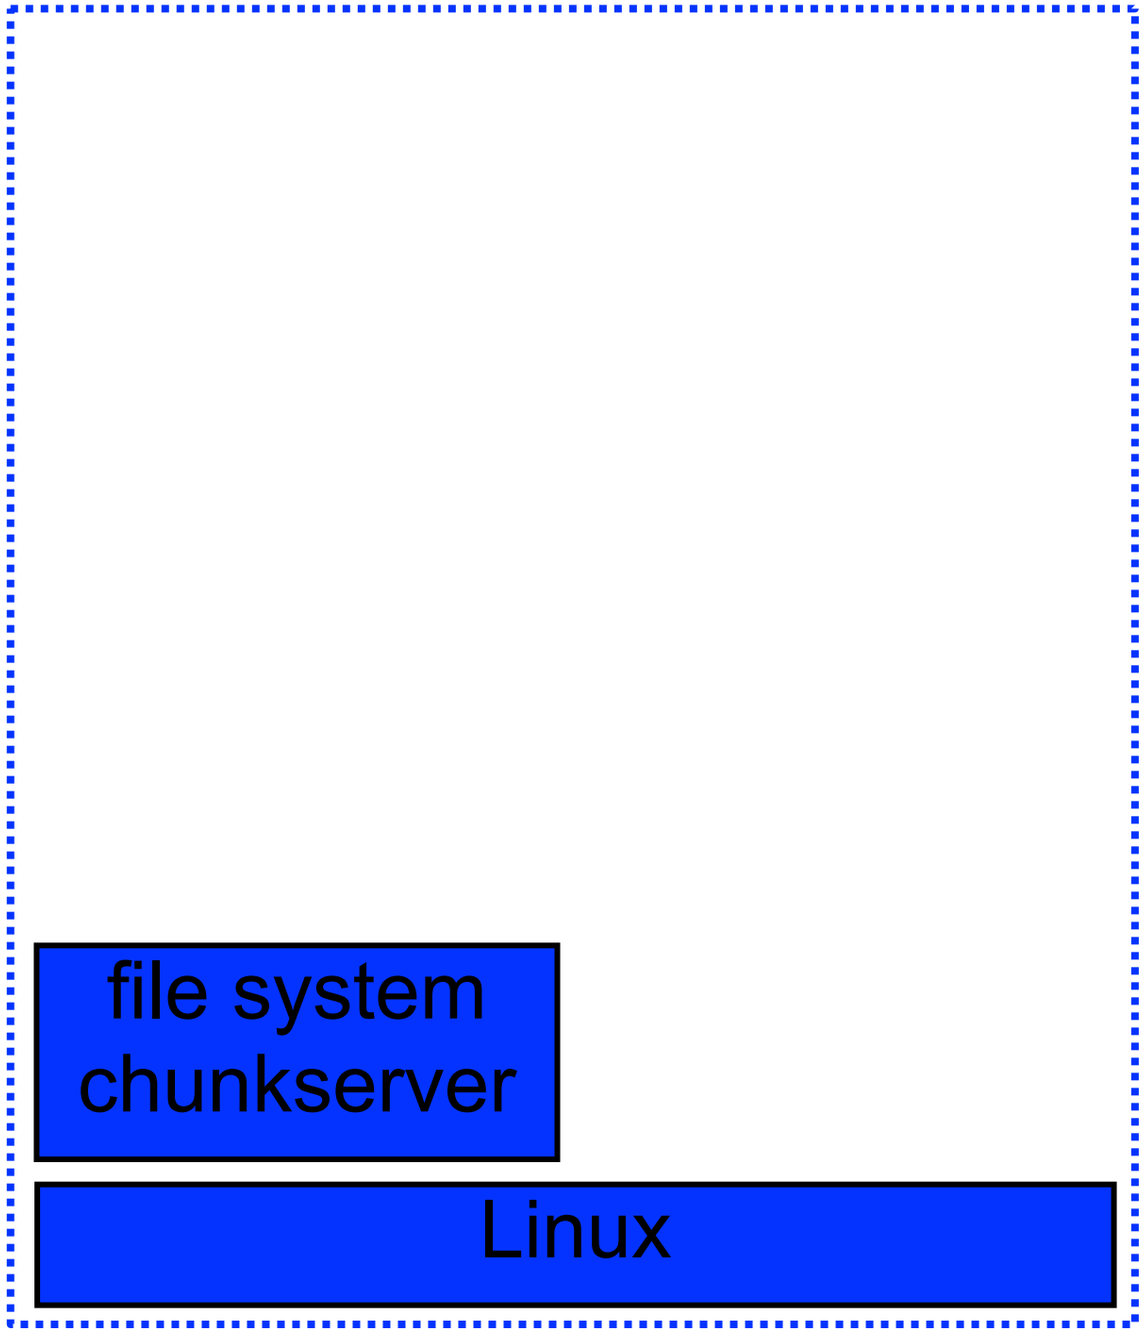

Google

Shared Environment

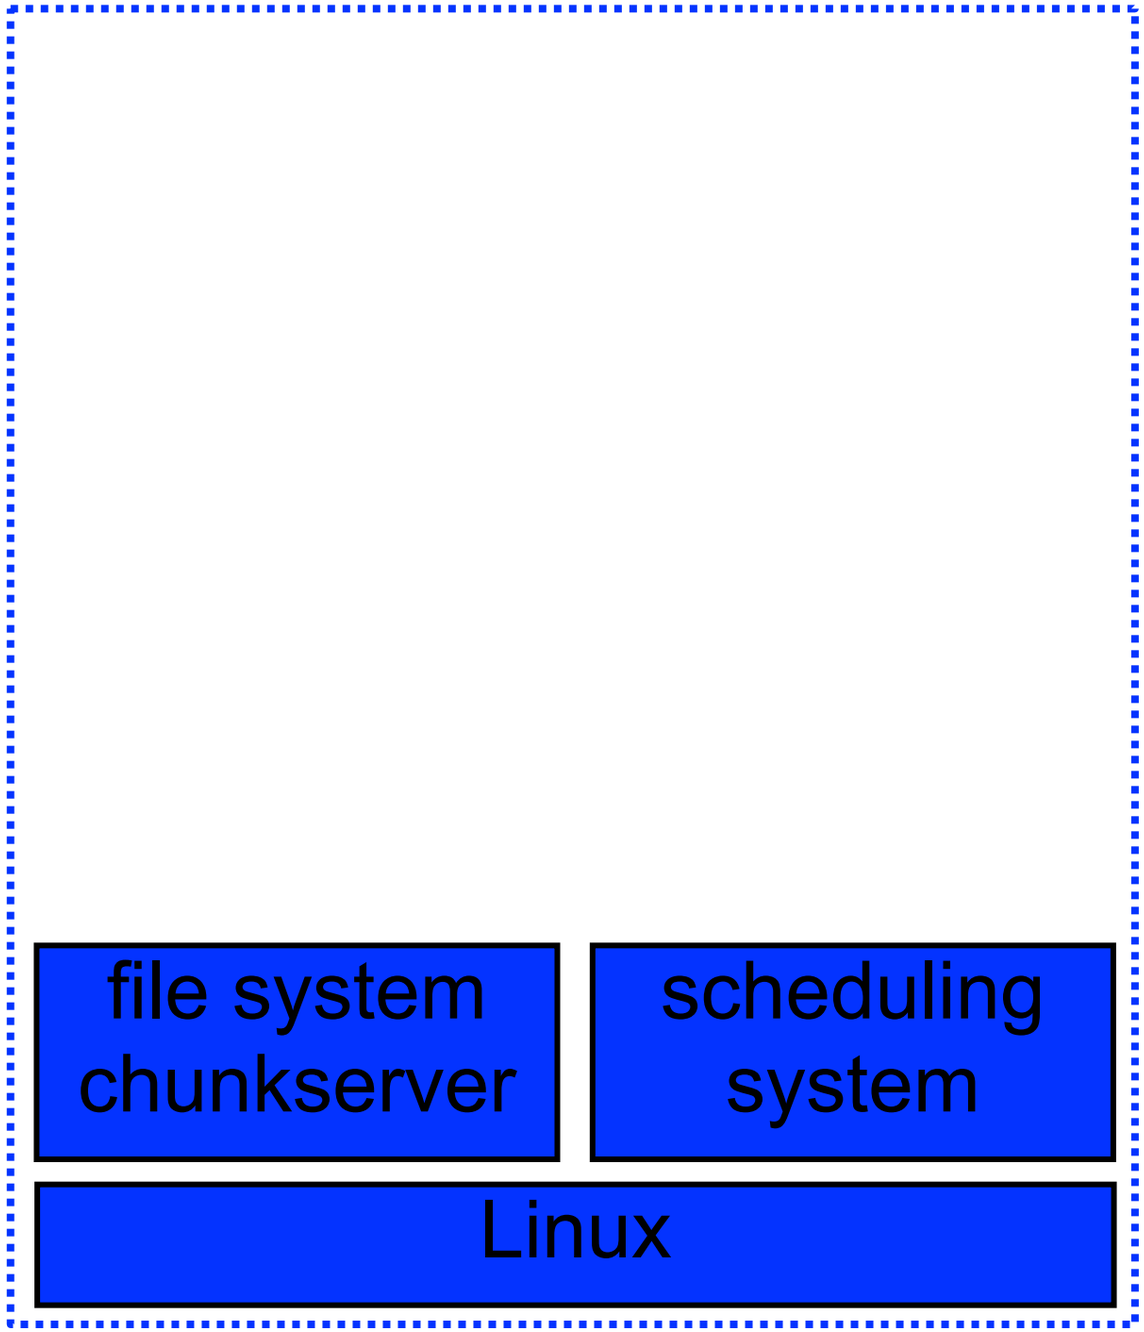

Google

Shared Environment

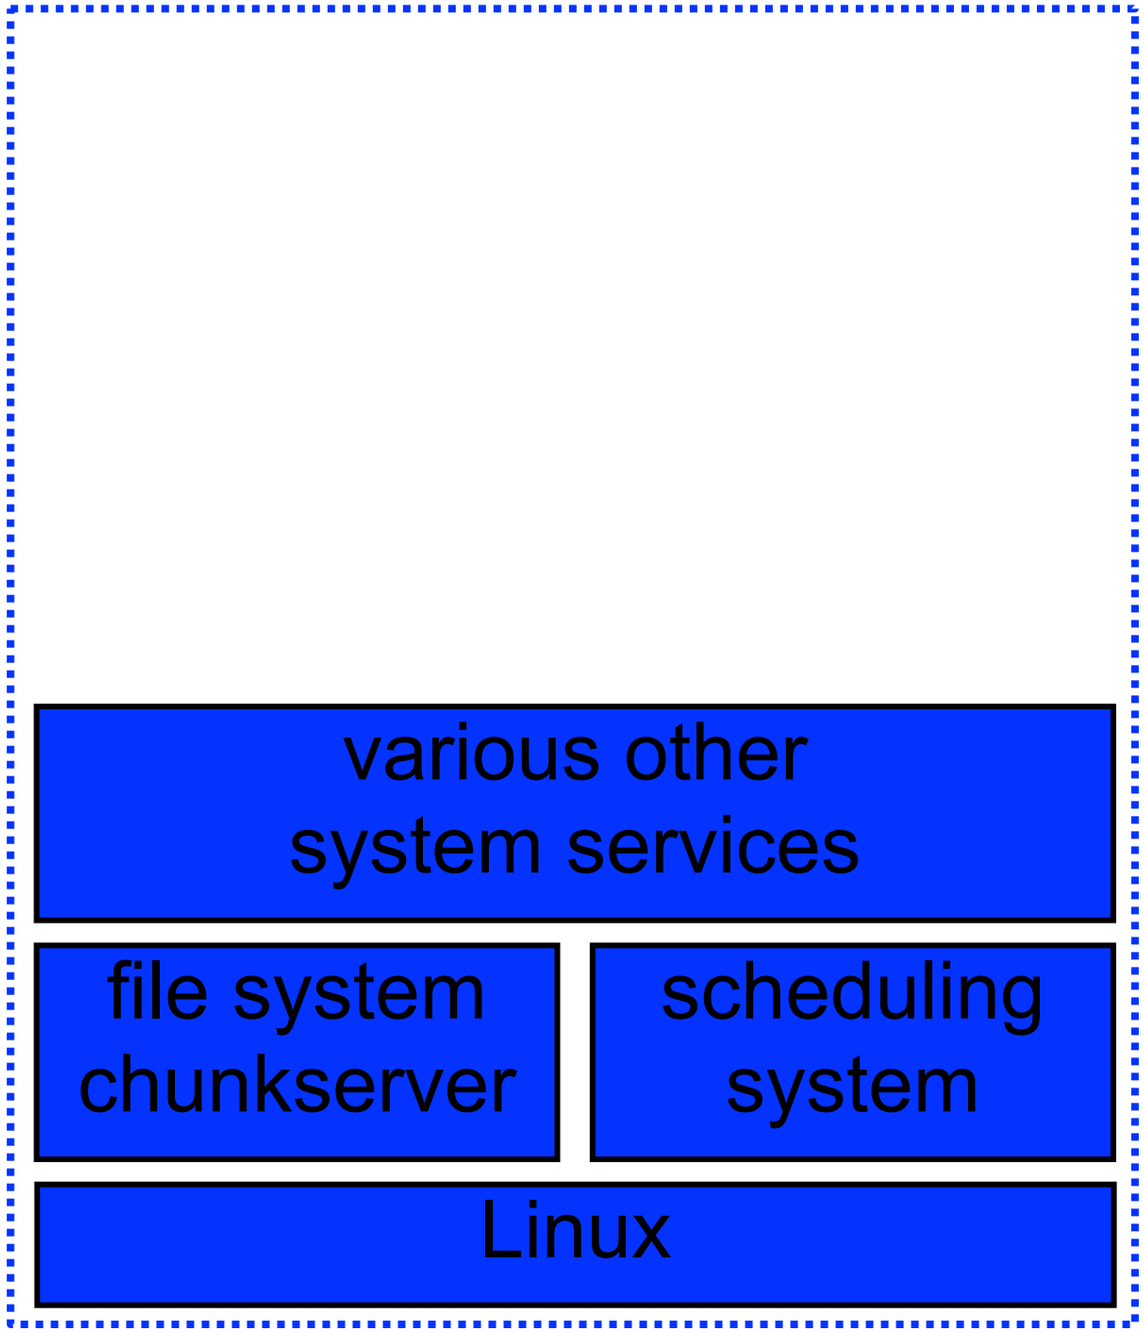

Google

Shared Environment

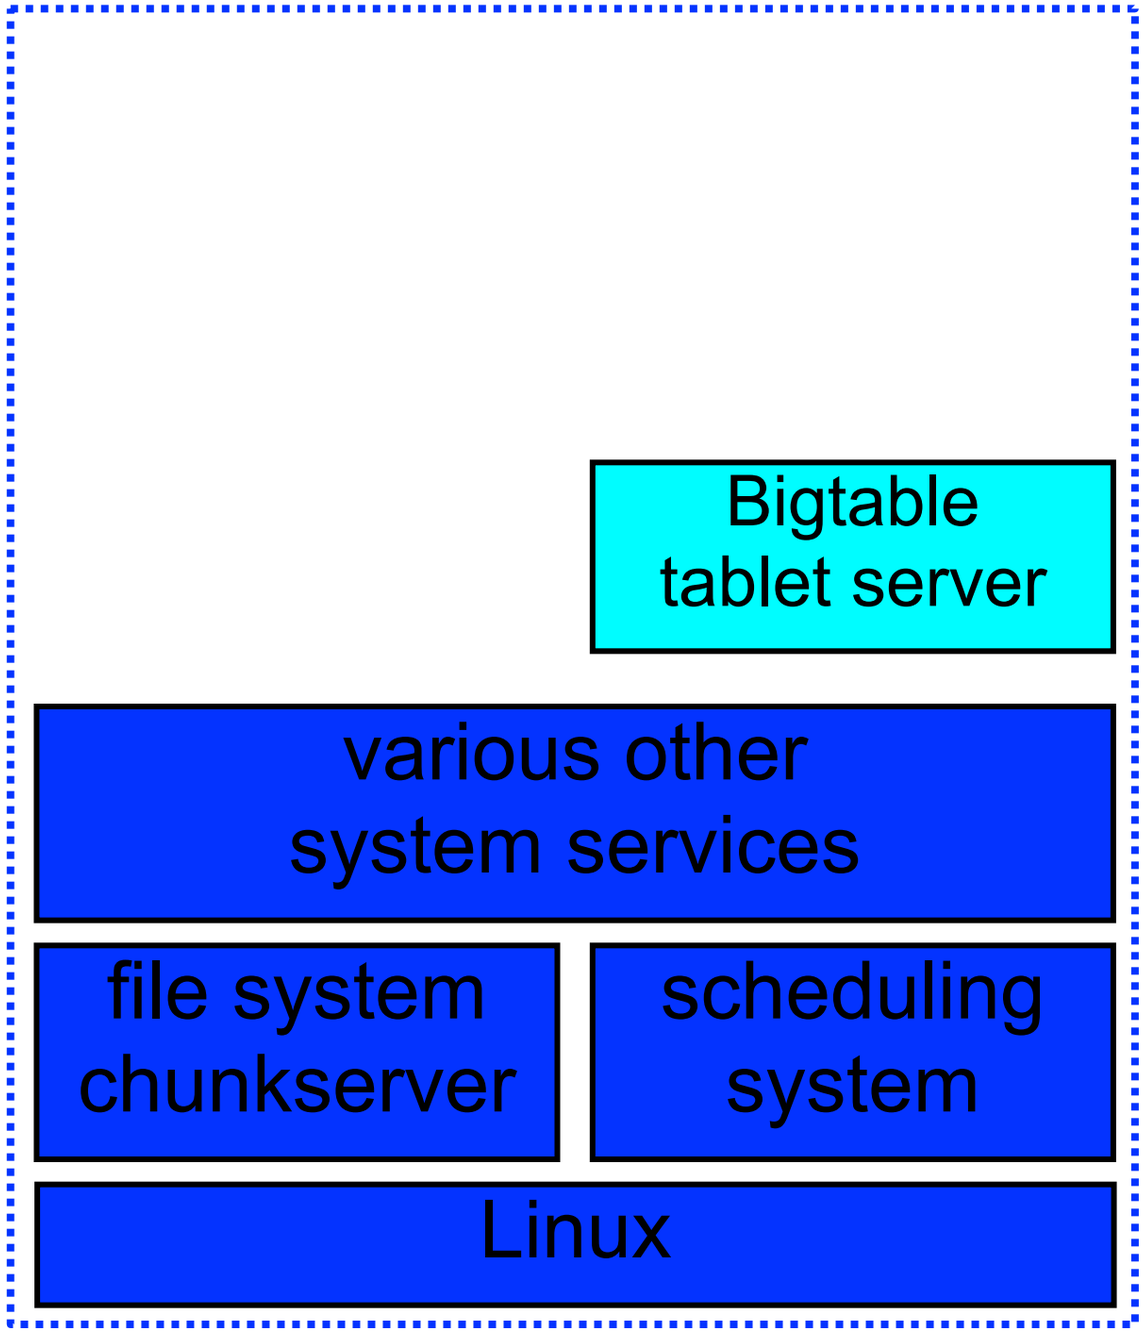

Google

Shared Environment

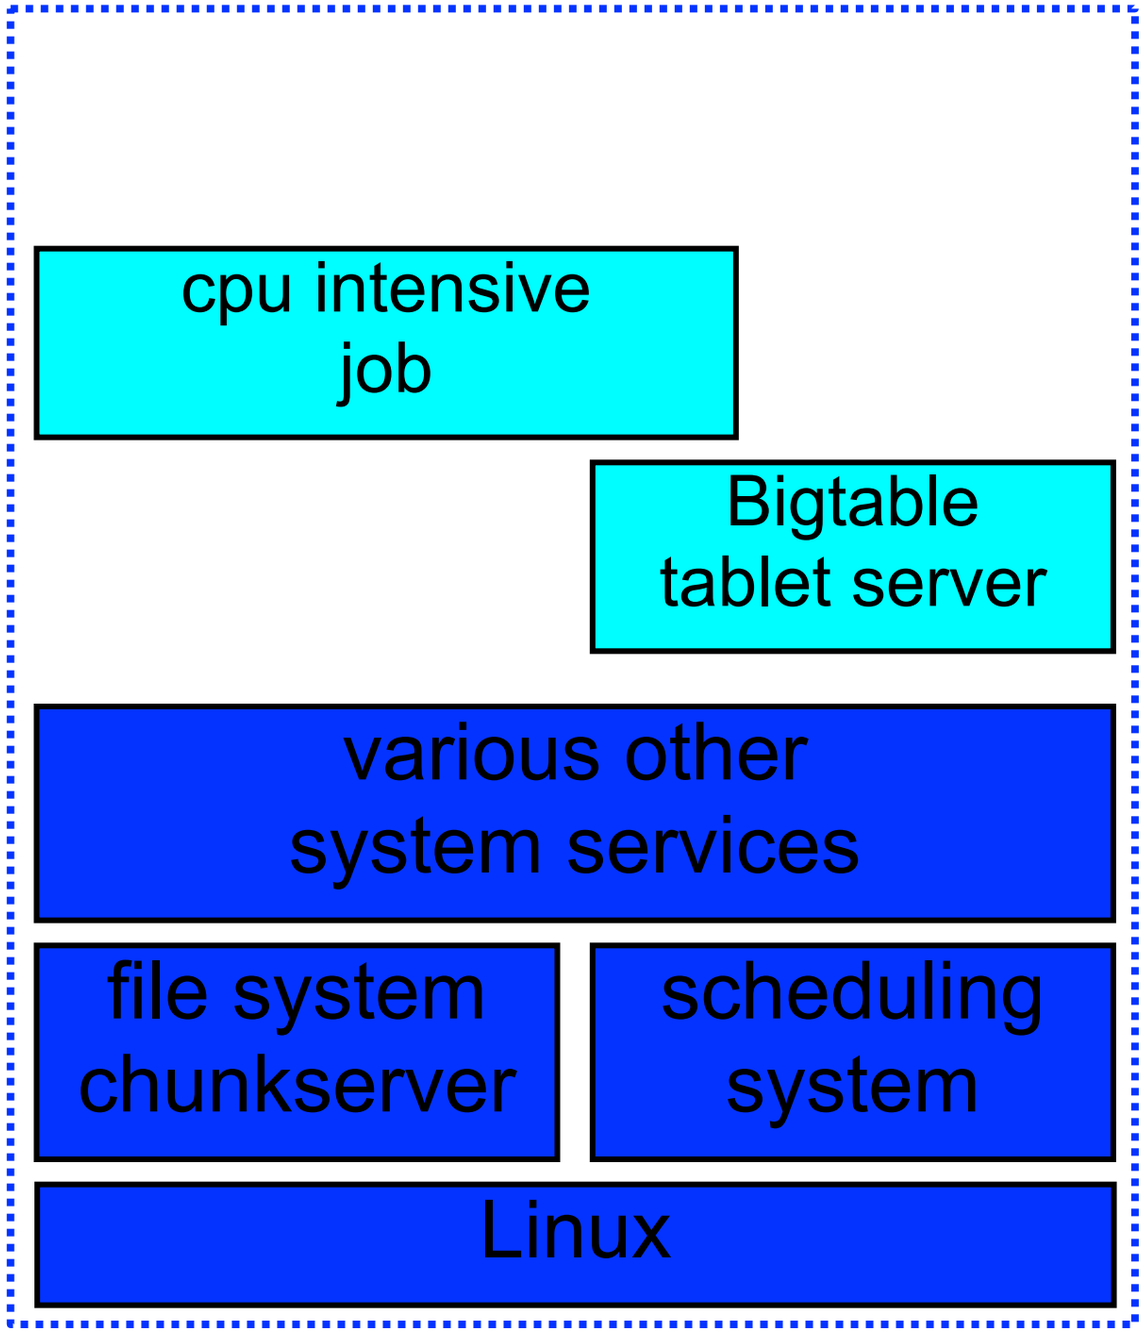

Google

Shared Environment

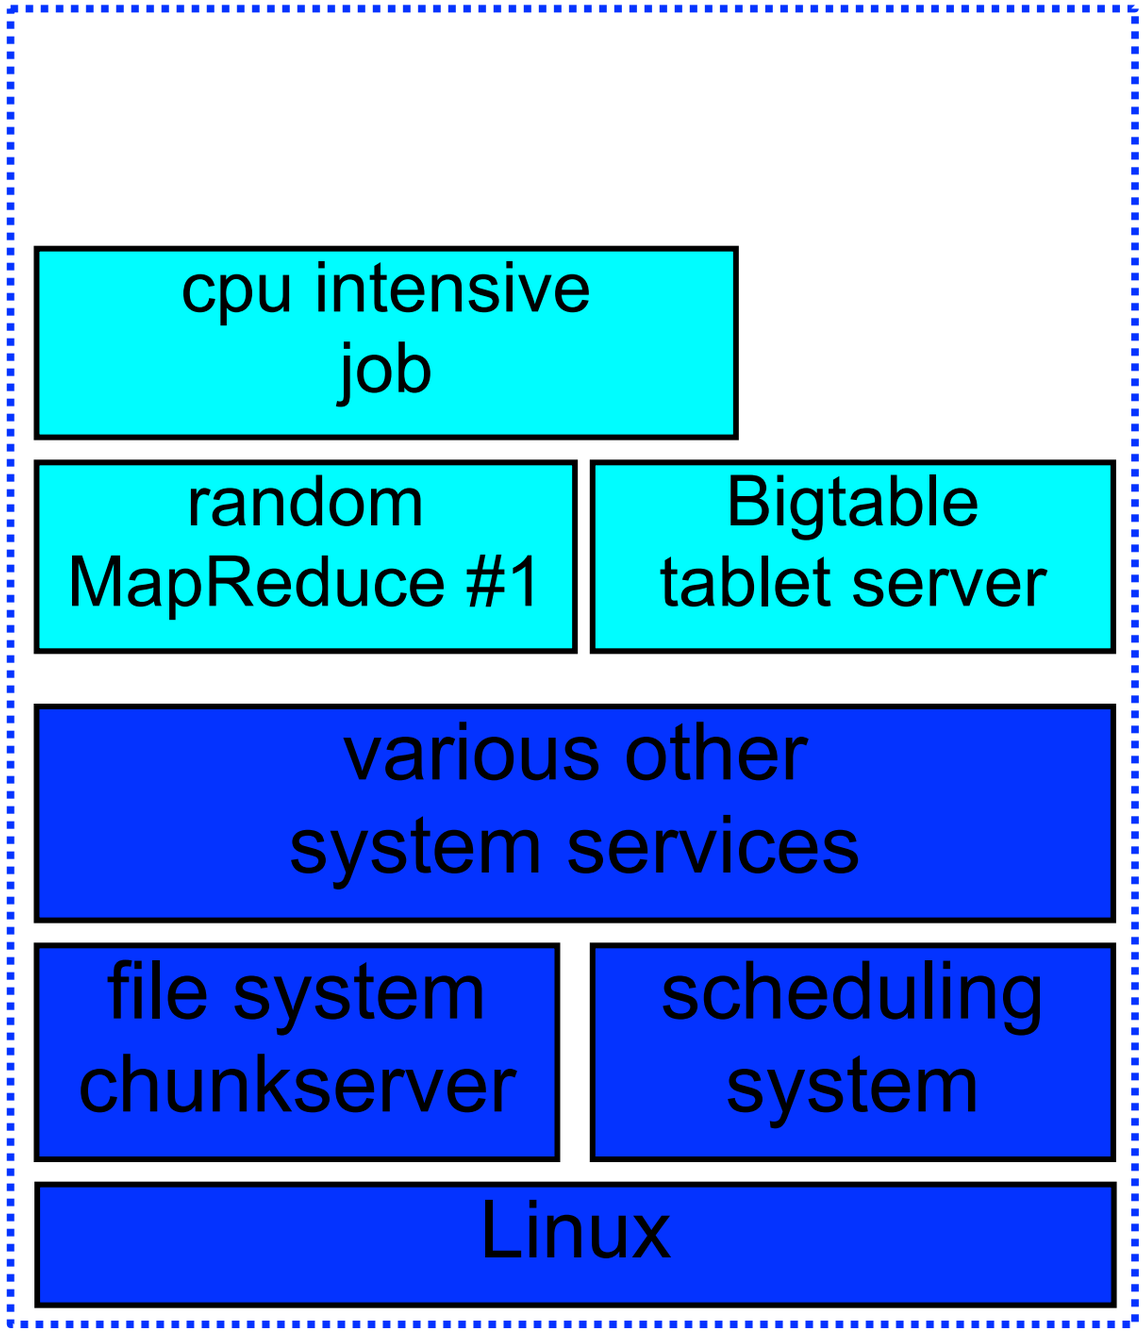

Google

Shared Environment

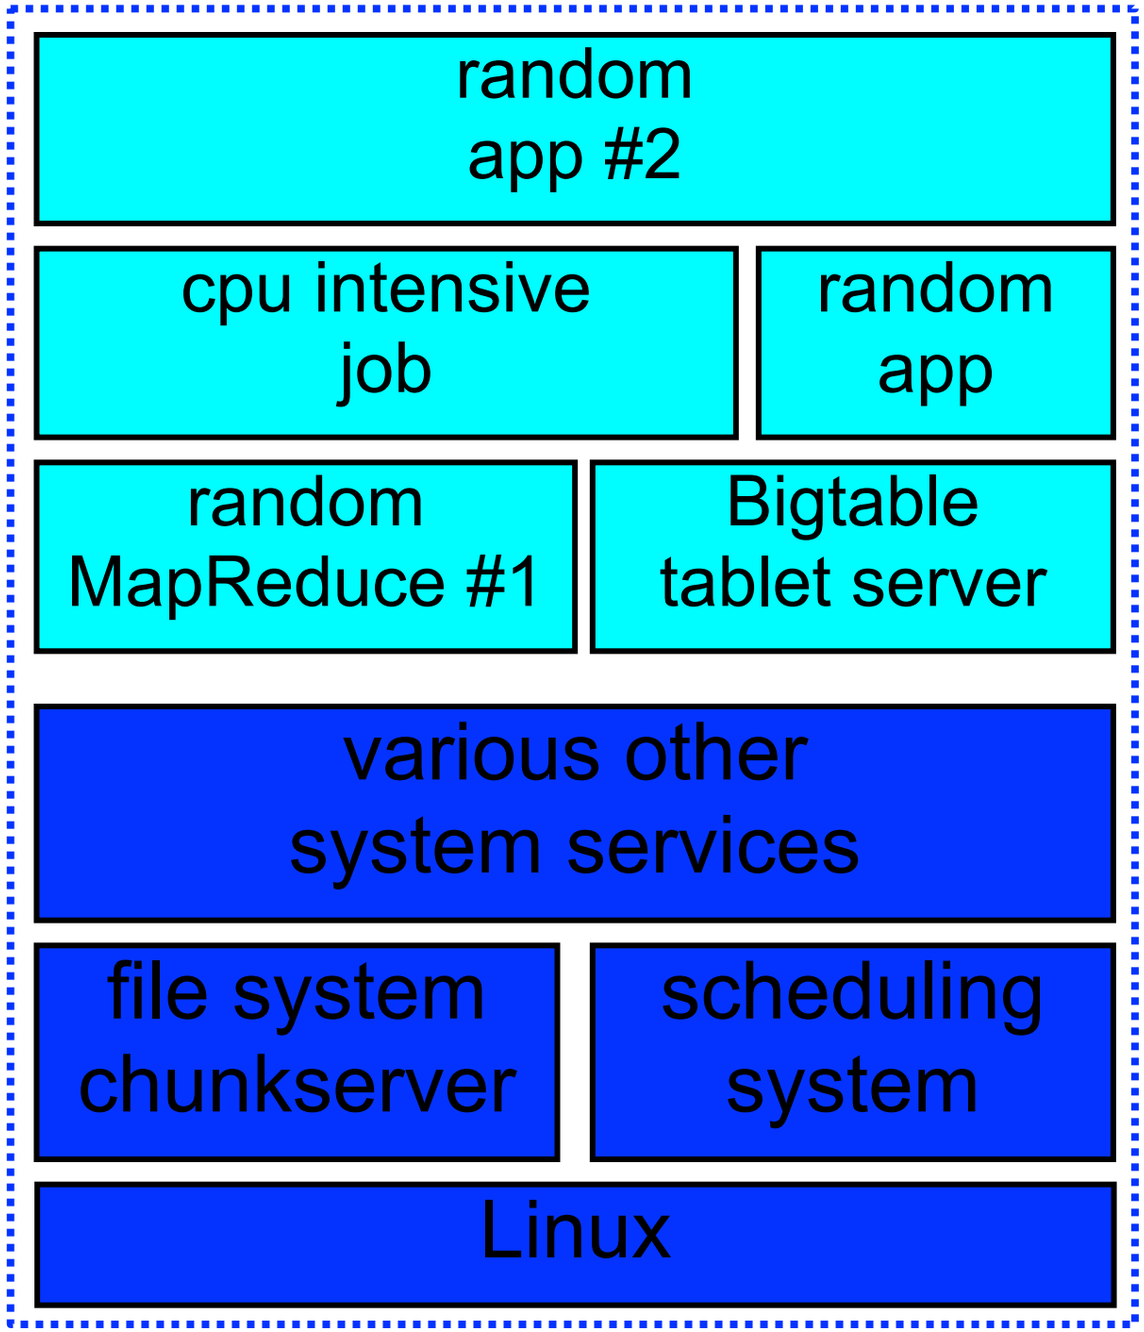

Google

Shared Environment

- Huge benefit: greatly increased utilization
- ... but hard to predict effects increase variability
  - network congestion
  - background activities
  - bursts of foreground activity
  - not just your jobs, but everyone else's jobs, too
  - not static: change happening constantly
- Exacerbated by large fanout systems

Google

The Problem with Shared Environments

Google

The Problem with Shared Environments

Google

The Problem with Shared Environments

- Server with 10 ms avg. but 1 sec 99%ile latency
  - touch 1 of these: 1% of requests take ≥1 sec
  - touch 100 of these: 63% of requests take ≥1 sec

Google

Tolerating Faults vs. Tolerating Variability

- Tolerating faults:
  - rely on extra resources
    - RAIDed disks, ECC memory, dist. system components, etc.
  - ***make a reliable whole out of unreliable parts***

- Tolerating variability:
  - use these same extra resources
  - ***make a predictable whole out of unpredictable parts***

- Times scales are very different:
  - variability: 1000s of disruptions/sec, scale of **milliseconds**
  - faults: 10s of failures per day, scale of **tens of seconds**

Google

Latency Tolerating Techniques

- Cross request adaptation
  - examine recent behavior
  - take action to improve latency of future requests
  - typically relate to balancing load across set of servers
  - time scale: 10s of seconds to minutes

- Within request adaptation
  - cope with slow subsystems in context of higher level request
  - time scale: right now, while user is waiting

Google

Backup Requests w/ Cross-Server Cancellation

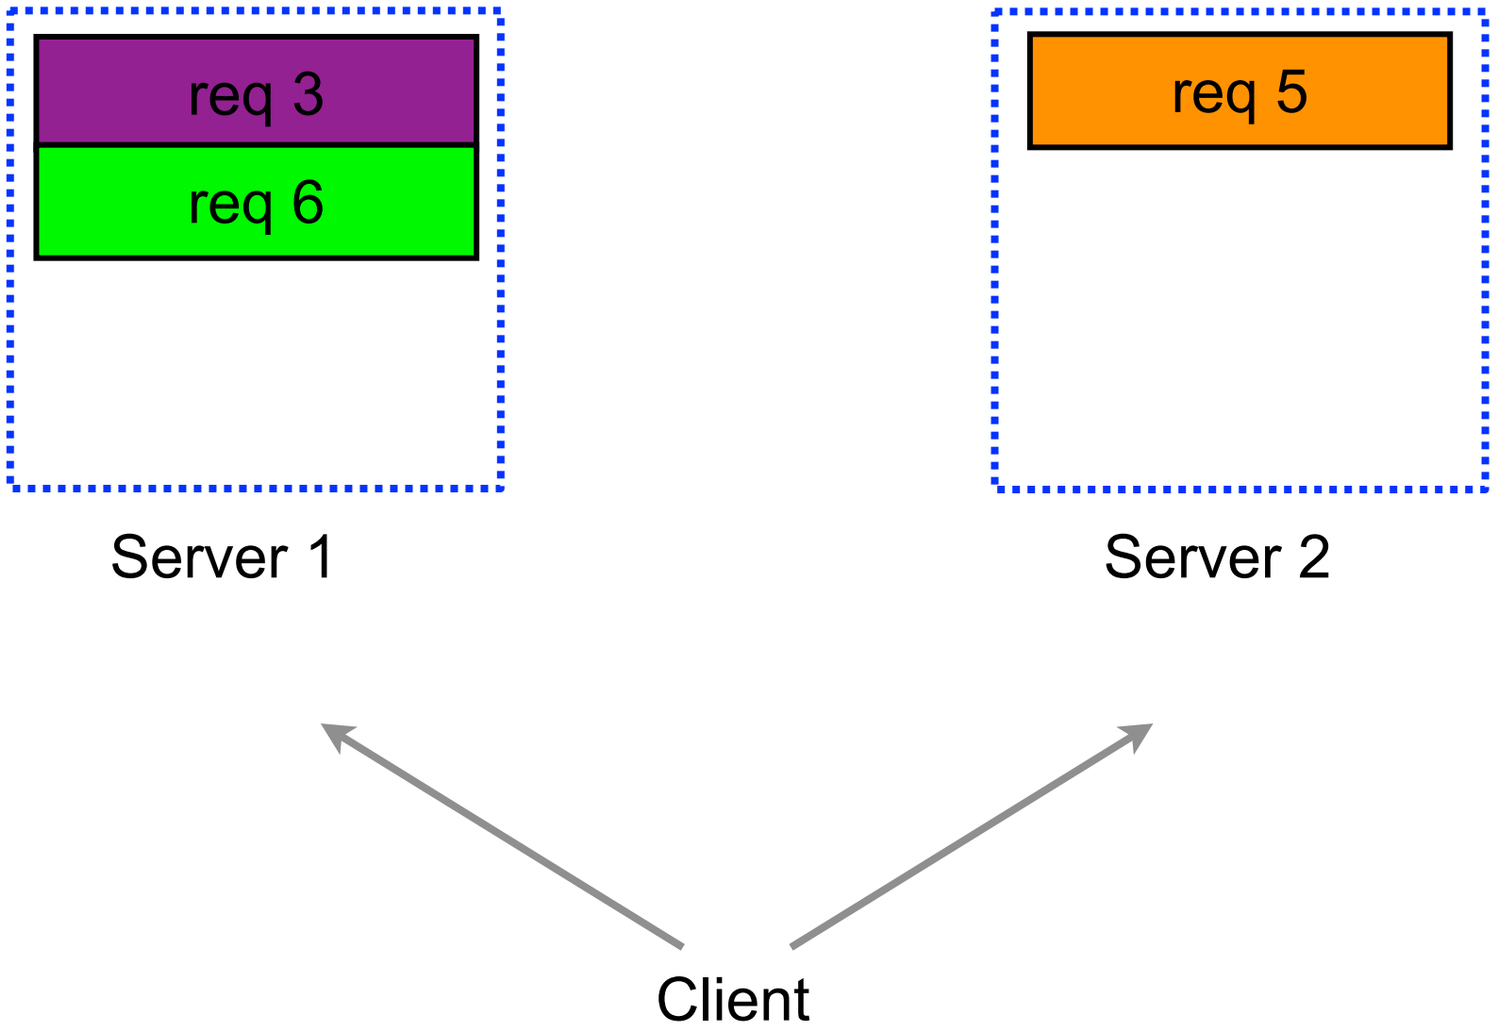

Similar to Michael Mitzenmacher's work on "The Power of Two Choices", except send to both, rather than just picking "best" one

Google

Backup Requests w/ Cross-Server Cancellation

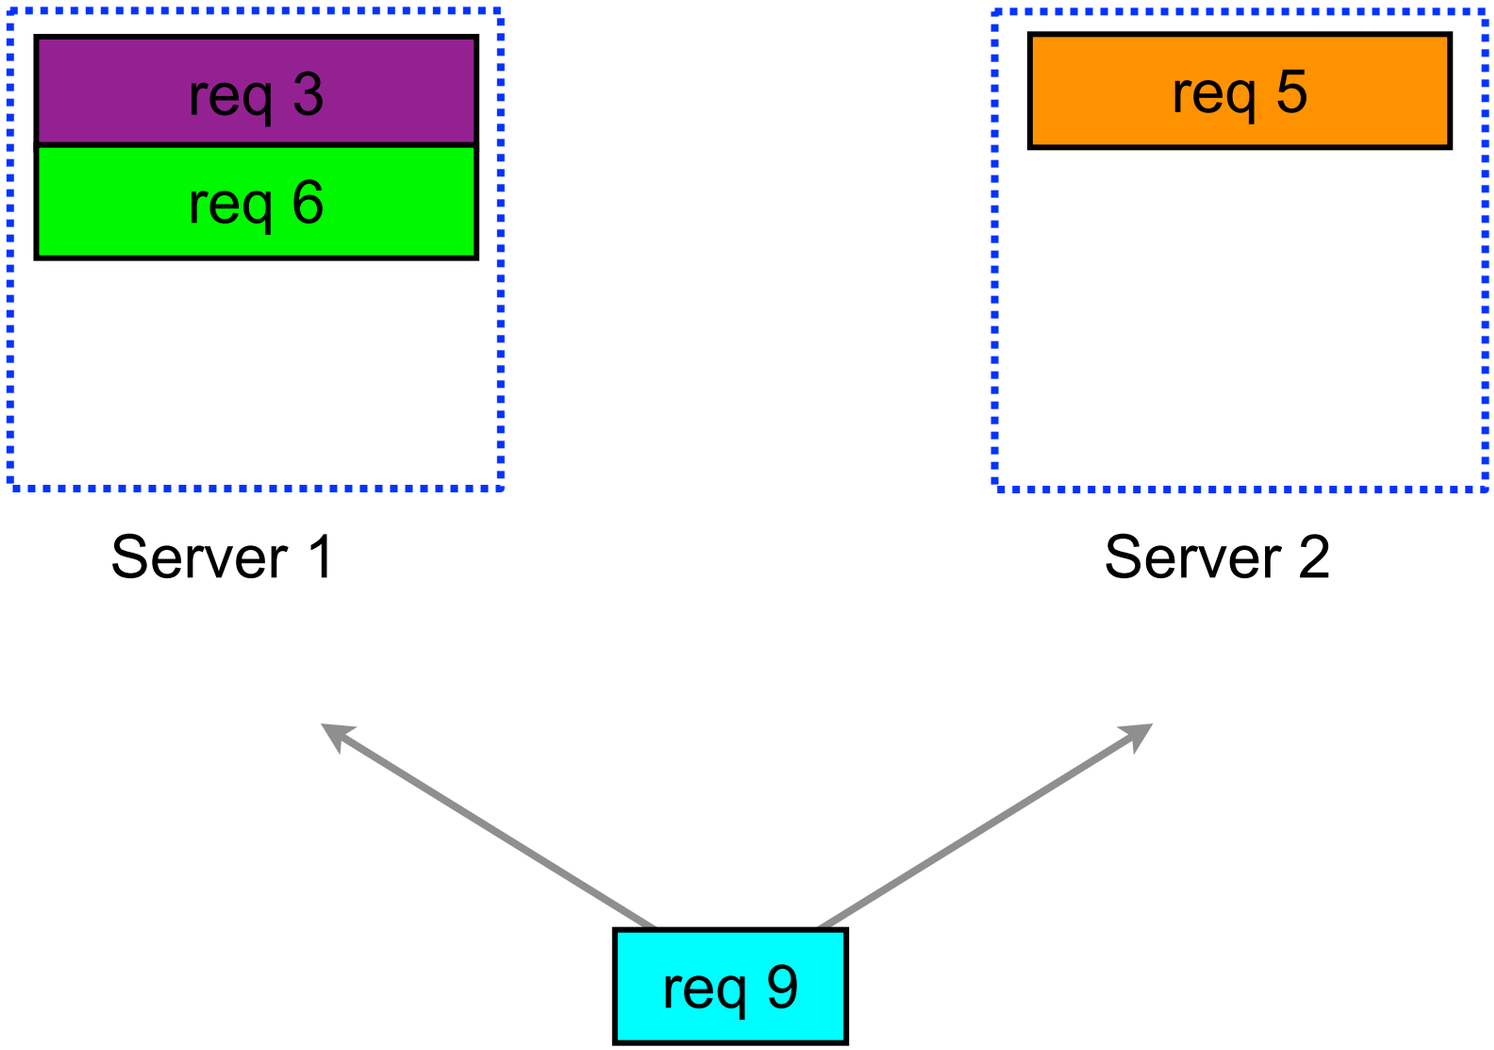

Similar to Michael Mitzenmacher's work on "The Power of Two Choices", except send to both, rather than just picking "best" one

Google

Backup Requests w/ Cross-Server Cancellation

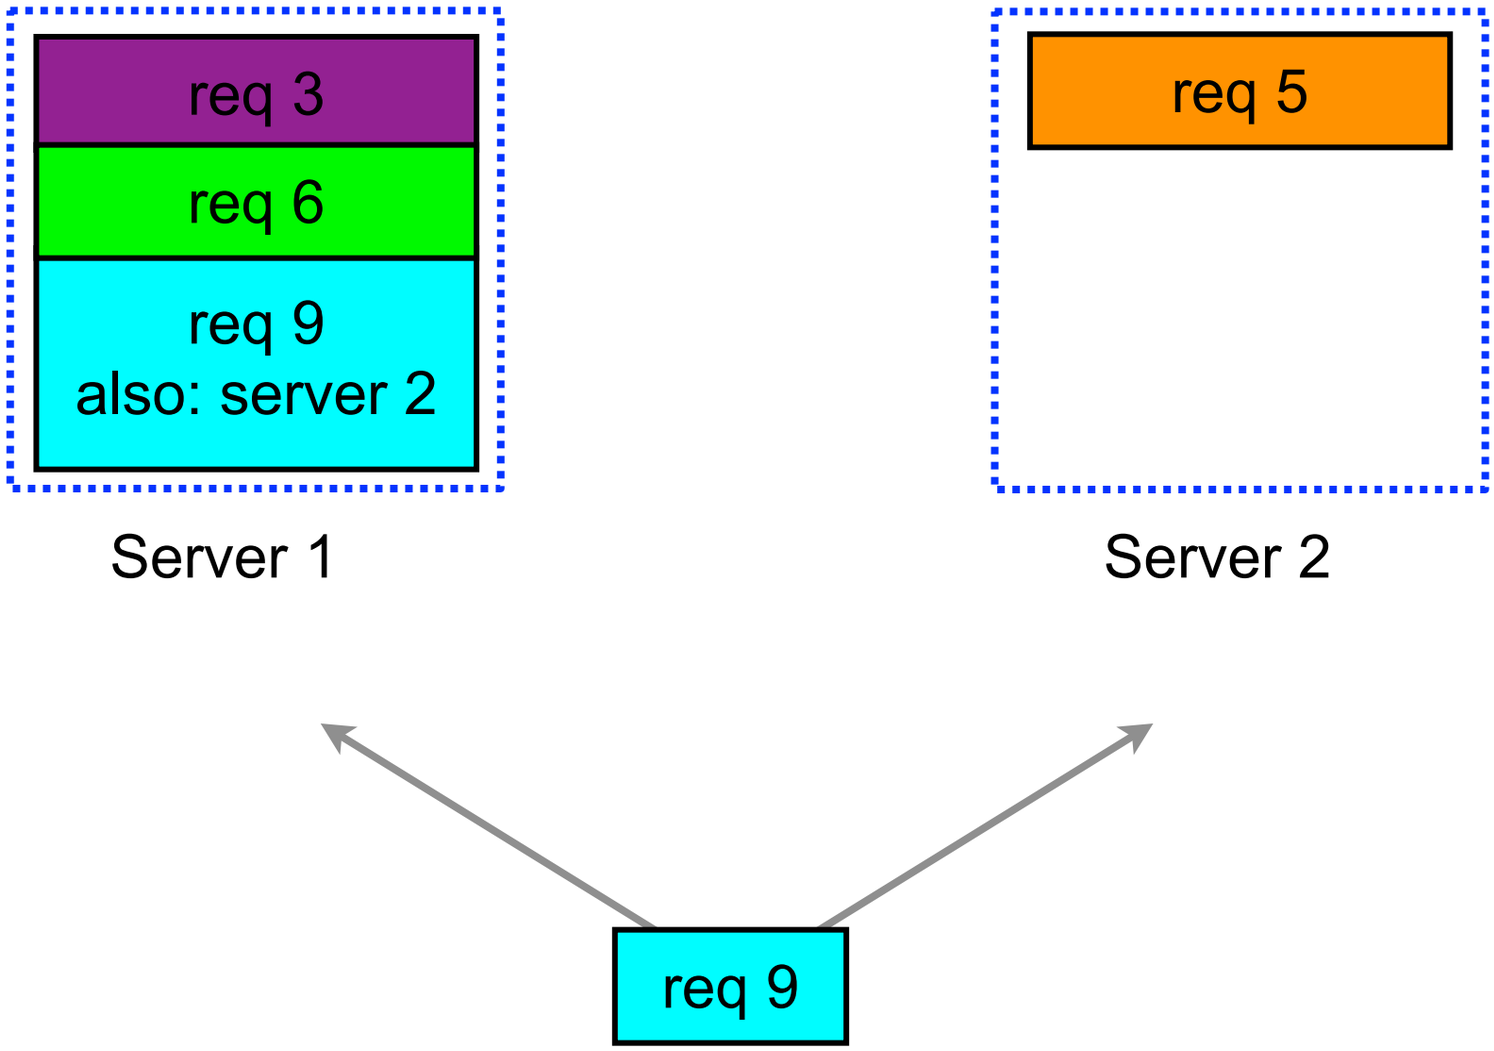

Similar to Michael Mitzenmacher's work on "The Power of Two Choices", except send to both, rather than just picking "best" one

Each request identifies other server(s) to which request might be sent

Google

Backup Requests w/ Cross-Server Cancellation

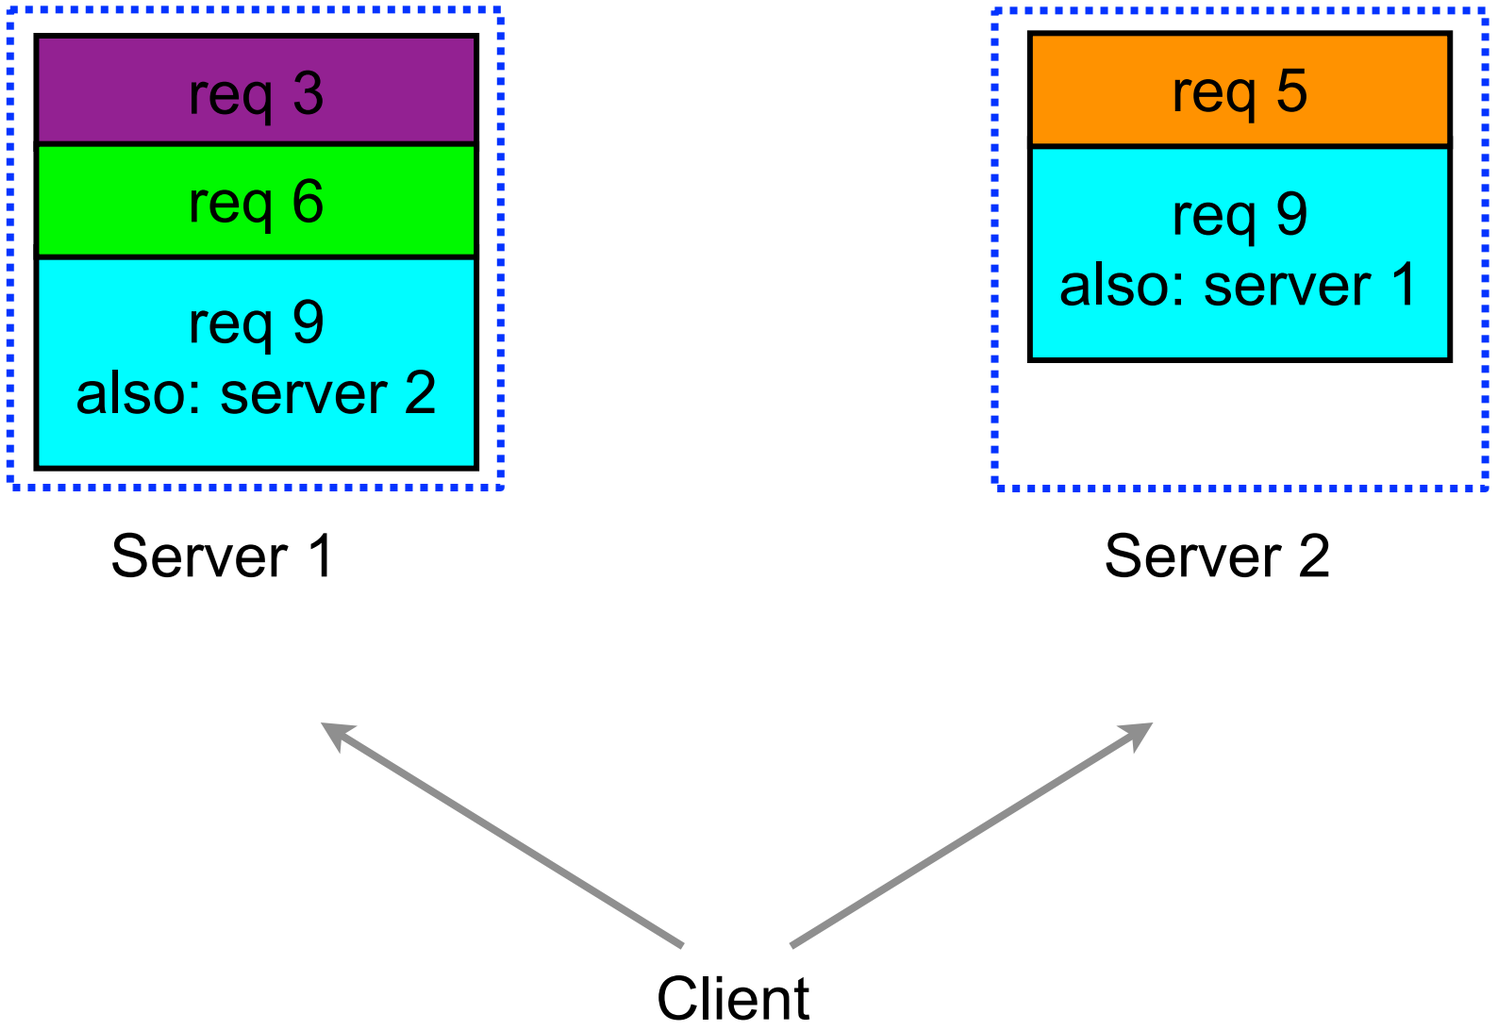

Similar to Michael Mitzenmacher's work on "The Power of Two Choices", except send to both, rather than just picking "best" one

Each request identifies other server(s) to which request might be sent

Google

Backup Requests w/ Cross-Server Cancellation

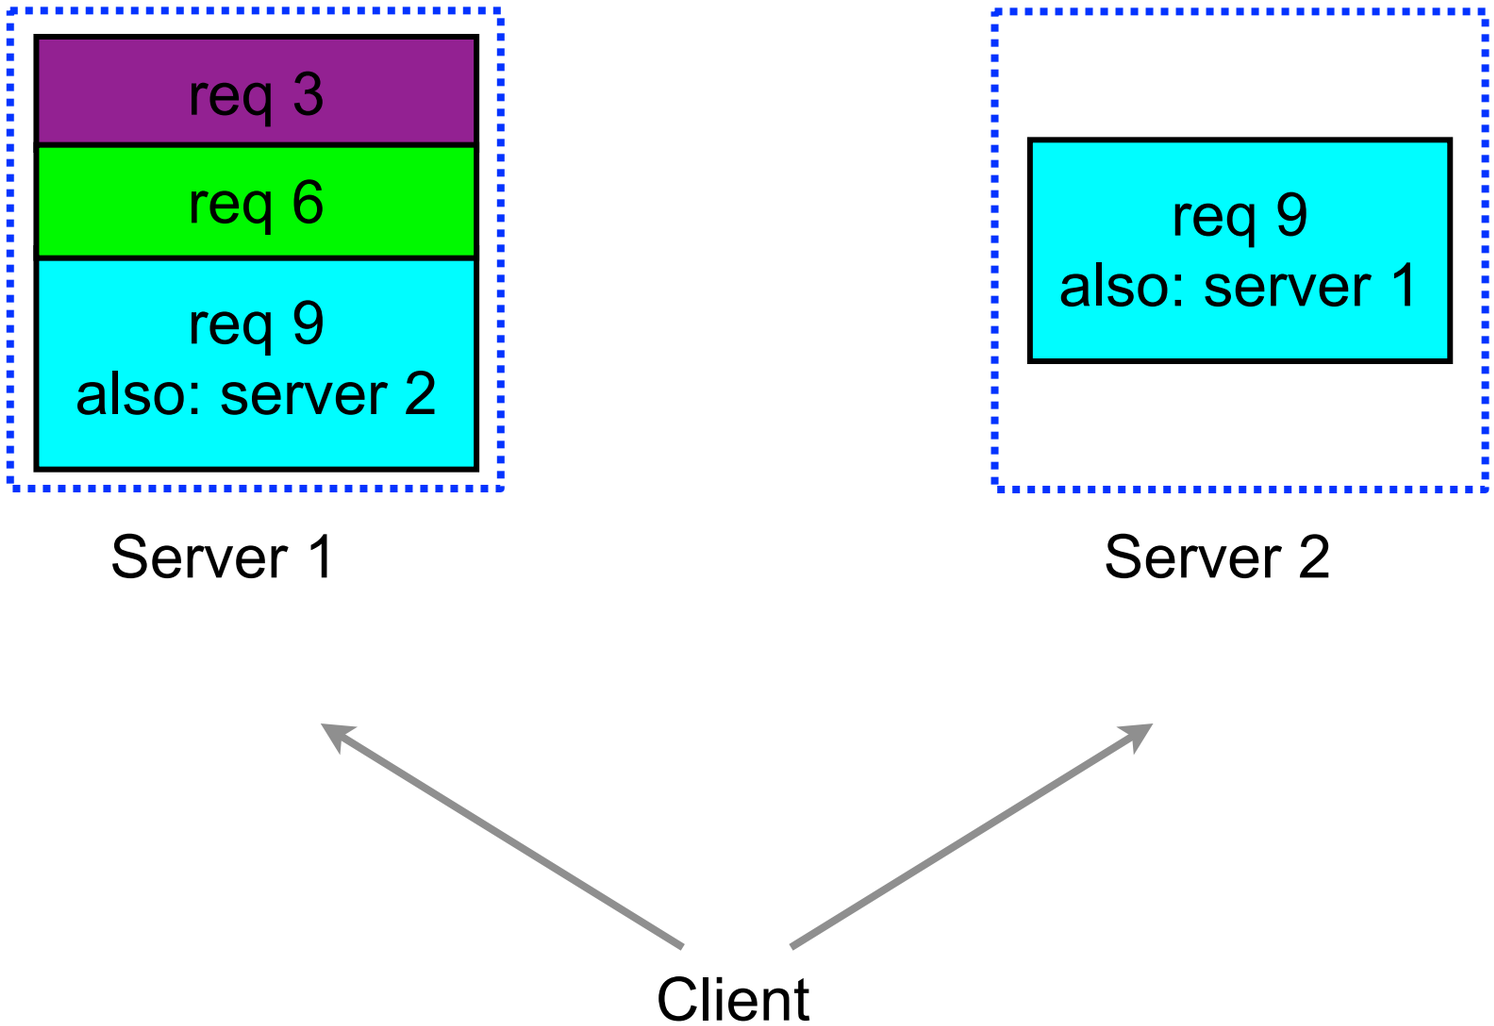

Similar to Michael Mitzenmacher's work on "The Power of Two Choices", except send to both, rather than just picking "best" one

Each request identifies other server(s) to which request might be sent

Google

Backup Requests w/ Cross-Server Cancellation

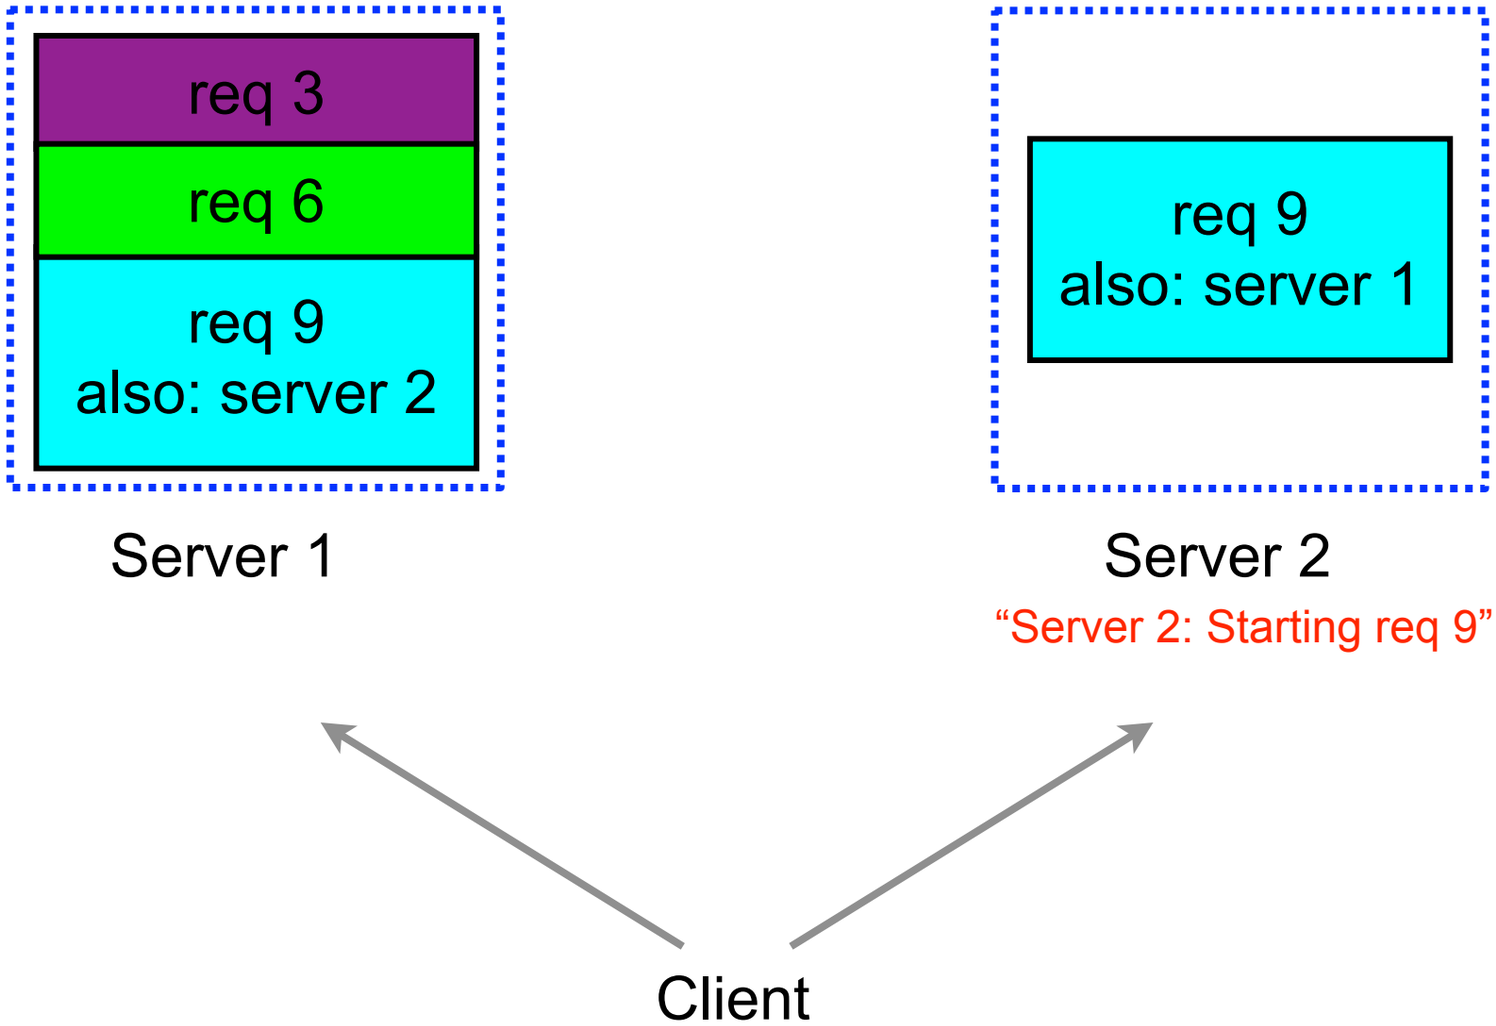

Similar to Michael Mitzenmacher's work on "The Power of Two Choices", except send to both, rather than just picking "best" one

Each request identifies other server(s) to which request might be sent

Google

Backup Requests w/ Cross-Server Cancellation

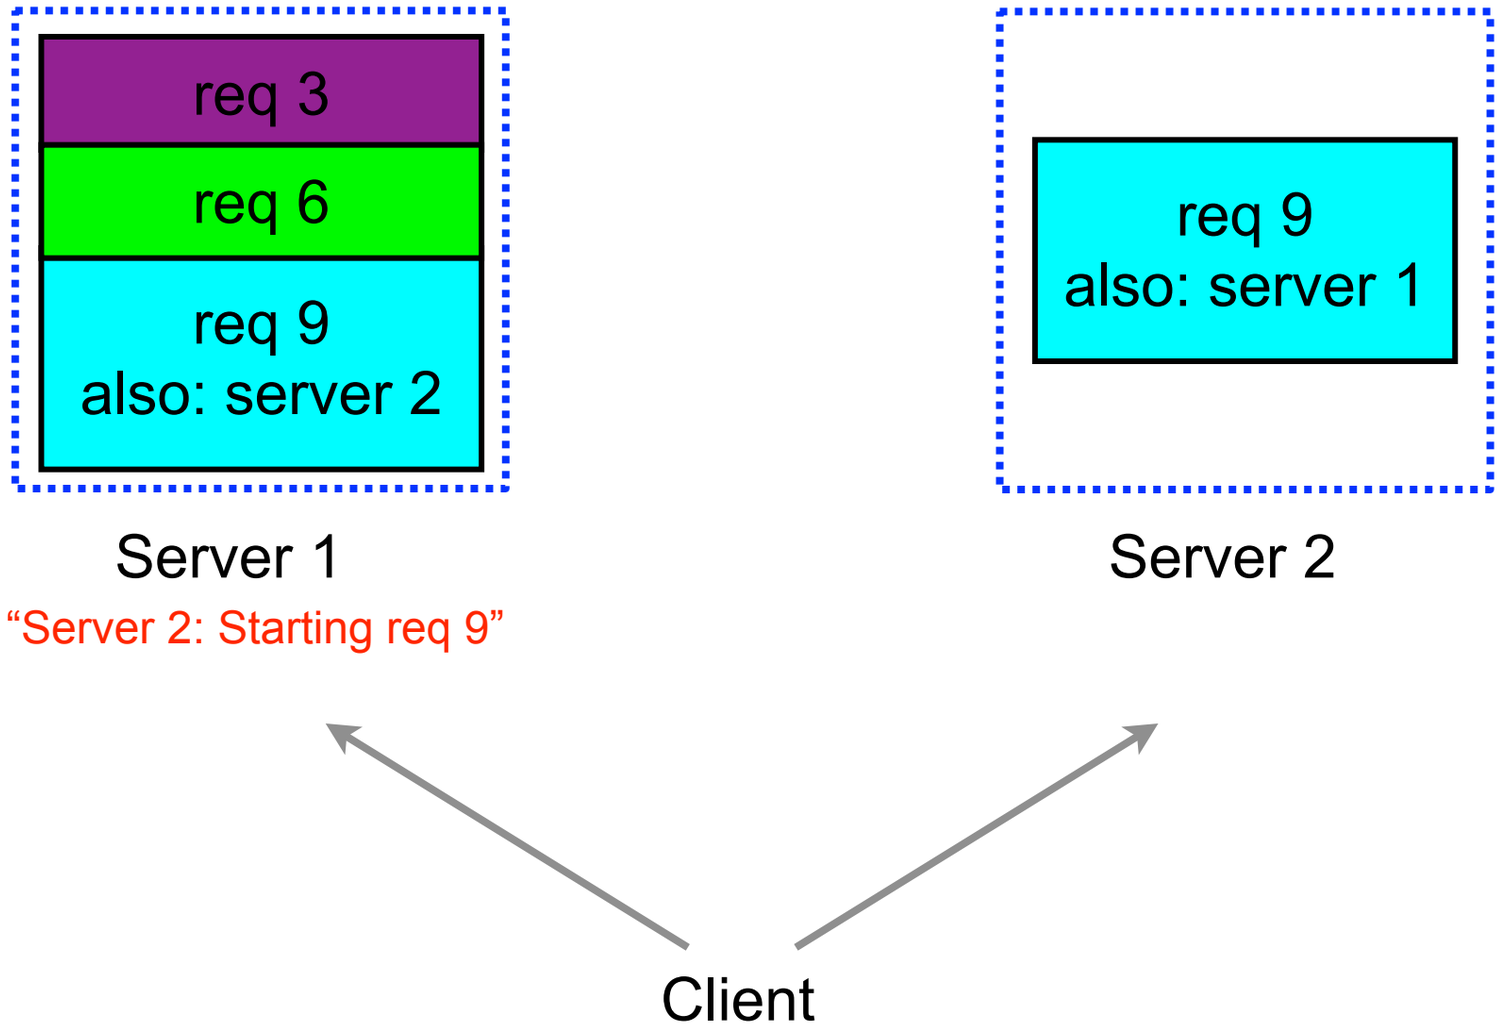

Similar to Michael Mitzenmacher's work on "The Power of Two Choices", except send to both, rather than just picking "best" one

Each request identifies other server(s) to which request might be sent

Google

Backup Requests w/ Cross-Server Cancellation

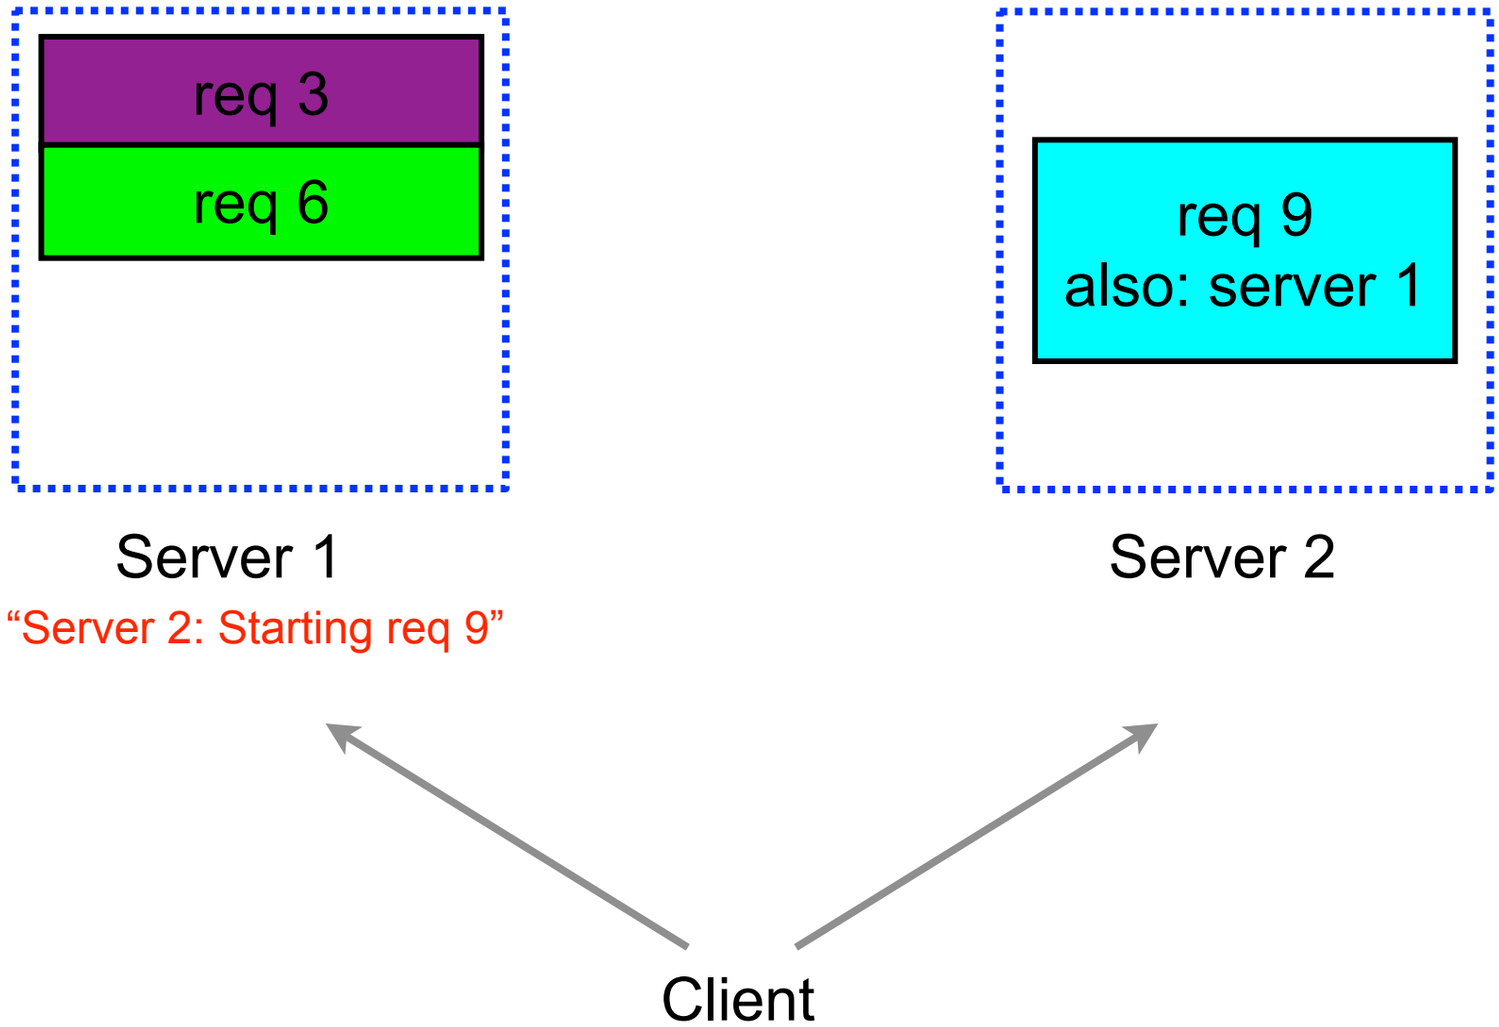

Similar to Michael Mitzenmacher's work on "The Power of Two Choices", except send to both, rather than just picking "best" one

Each request identifies other server(s) to which request might be sent

Google

Backup Requests w/ Cross-Server Cancellation

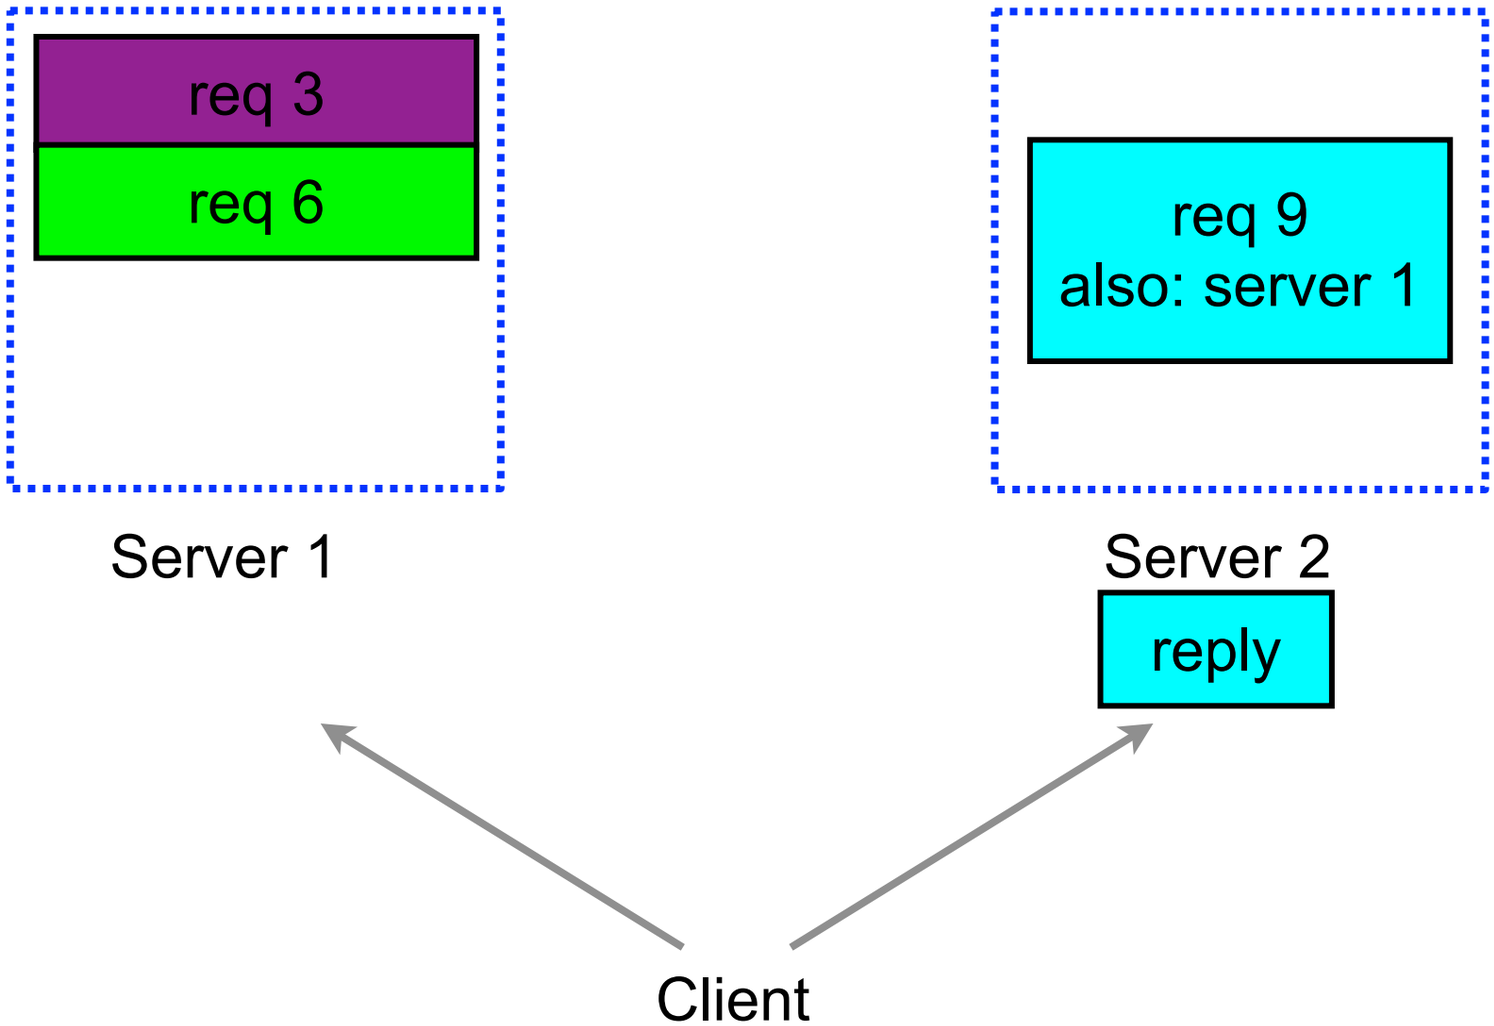

Similar to Michael Mitzenmacher's work on "The Power of Two Choices", except send to both, rather than just picking "best" one

Each request identifies other server(s) to which request might be sent

Backup Requests w/ Cross-Server Cancellation

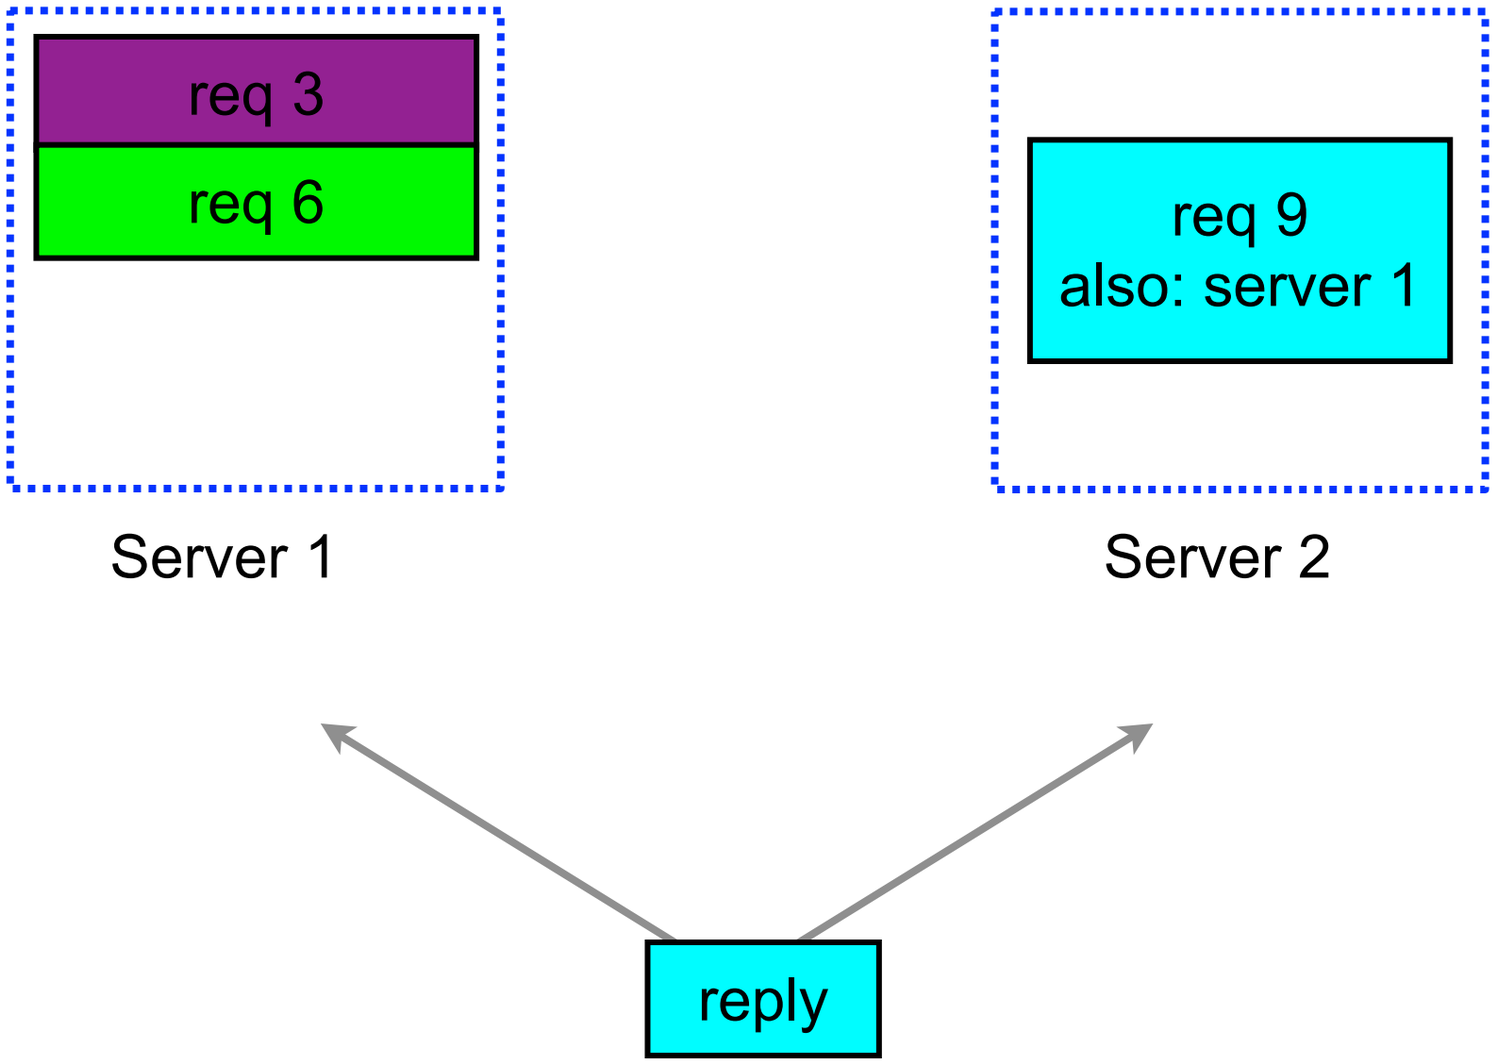

Similar to Michael Mitzenmacher's work on "The Power of Two Choices", except send to both, rather than just picking "best" one

Each request identifies other server(s) to which request might be sent

Google

Backup Requests: Bad Case

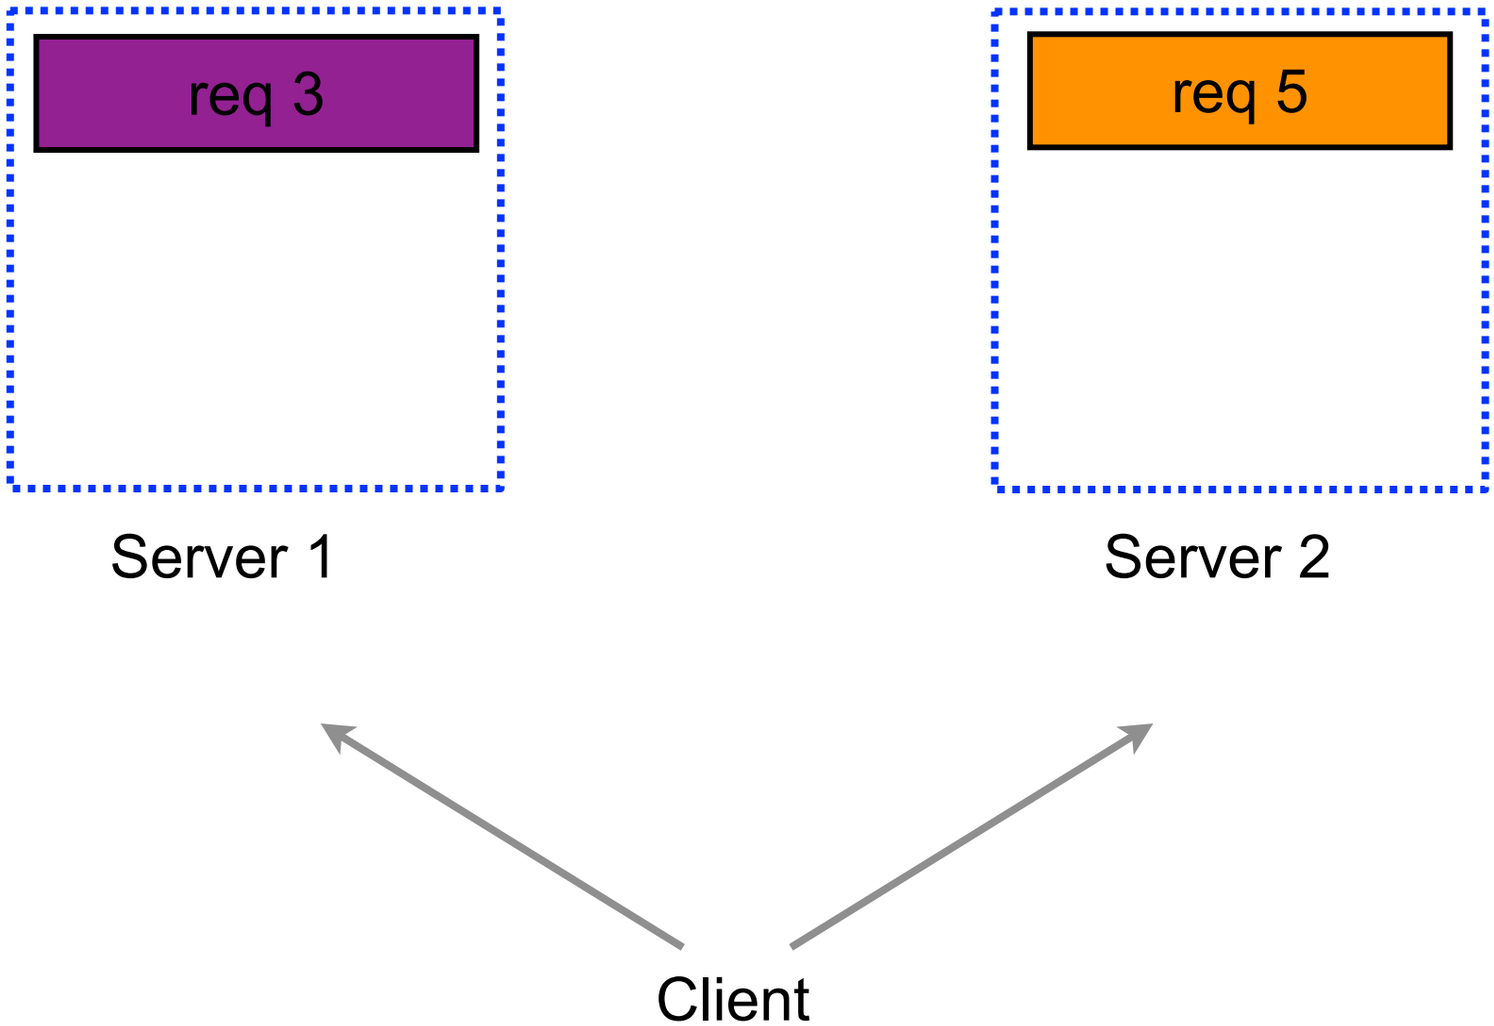

Google

Backup Requests: Bad Case

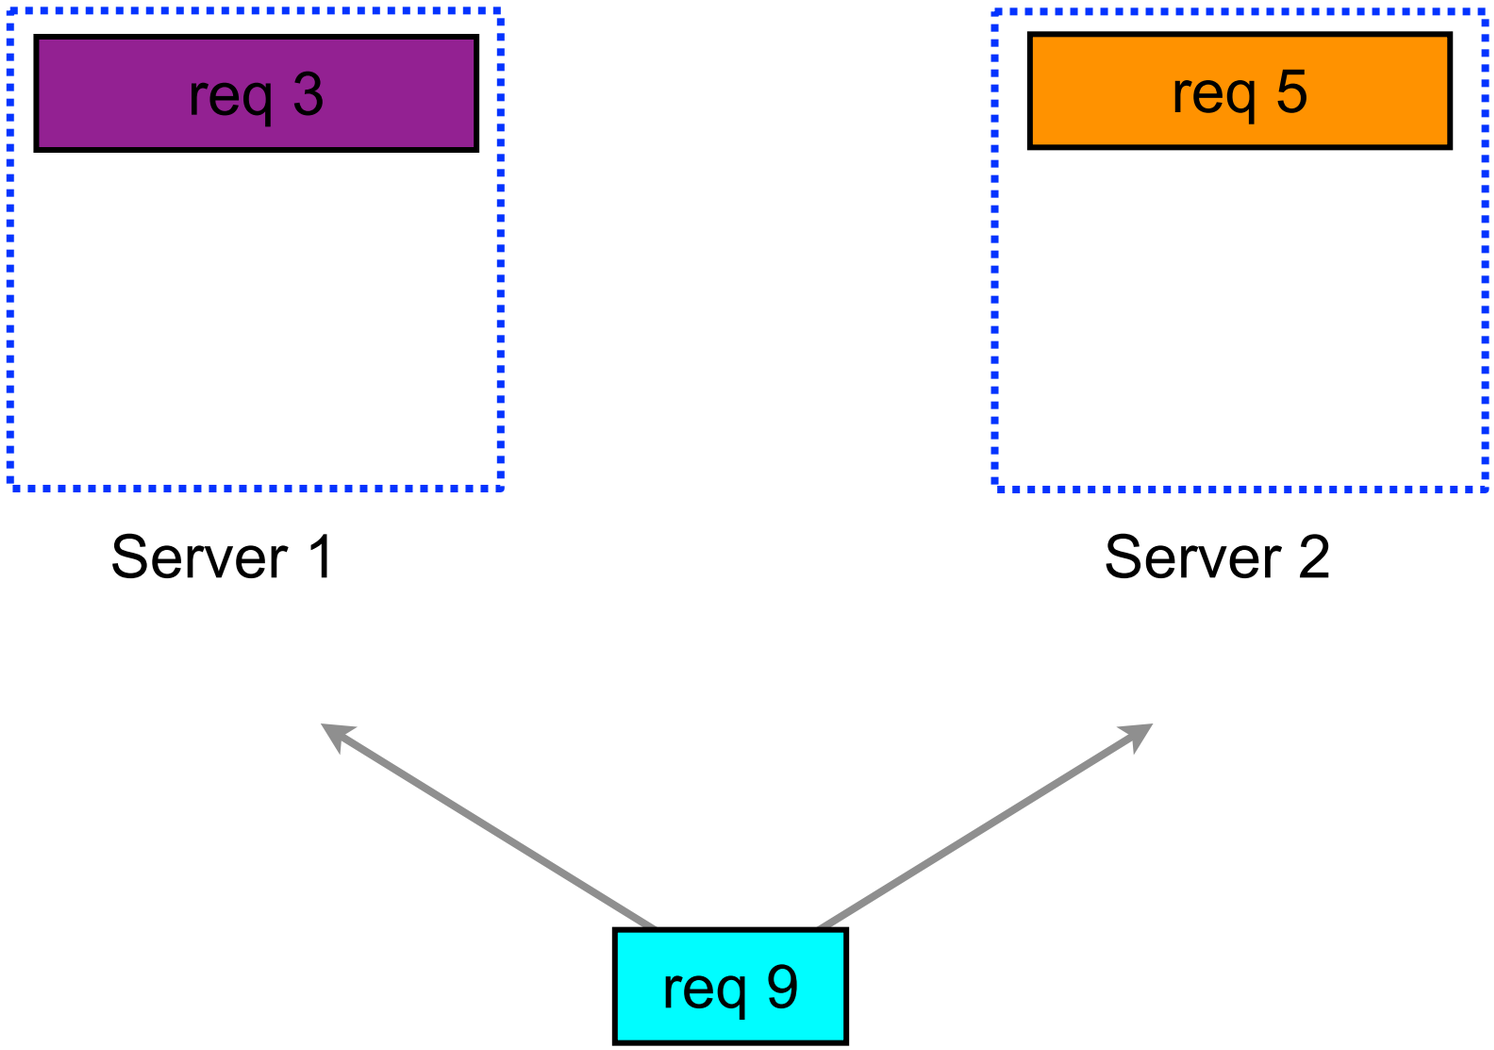

Google

Backup Requests: Bad Case

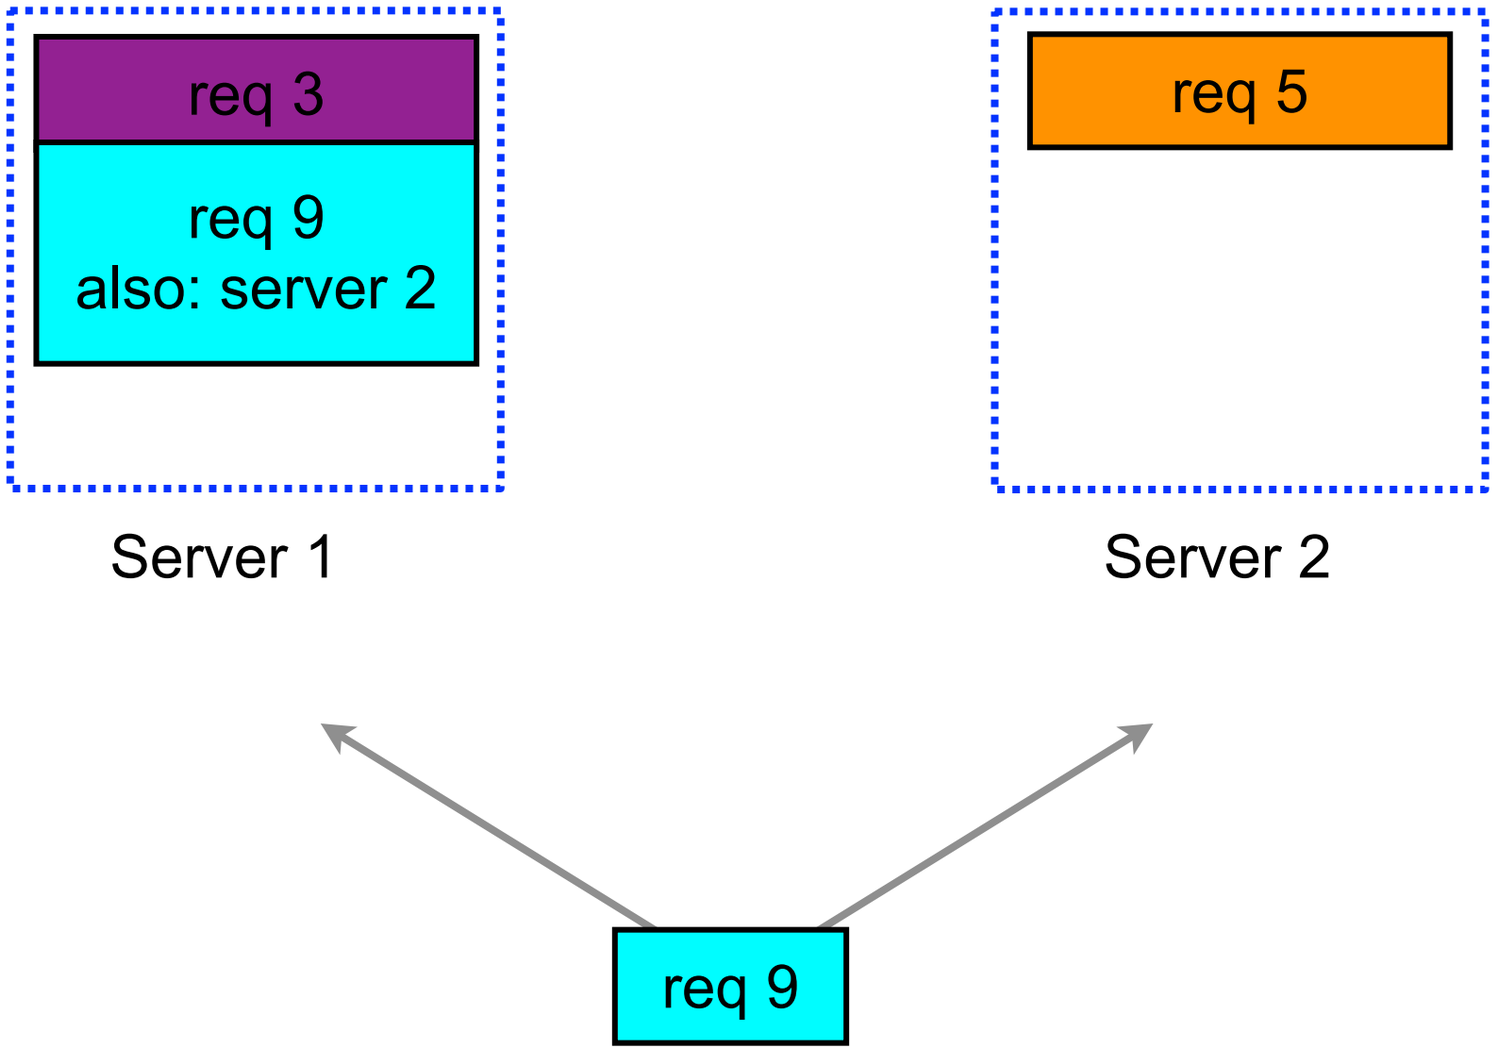

Google

Backup Requests: Bad Case

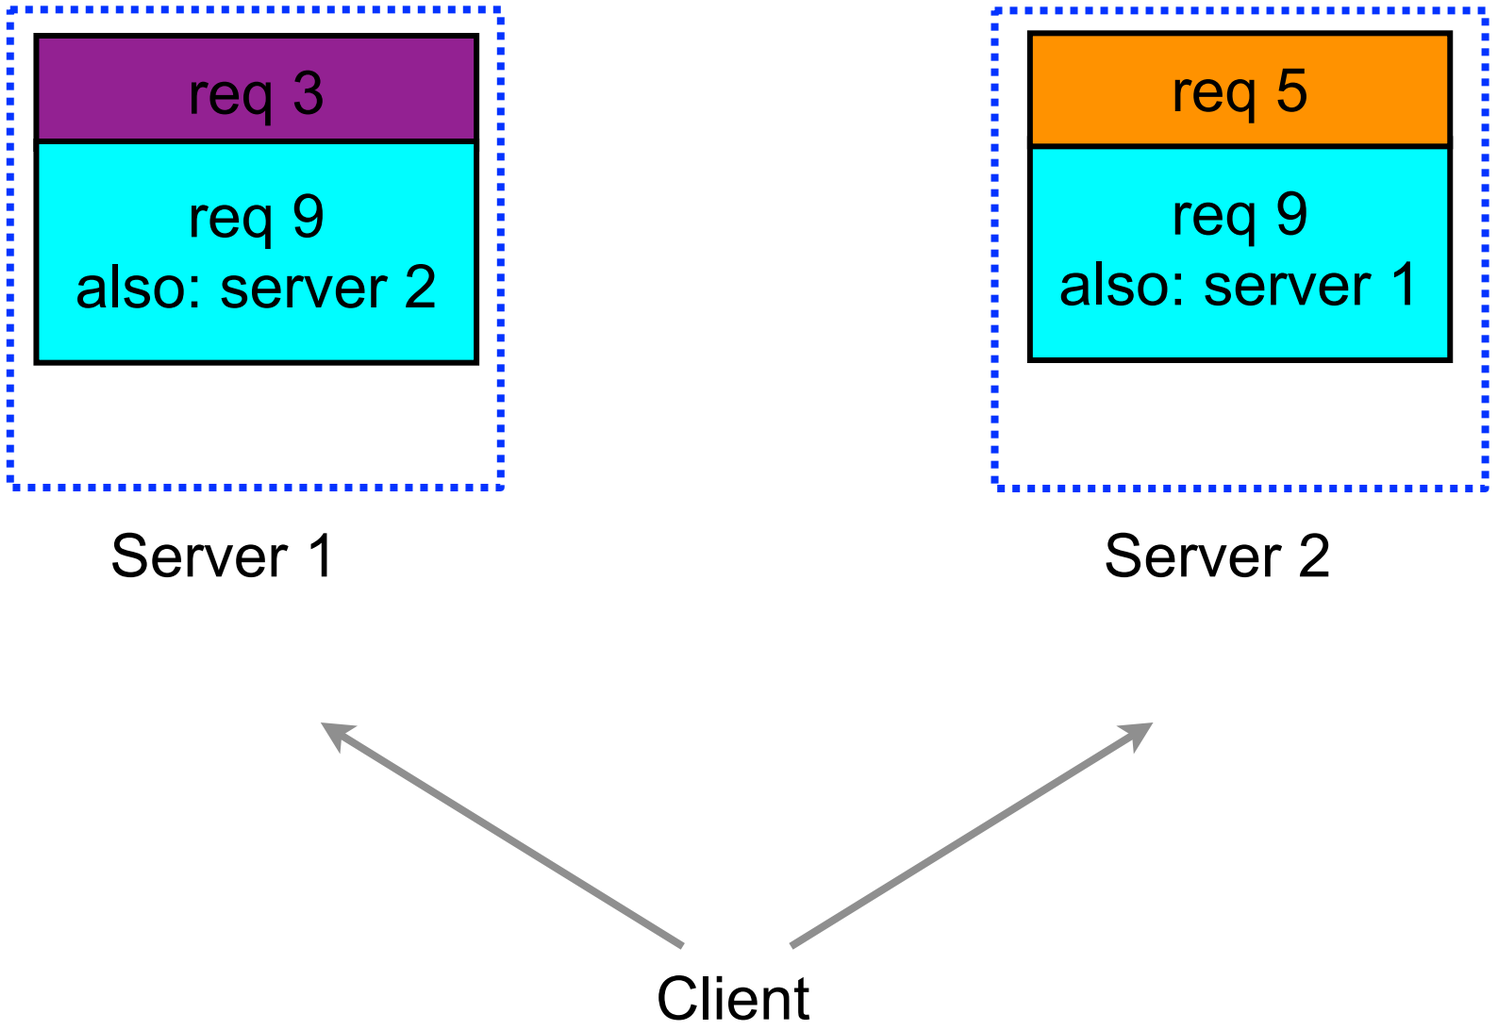

Google

Backup Requests: Bad Case

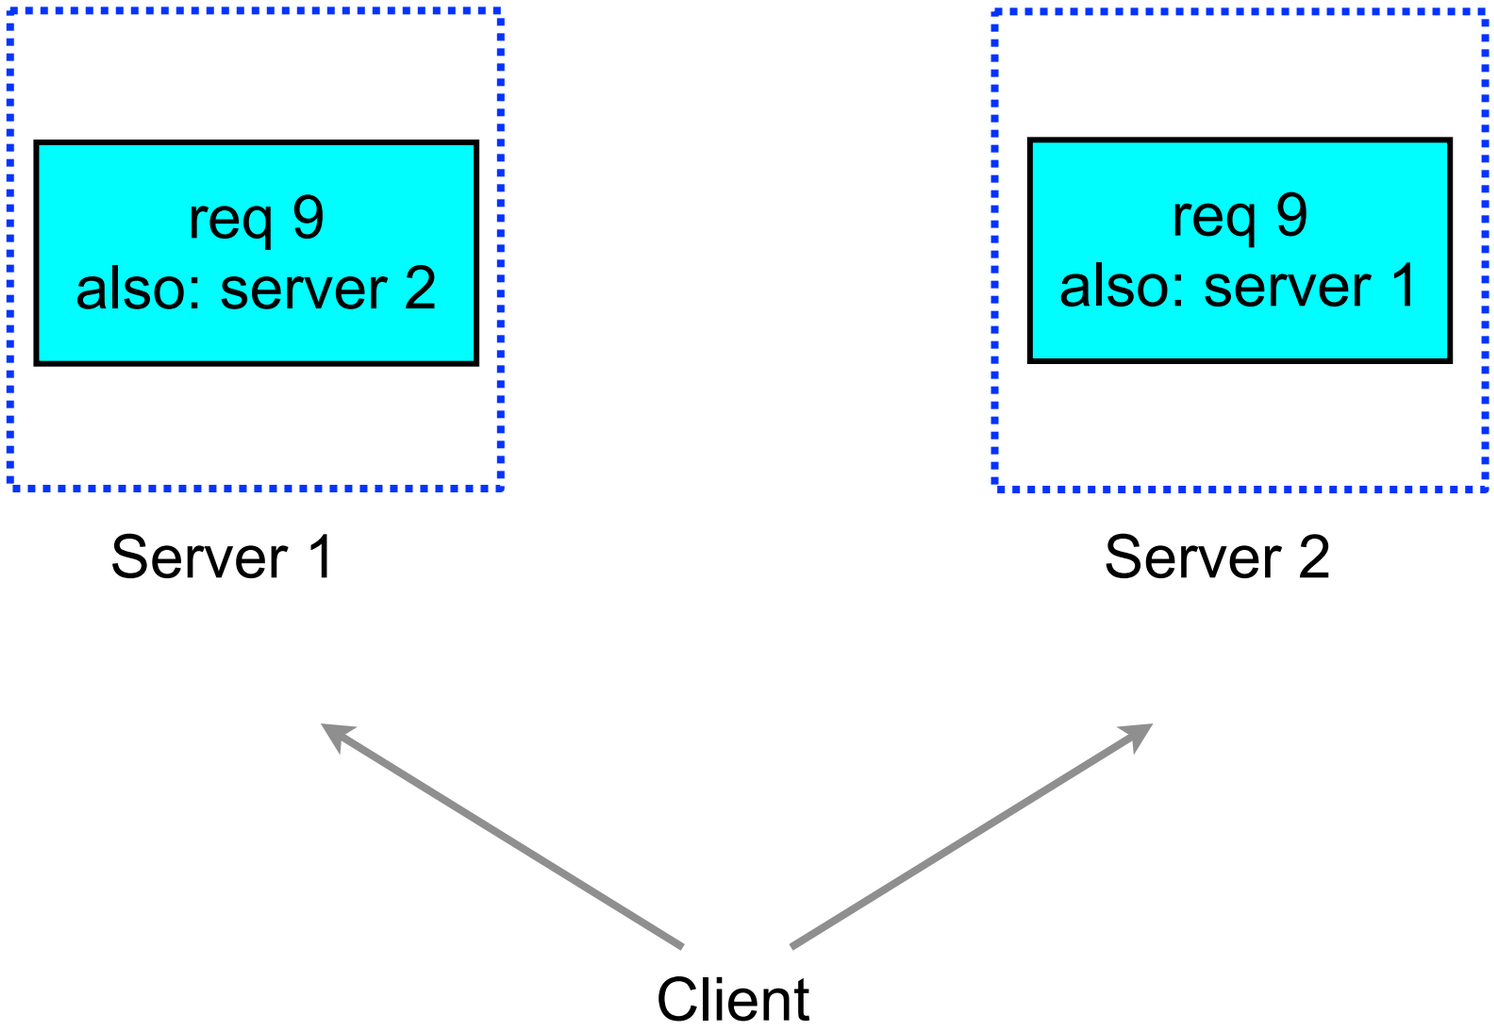

Google

Backup Requests: Bad Case

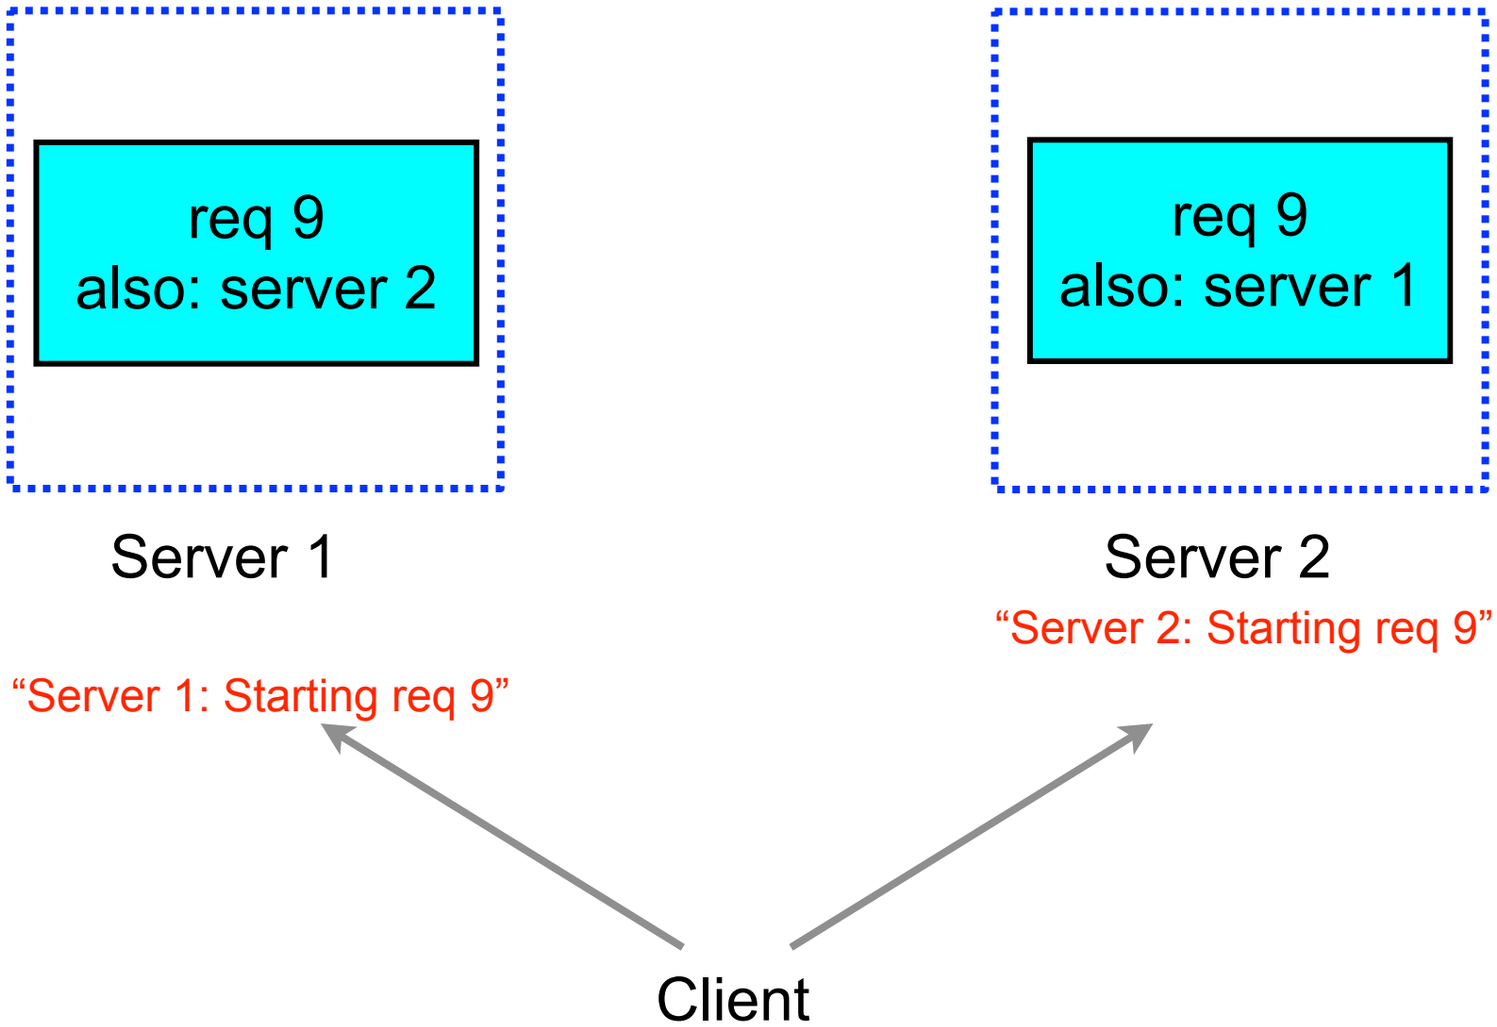

Google

Backup Requests: Bad Case

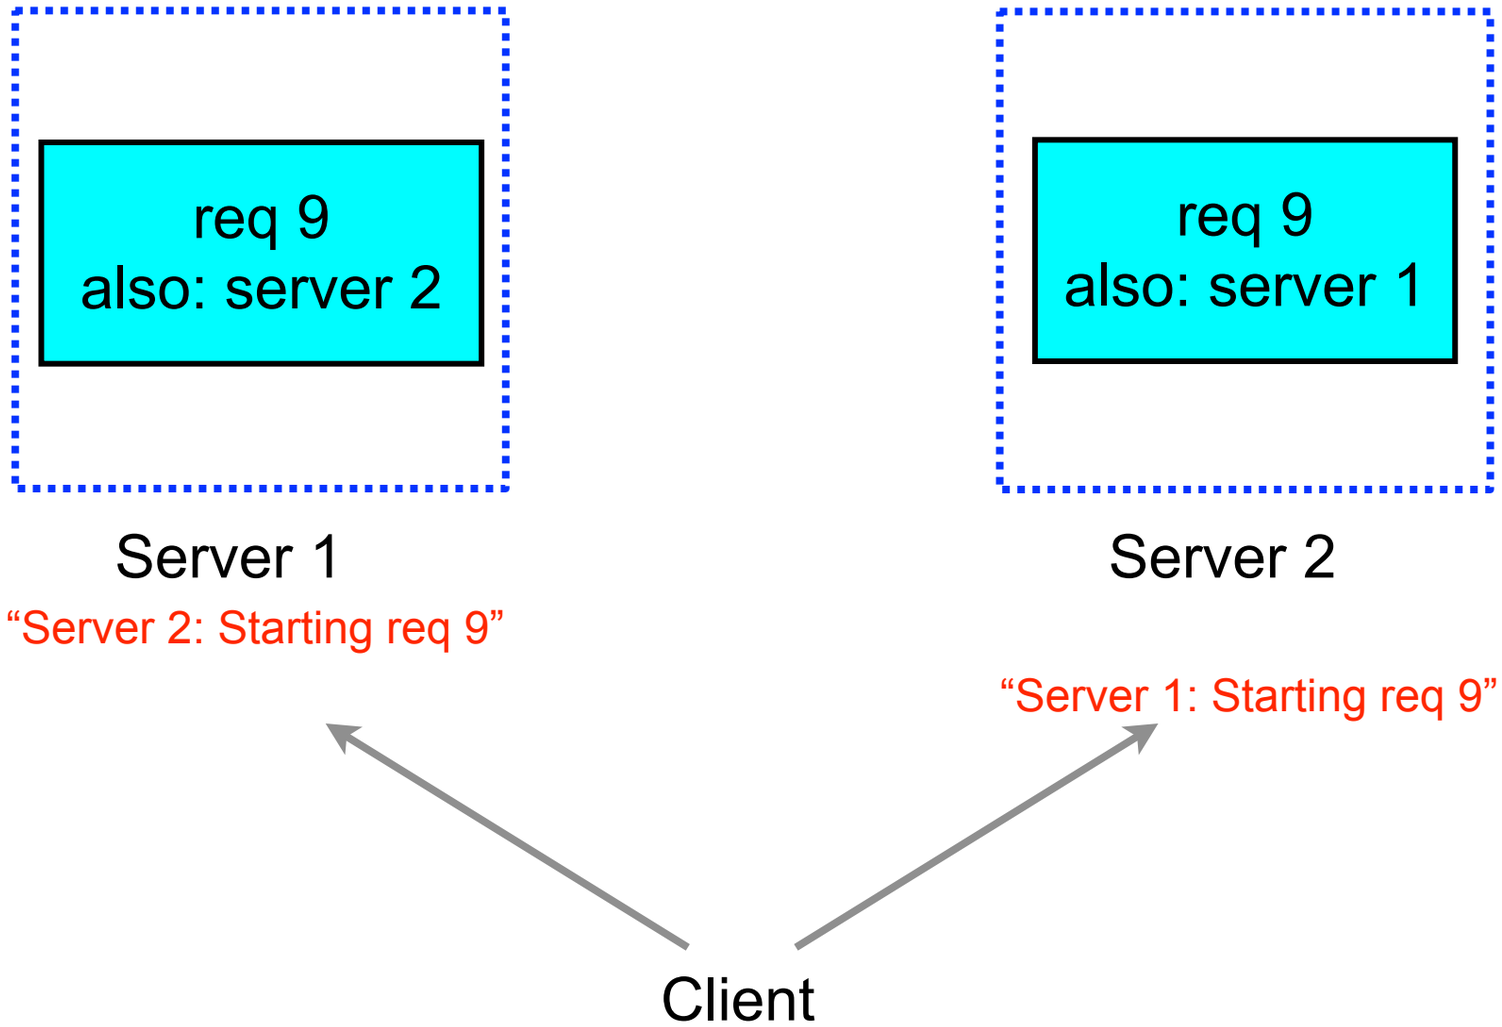

Backup Requests: Bad Case

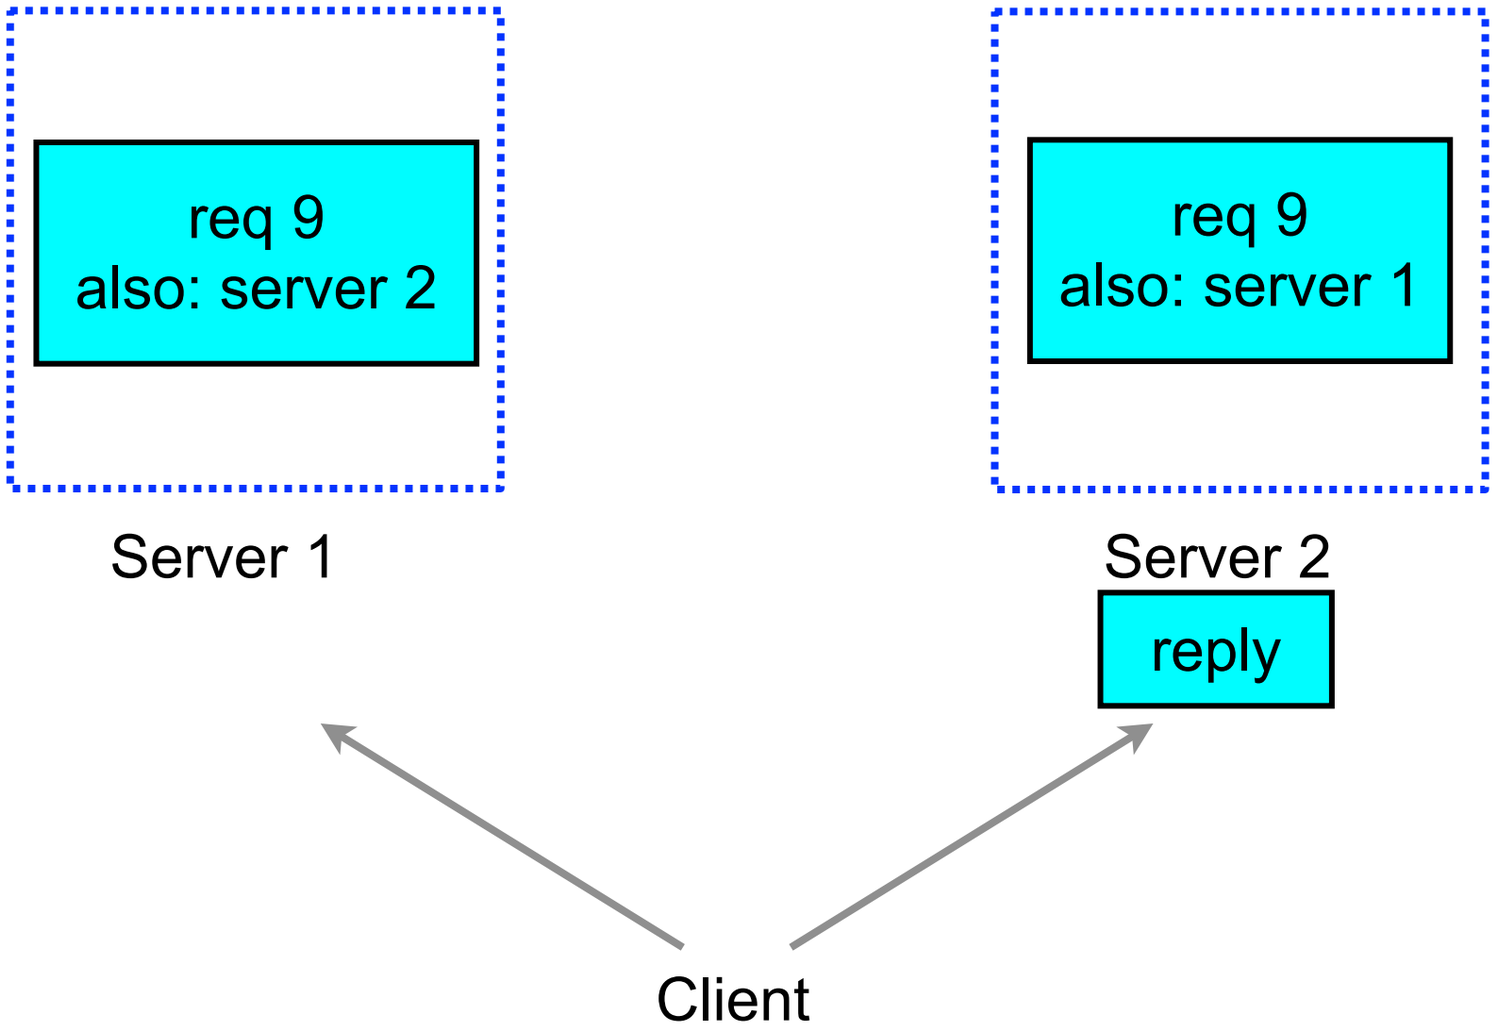

Google

Backup Requests: Bad Case

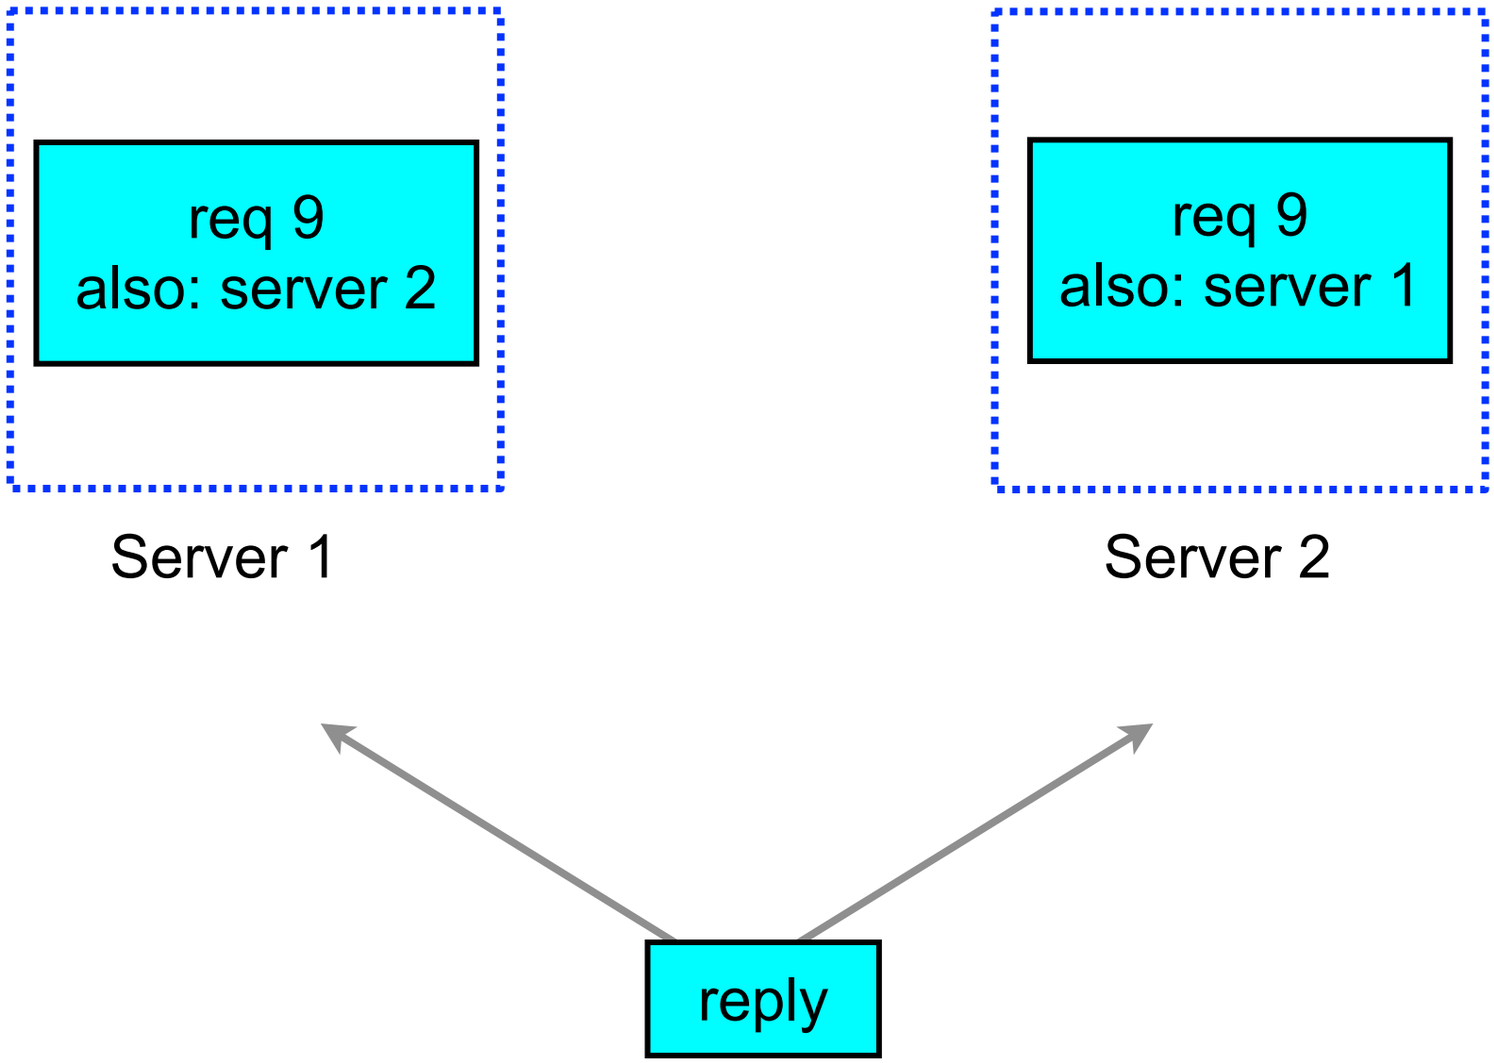

Google

Backup Requests: Bad Case

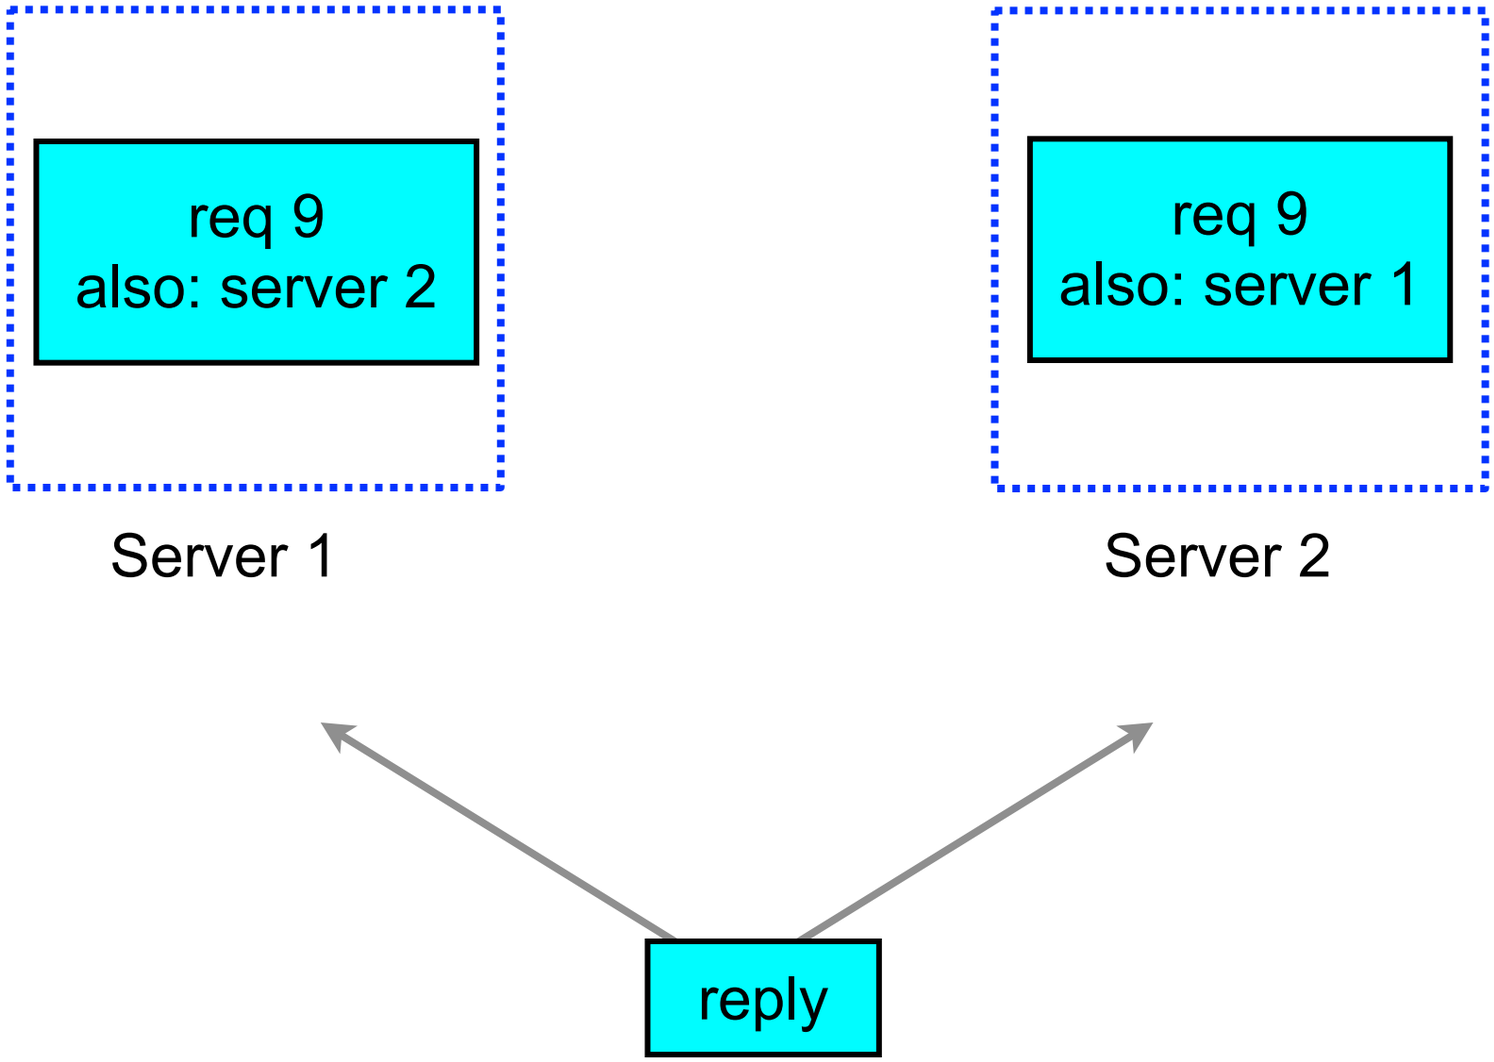

Likelihood of this bad case is reduced with lower latency networks

Google

Backup Requests w/ Cross-Server Cancellation

- Read operations in distributed file system client
  - send request to first replica
  - wait 2 ms, and send to second replica
  - servers cancel request on other replica when starting read
- Measure higher-level monitoring ops that touch disk

Google

Backup Requests w/ Cross-Server Cancellation

- Read operations in distributed file system client
  - send request to first replica
  - wait 2 ms, and send to second replica
  - servers cancel request on other replica when starting read
- Measure higher-level monitoring ops that touch disk

Cluster state	Policy	50%ile	90%ile	99%ile	99.9%ile
Mostly idle	No backups	19 ms	38 ms	67 ms	98 ms
	Backup after 2 ms	16 ms	28 ms	38 ms	51 ms

Google

Backup Requests w/ Cross-Server Cancellation

- Read operations in distributed file system client
  - send request to first replica
  - wait 2 ms, and send to second replica
  - servers cancel request on other replica when starting read
- Measure higher-level monitoring ops that touch disk

-43%

Cluster state	Policy	50%ile	90%ile	99%ile	99.9%ile
Mostly idle	No backups	19 ms	38 ms	67 ms	98 ms
	Backup after 2 ms	16 ms	28 ms	38 ms	51 ms

Google

Backup Requests w/ Cross-Server Cancellation

- Read operations in distributed file system client
  - send request to first replica
  - wait 2 ms, and send to second replica
  - servers cancel request on other replica when starting read
- Measure higher-level monitoring ops that touch disk

Cluster state	Policy	50%ile	90%ile	99%ile	99.9%ile
Mostly idle	No backups	19 ms	38 ms	67 ms	98 ms
	Backup after 2 ms	16 ms	28 ms	38 ms	51 ms
+Terasort	No backups	24 ms	56 ms	108 ms	159 ms
	Backup after 2 ms	19 ms	35 ms	67 ms	108 ms

Google

Backup Requests w/ Cross-Server Cancellation

- Read operations in distributed file system client
  - send request to first replica
  - wait 2 ms, and send to second replica
  - servers cancel request on other replica when starting read
- Measure higher-level monitoring ops that touch disk

-38%

Cluster state	Policy	50%ile	90%ile	99%ile	99.9%ile
Mostly idle	No backups	19 ms	38 ms	67 ms	98 ms
	Backup after 2 ms	16 ms	28 ms	38 ms	51 ms
+Terasort	No backups	24 ms	56 ms	108 ms	159 ms
	Backup after 2 ms	19 ms	35 ms	67 ms	108 ms

Google

Backup Requests w/ Cross-Server Cancellation

- Read operations in distributed file system client
  - send request to first replica
  - wait 2 ms, and send to second replica
  - servers cancel request on other replica when starting read
- Measure higher-level monitoring ops that touch disk

Cluster state	Policy	50%ile	90%ile	99%ile	99.9%ile
Mostly idle	No backups	19 ms	38 ms	67 ms	98 ms
	Backup after 2 ms	16 ms	28 ms	38 ms	51 ms
+Terasort	No backups	24 ms	56 ms	108 ms	159 ms
	Backup after 2 ms	19 ms	35 ms	67 ms	108 ms

Backups cause about ~1% extra disk reads

Google

Backup Requests w/ Cross-Server Cancellation

- Read operations in distributed file system client
  - send request to first replica
  - wait 2 ms, and send to second replica
  - servers cancel request on other replica when starting read
- Measure higher-level monitoring ops that touch disk

Cluster state	Policy	50%ile	90%ile	99%ile	99.9%ile
Mostly idle	No backups	19 ms	38 ms	67 ms	98 ms
	Backup after 2 ms	16 ms	28 ms	38 ms	51 ms
+Terasort	No backups	24 ms	56 ms	108 ms	159 ms
	Backup after 2 ms	19 ms	35 ms	67 ms	108 ms

Google

Backup Requests w/ Cross-Server Cancellation

- Read operations in distributed file system client
  - send request to first replica
  - wait 2 ms, and send to second replica
  - servers cancel request on other replica when starting read
- Measure higher-level monitoring ops that touch disk

Cluster state	Policy	50%ile	90%ile	99%ile	99.9%ile
Mostly idle	No backups	19 ms	38 ms	67 ms	98 ms
	Backup after 2 ms	16 ms	28 ms	38 ms	51 ms
+Terasort	No backups	24 ms	56 ms	108 ms	159 ms
	Backup after 2 ms	19 ms	35 ms	67 ms	108 ms

Backups w/big sort job gives same read latencies as no backups w/ idle cluster!

Google

Cluster-Level Services

- Our earliest systems made things easier within a cluster:
  - GFS/Colossus: reliable cluster-level file system
  - MapReduce: reliable large-scale computations
  - Cluster scheduling system: abstracted individual machines
  - BigTable: automatic scaling of higher-level structured storage

Cluster-Level Services

- Our earliest systems made things easier within a cluster:
  - GFS/Colossus: reliable cluster-level file system
  - MapReduce: reliable large-scale computations
  - Cluster scheduling system: abstracted individual machines
  - BigTable: automatic scaling of higher-level structured storage

- Solve many problems, but leave many cross-cluster issues to human-level operators
  - different copies of same dataset have different names
  - moving or deploying new service replicas is labor intensive

Google

Dependencies

- Google Docs depends on 50+ services
- Want to perform high-level operations:
  - run it in half a dozen clusters
  - release a new version
  - fix it on the fly in an emergency
  - move some (possibly shared) sub-services to another cluster

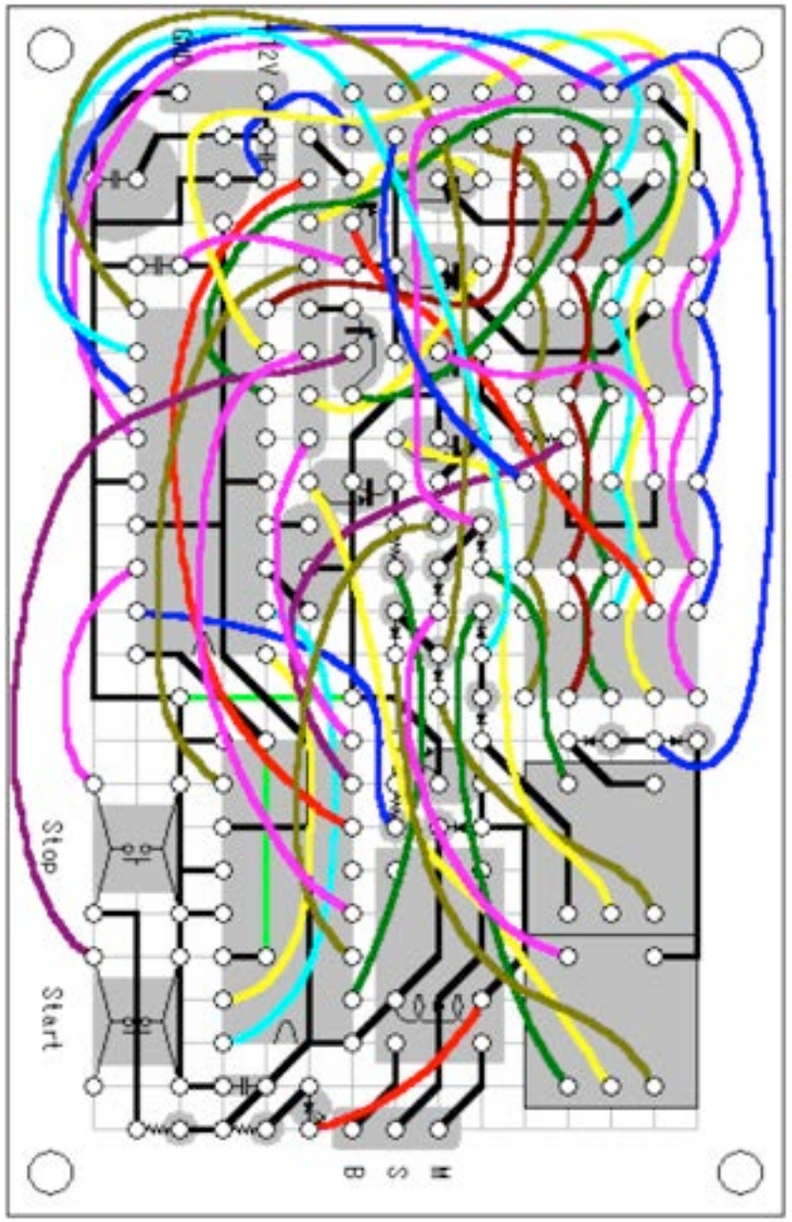

Google

Dependencies

- Google Docs depends on 50+ services
- Want to perform high-level operations:
  - run it in half a dozen clusters
  - release a new version
  - fix it on the fly in an emergency
  - move some (possibly shared) sub-services to another cluster

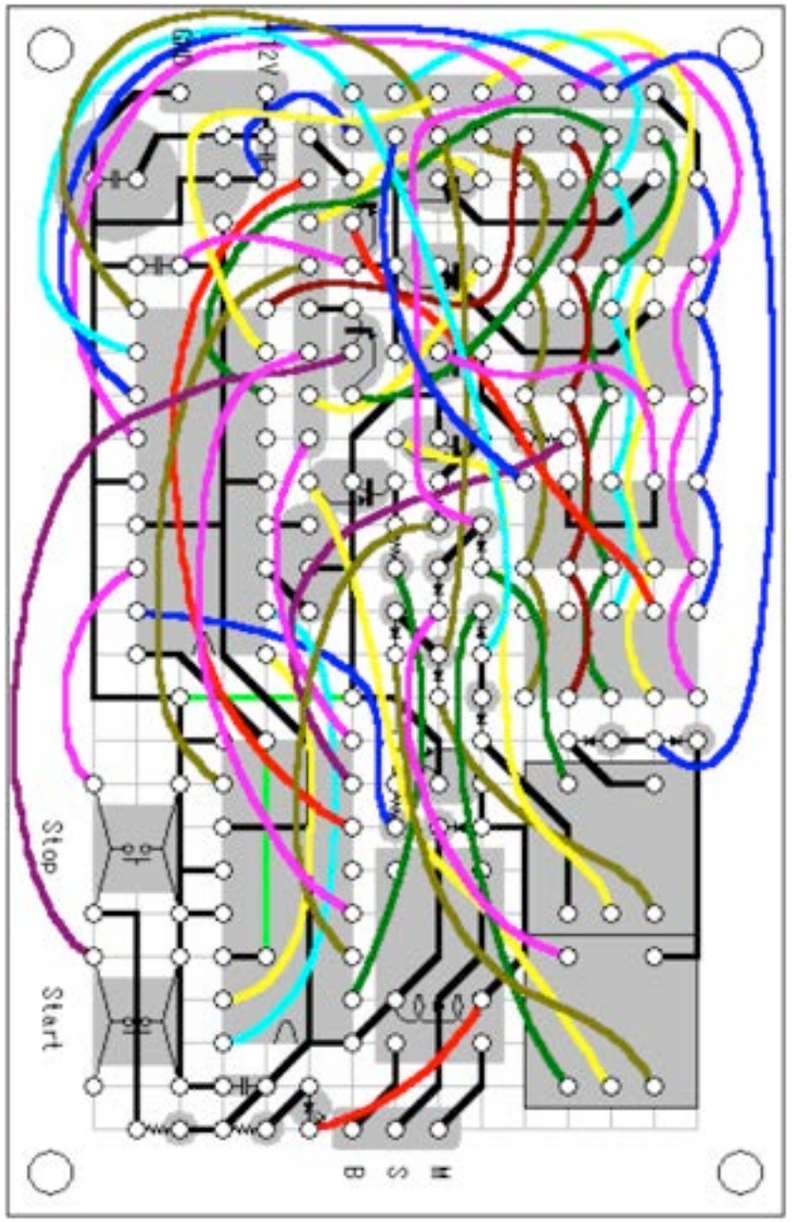

Important that scheduling service has worldwide view

Google

World-Wide Systems

- Future scale: ~10^6 to 10^7 machines, spread at 100s to 1000s of locations around the world, ~10^9 client machines

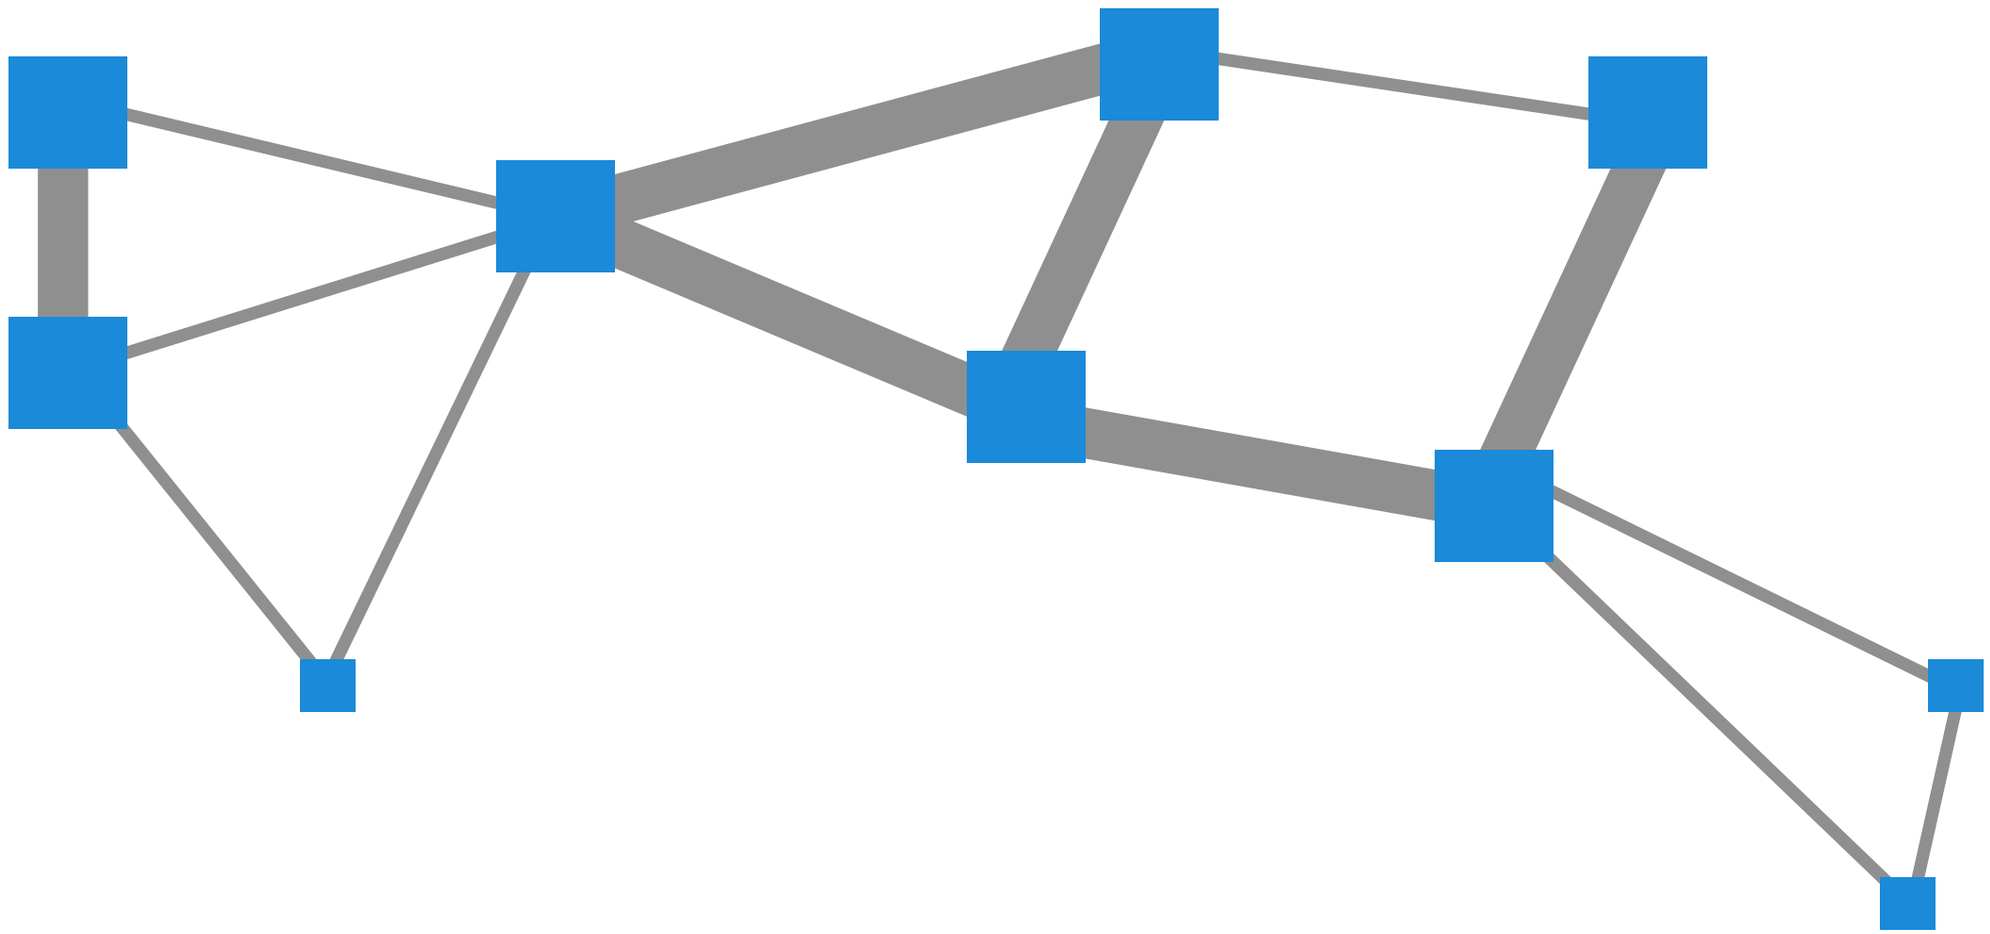

  – zones of semi-autonomous control
  – consistency after disconnected operation
  – power adaptivity

Google

Adaptivity in World-Wide Systems

- Challenge: automatic, dynamic world-wide placement of data & computation to minimize latency and/or cost, given constraints on:
  - bandwidth
  - packet loss
  - power
  - resource usage
  - failure modes
  - ...
- Users specify high-level desires:

  "Store this data on at least 2 disks in EU, 2 in U.S. & 1 in Asia"

  "99%ile latency for accessing this data should be <50ms"

Google

Higher Level Systems

- Systems that provide high level of abstraction that "just works" are incredibly valuable:
  - GFS, MapReduce, BigTable, world-wide systems, transparent variability reduction techniques, etc.

- Can we build high-level systems that just work in other domains like machine learning?

Scaling Deep Learning

- Much of Google is working on approximating AI. AI is hard
  - Many people at Google spend countless person-years hand-engineering complex features to feed as input to machine learning algorithms

- Is there a better way?

- Deep Learning: Use very large scale brain simulations
  - improve many Google applications
  - make significant advances towards perceptual AI

Deep Learning

- Algorithmic approach
  - automatically learn high-level representations from raw data
  - can learn from both labeled and unlabeled data

- Recent academic deep learning results improve on state-of-the-art in many areas:
  - images, video, speech, NLP, ...
  - ... using modest model sizes (<= ~50M parameters)

- We want to scale this approach up to much bigger models
  - currently: ~2B parameters, want ~10B-100B parameters
  - general approach: parallelize at many levels

Deep Networks

Input Image
(or video)

Deep Networks

Input Image
(or video)

Deep Networks

Deep Networks

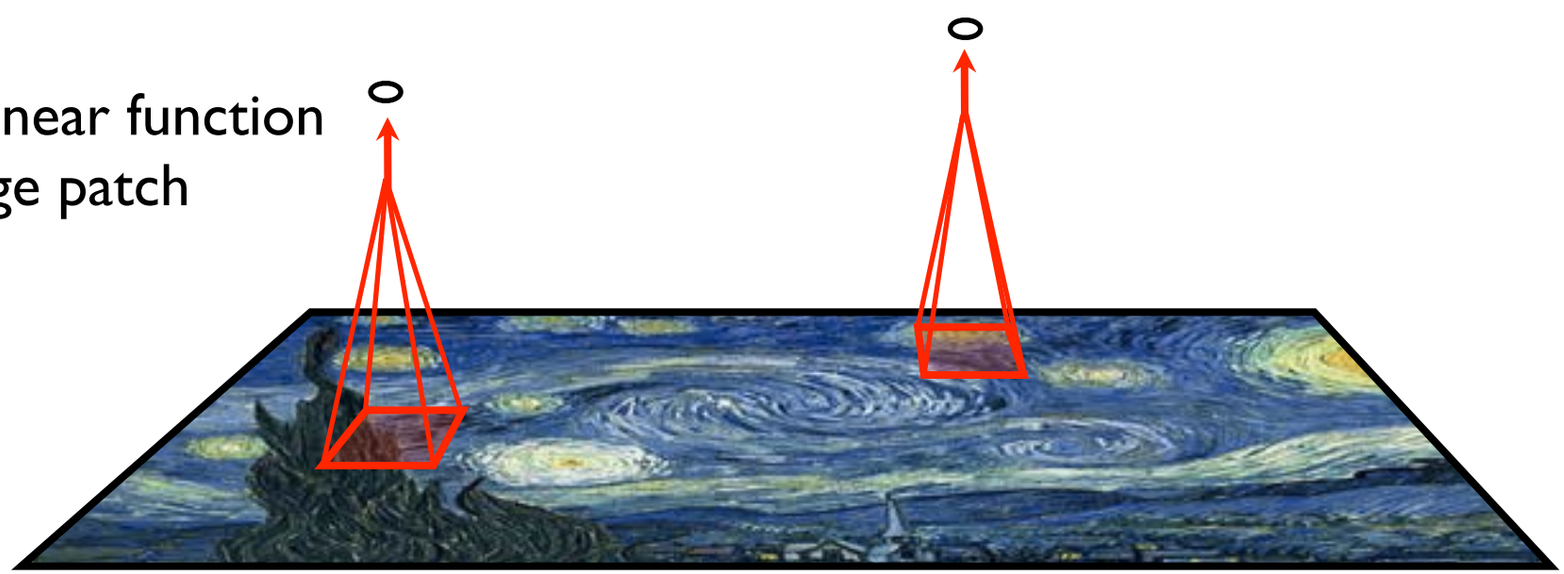

Deep Networks

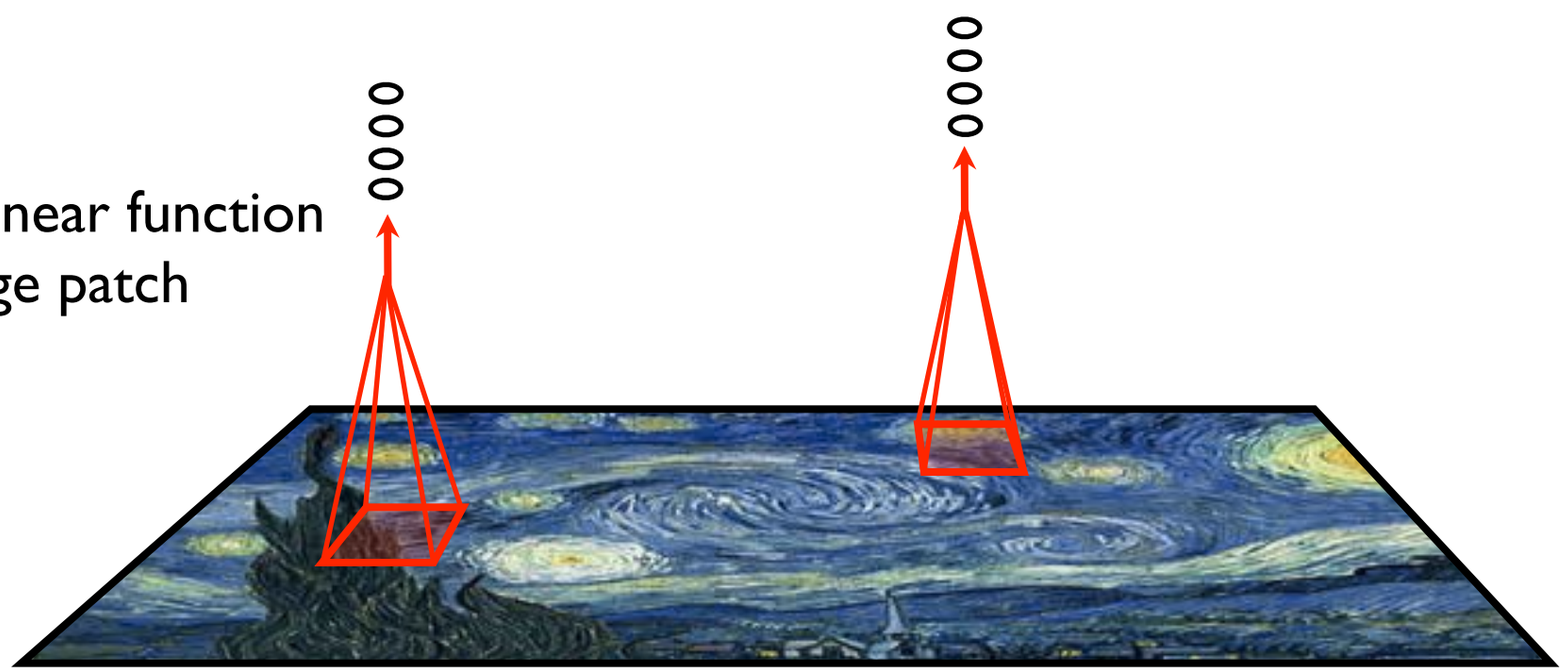

Deep Networks

Many responses at a single location.
In many models these are independent,
but some allow strong nonlinear interactions

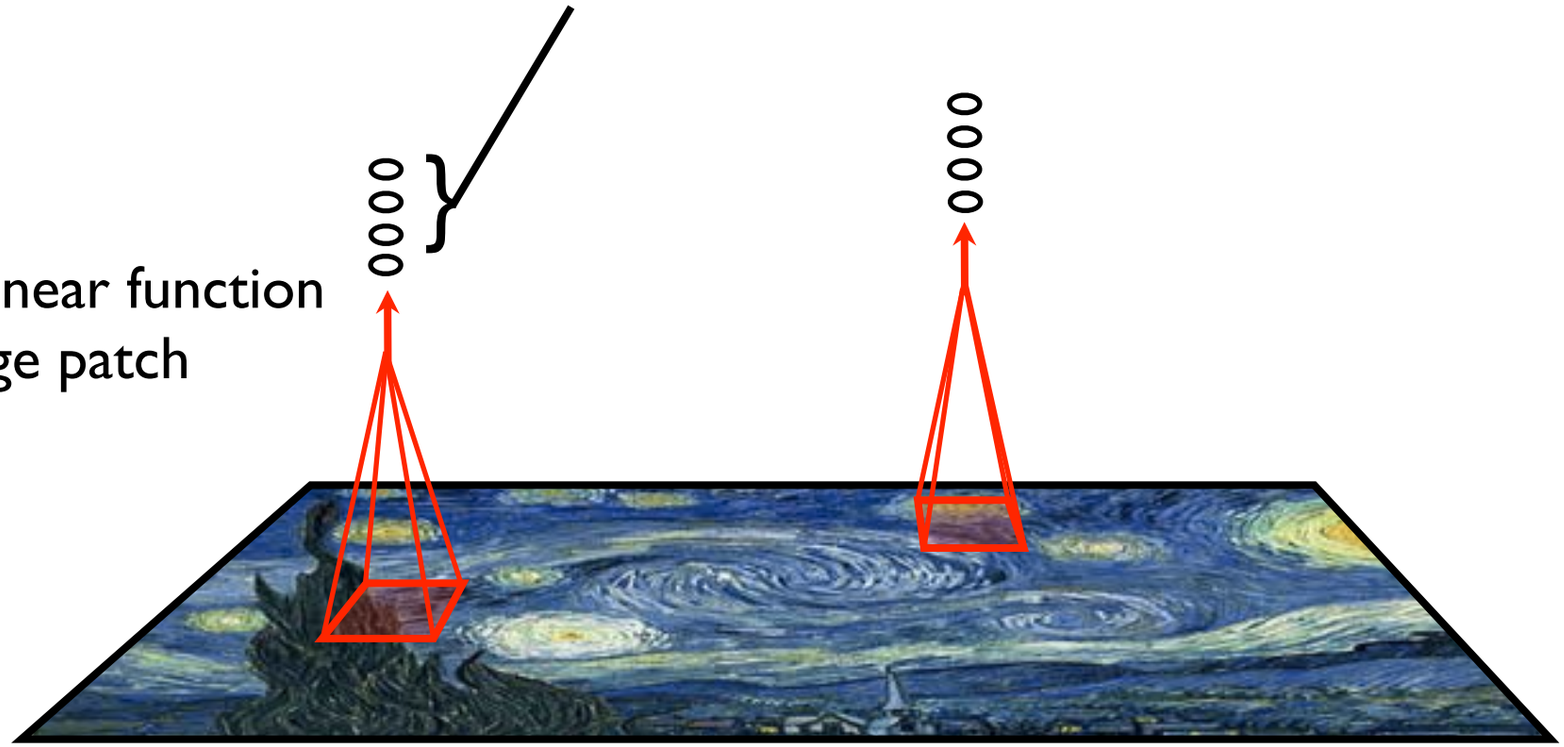

Some scalar, nonlinear function
of local image patch

Input Image
(or video)

Deep Networks

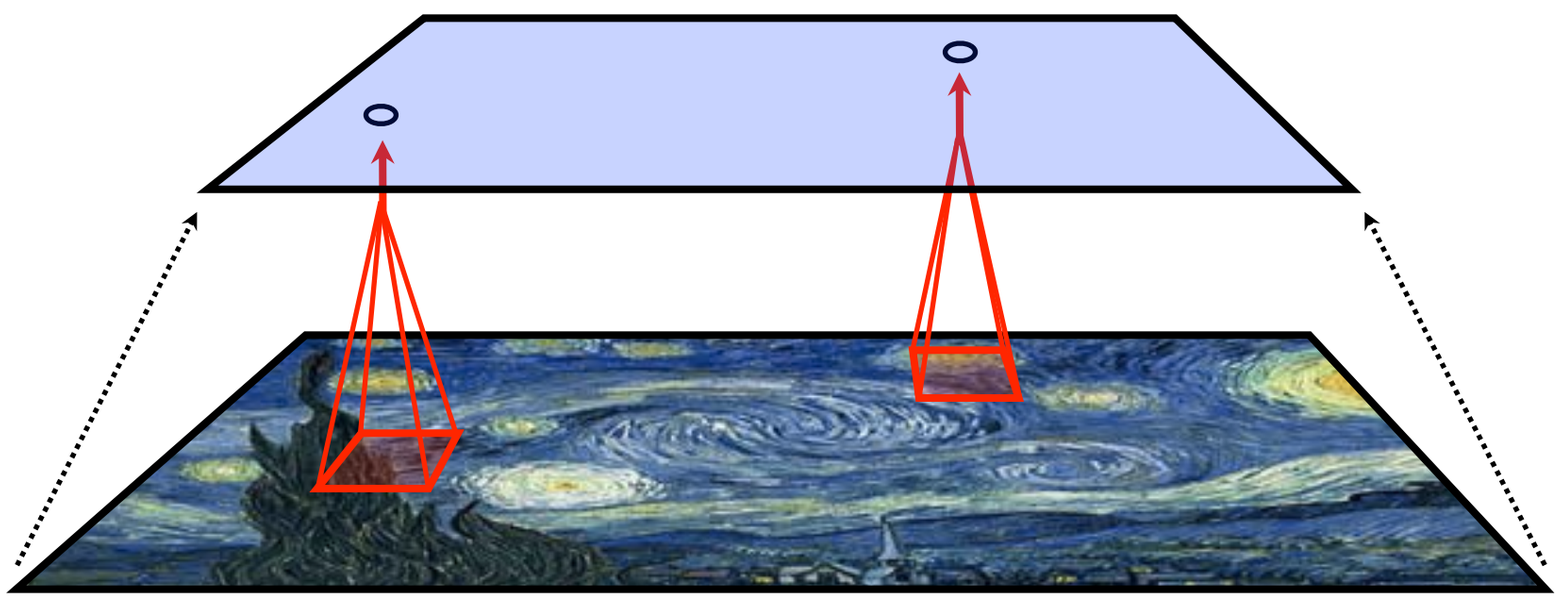

Input Image
(or video)

Deep Networks

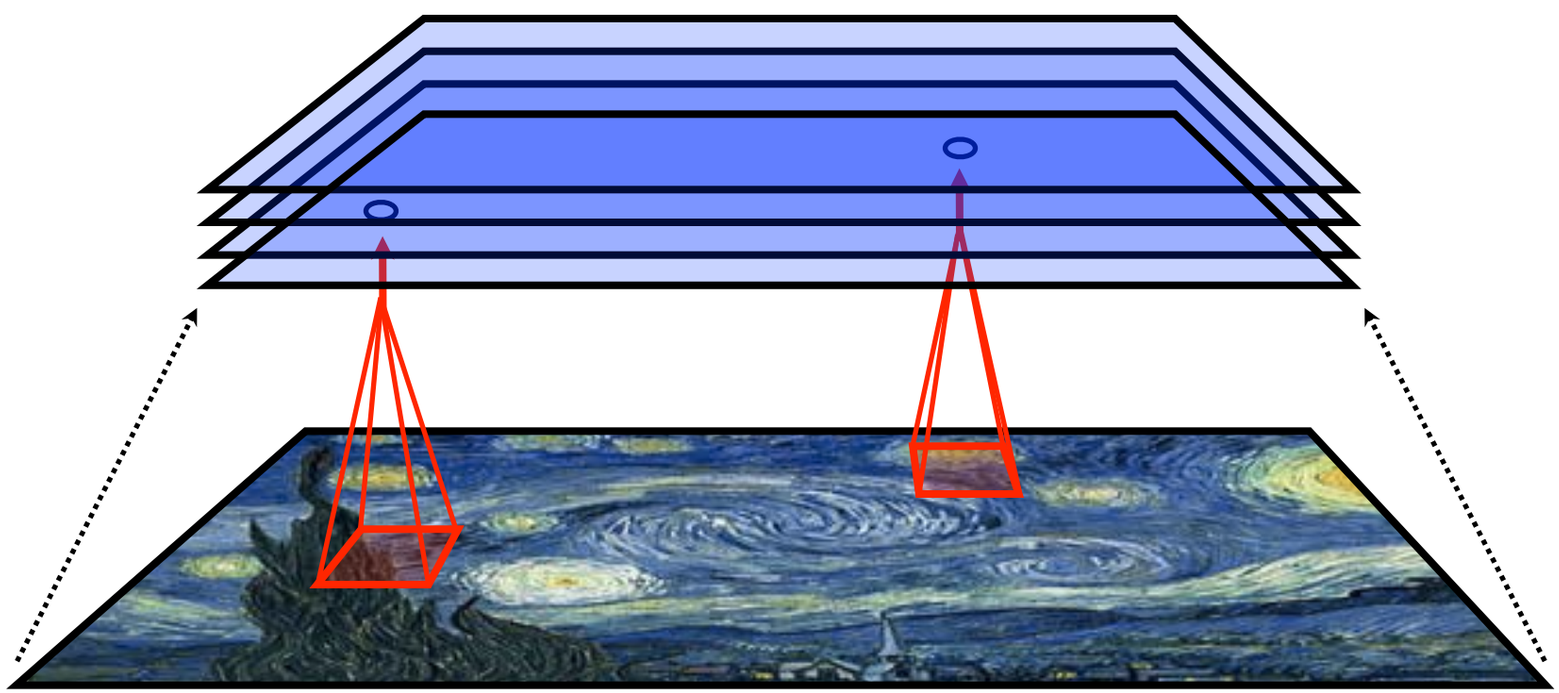

Input Image
(or video)

Deep Networks

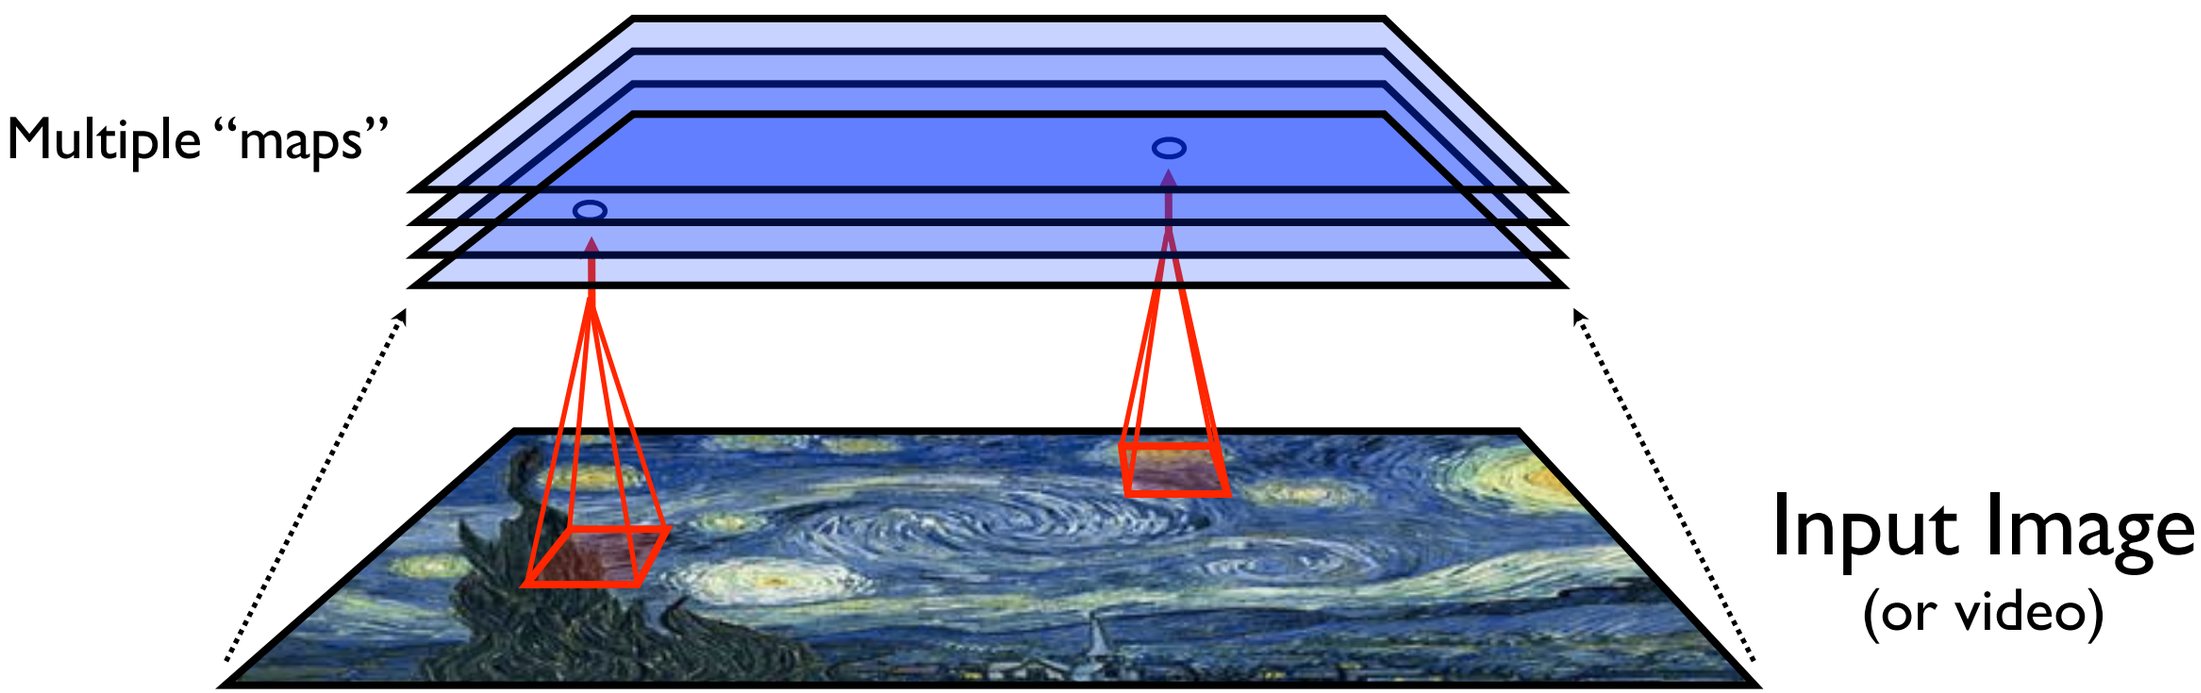

Deep Networks

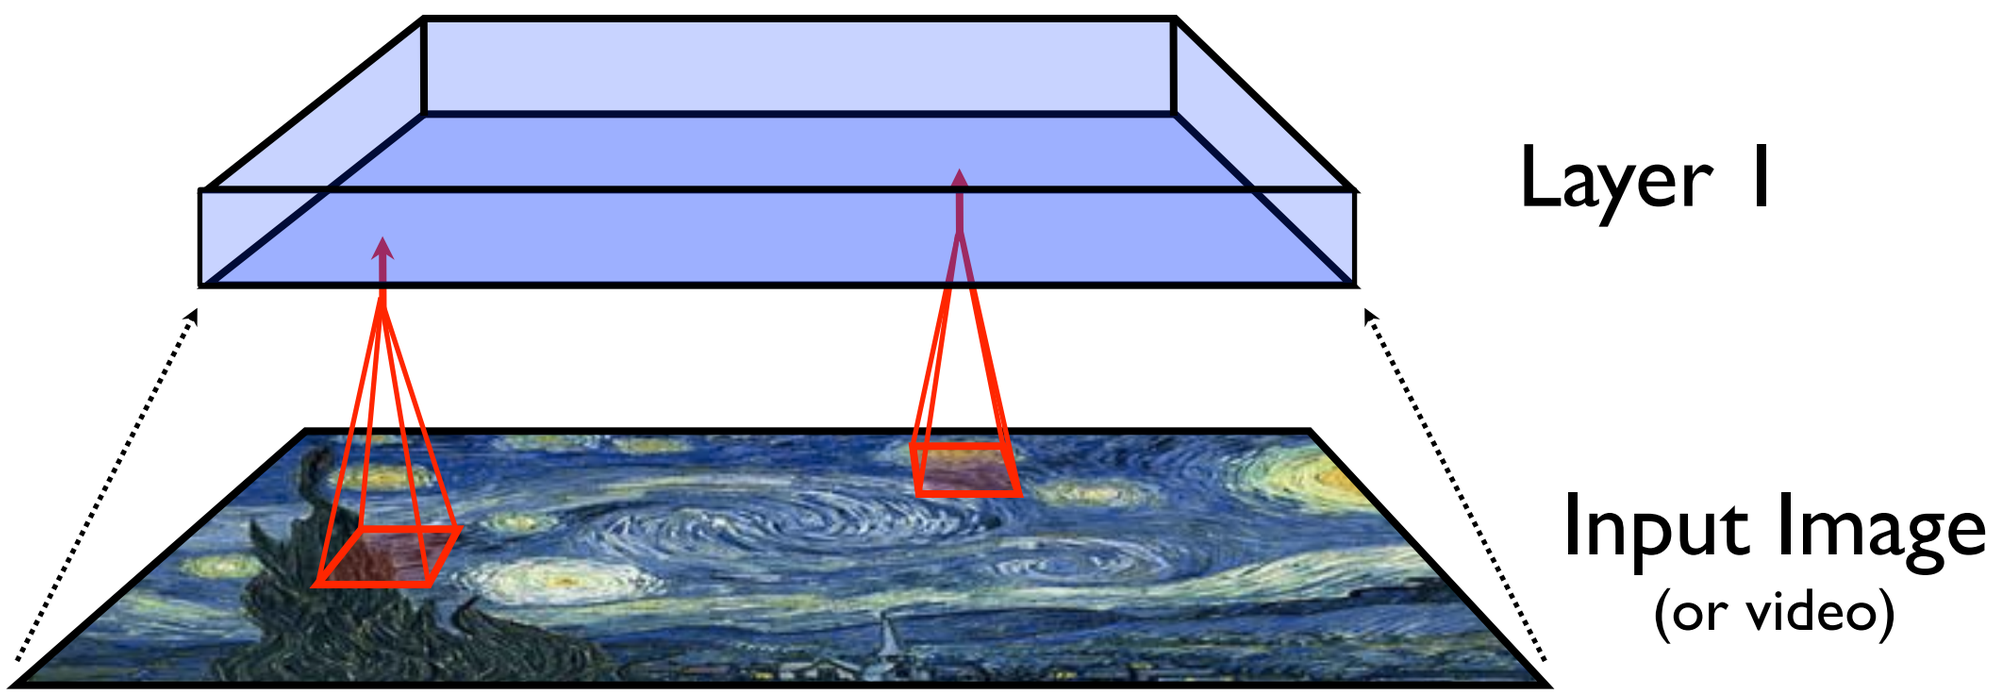

Unsupervised Training

Core idea: try to reconstruct input from just the learned representation

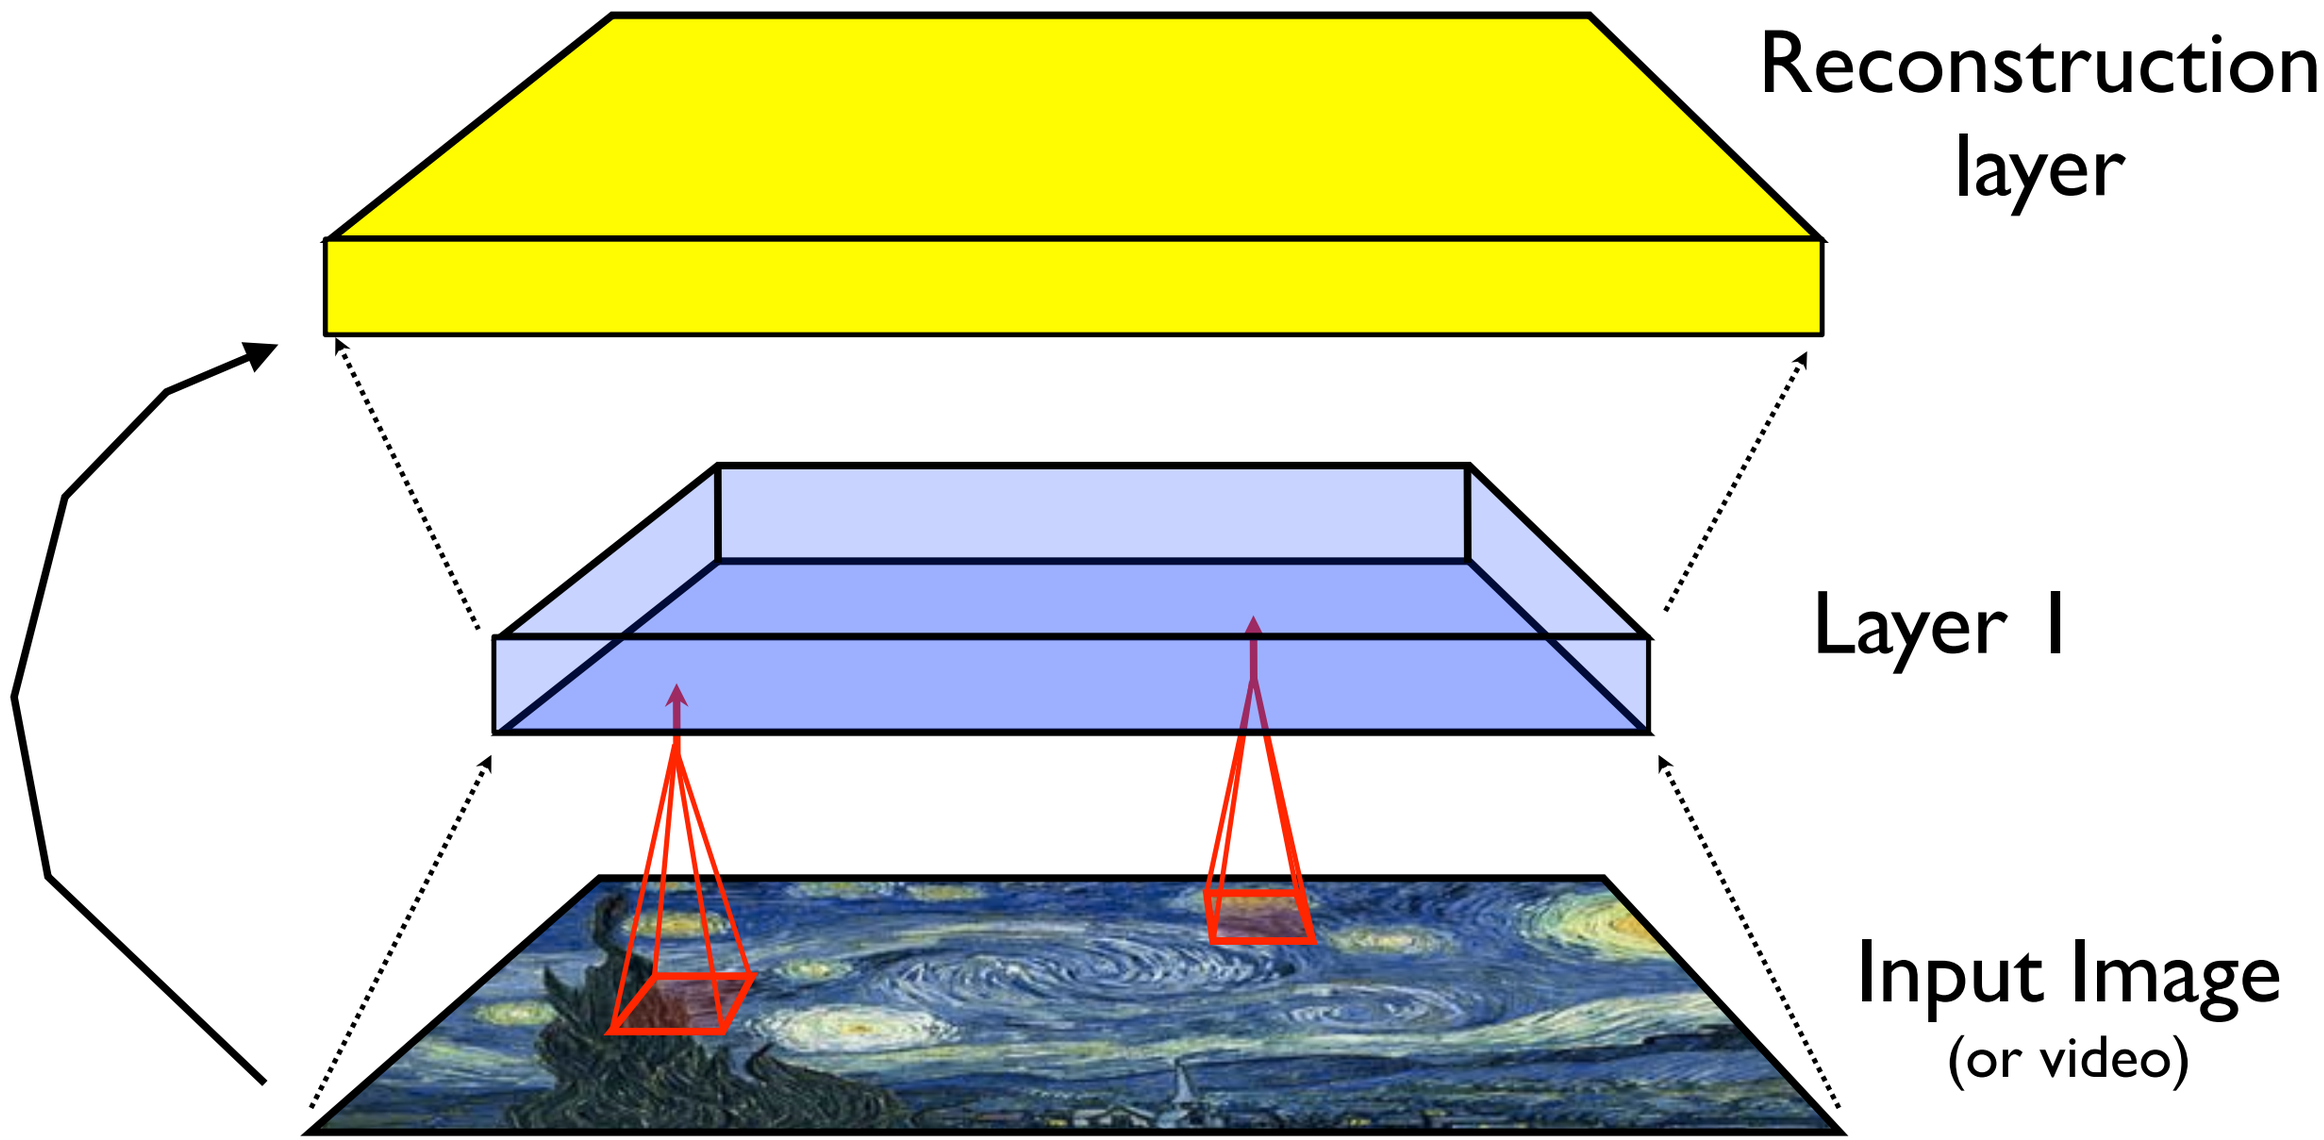

Due to Geoff Hinton, Yoshua Bengio, Andrew Ng, and others

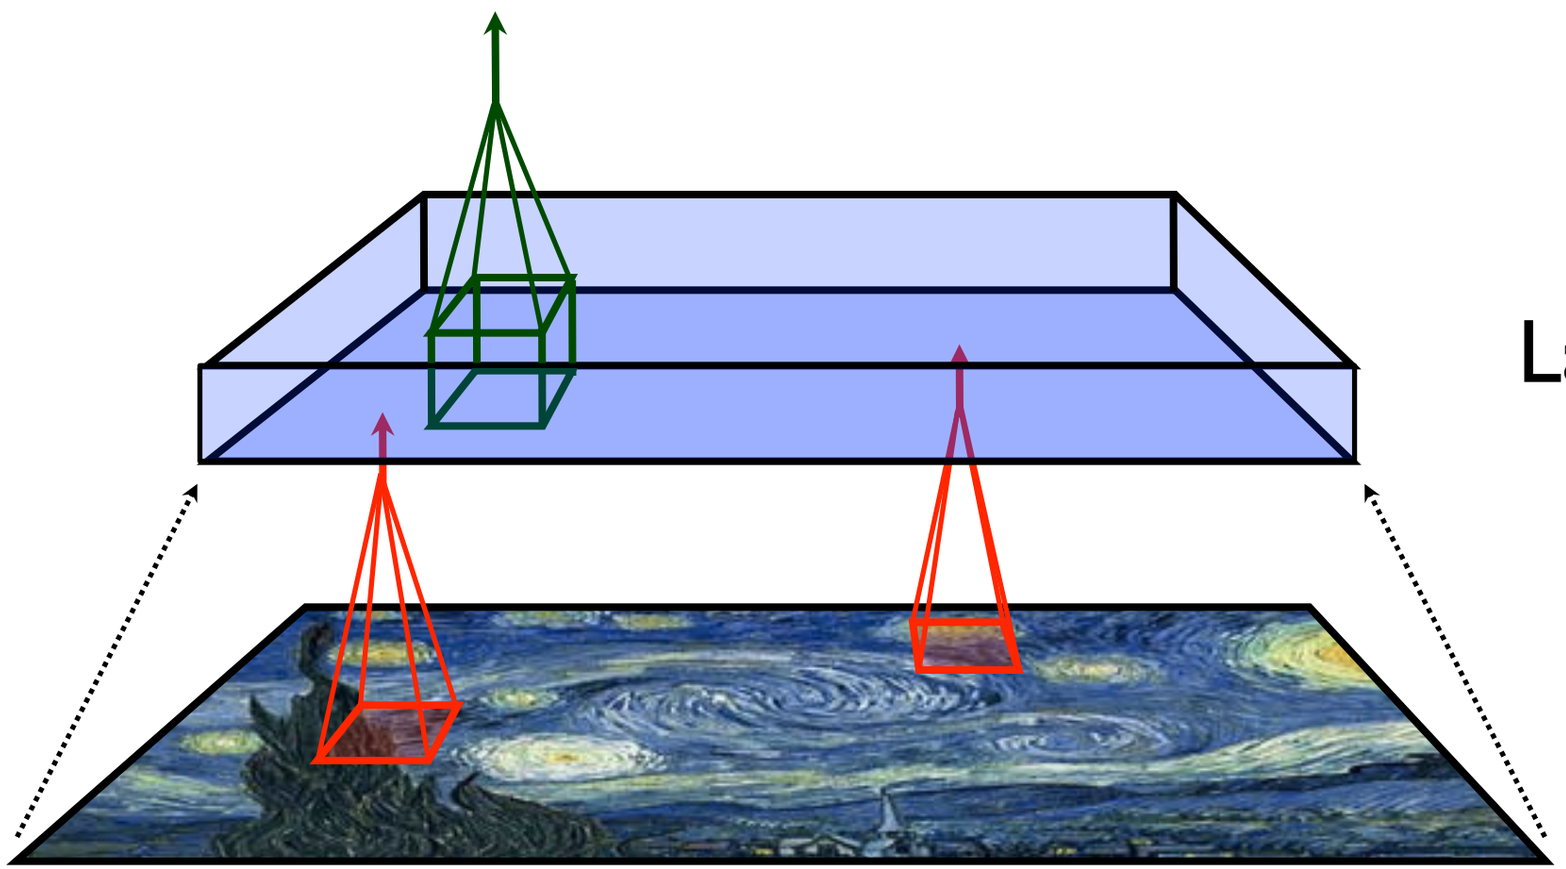
Layer 1
Input Image
(or video)

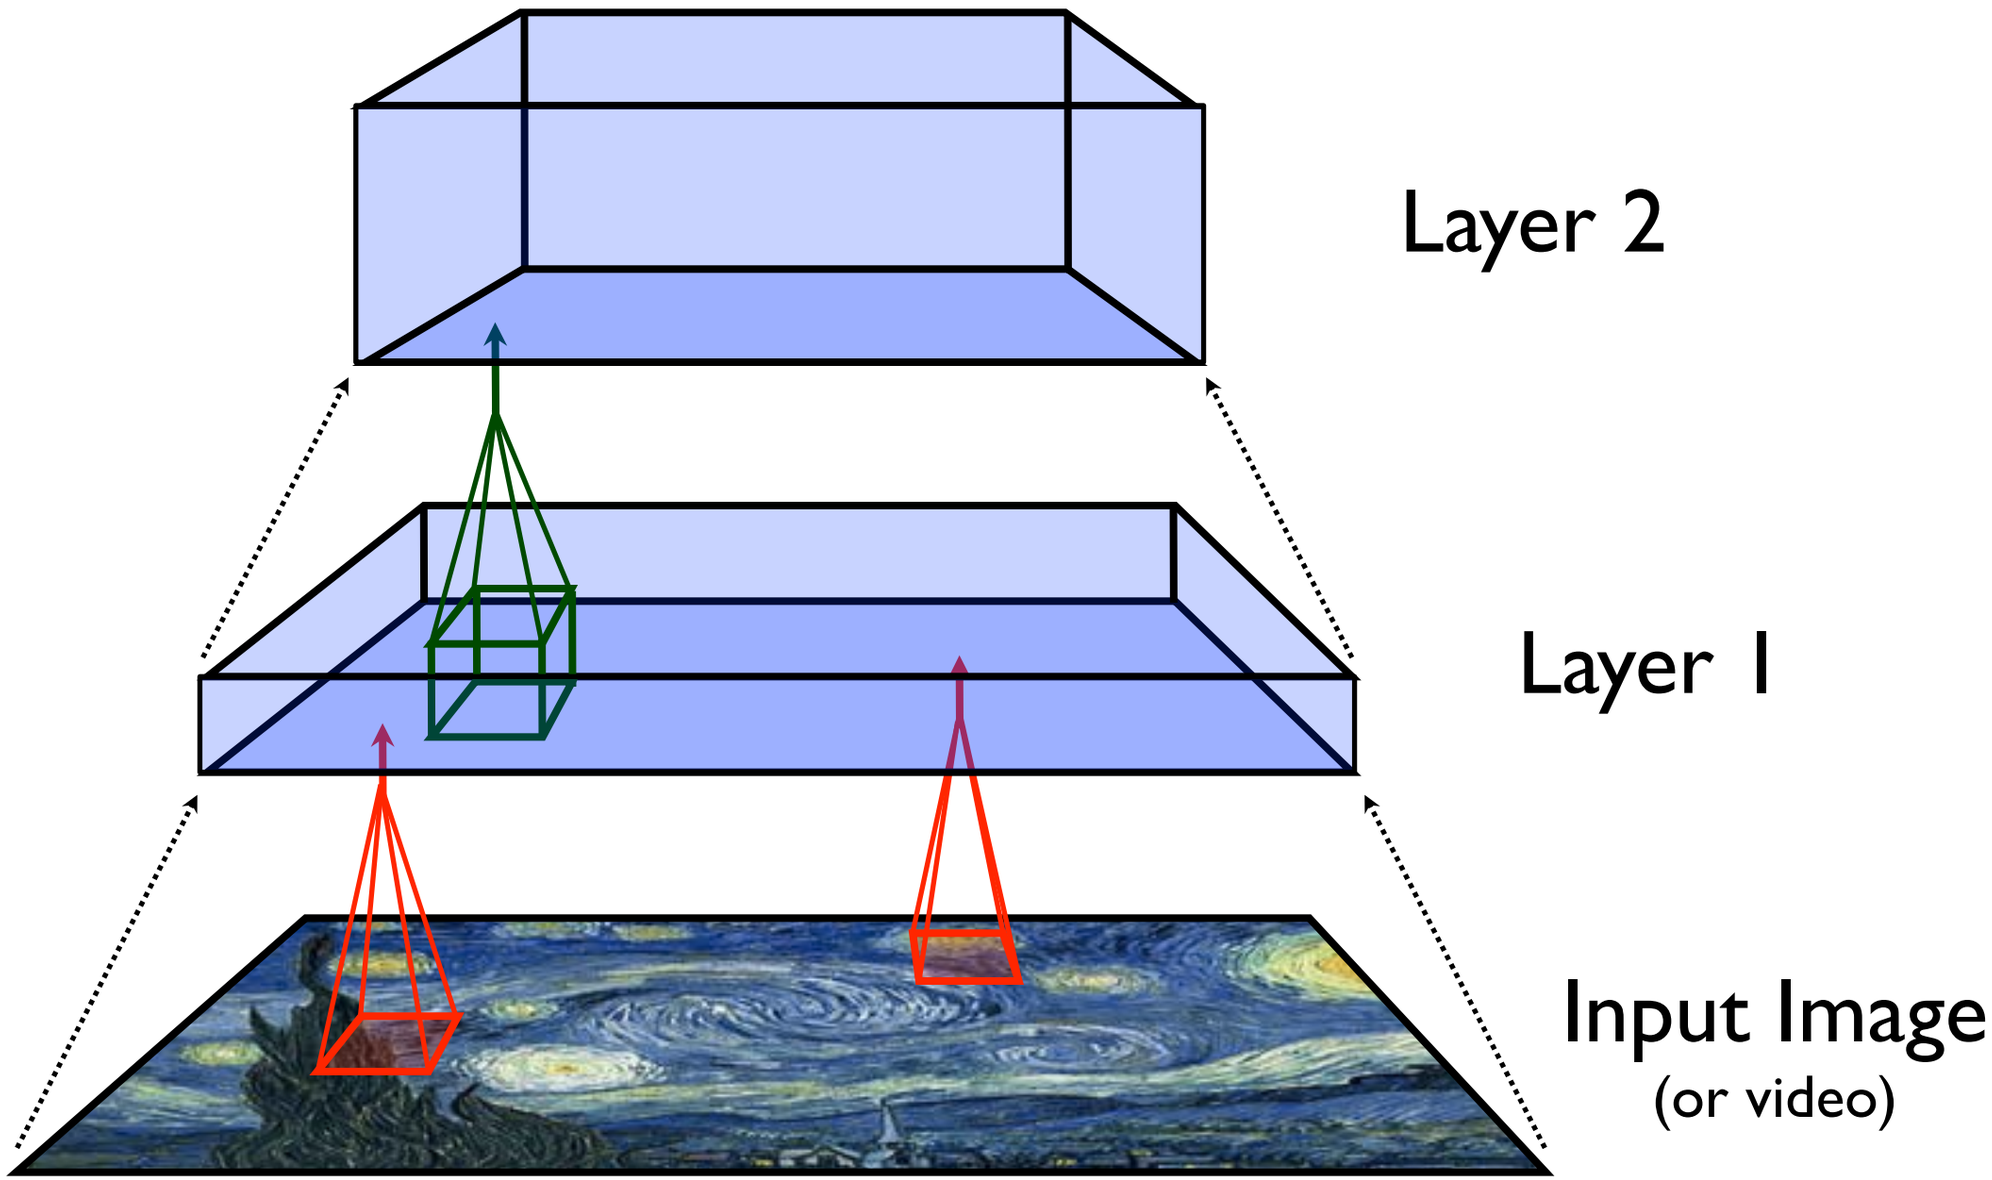
Layer 2
Layer 1
Input Image
(or video)

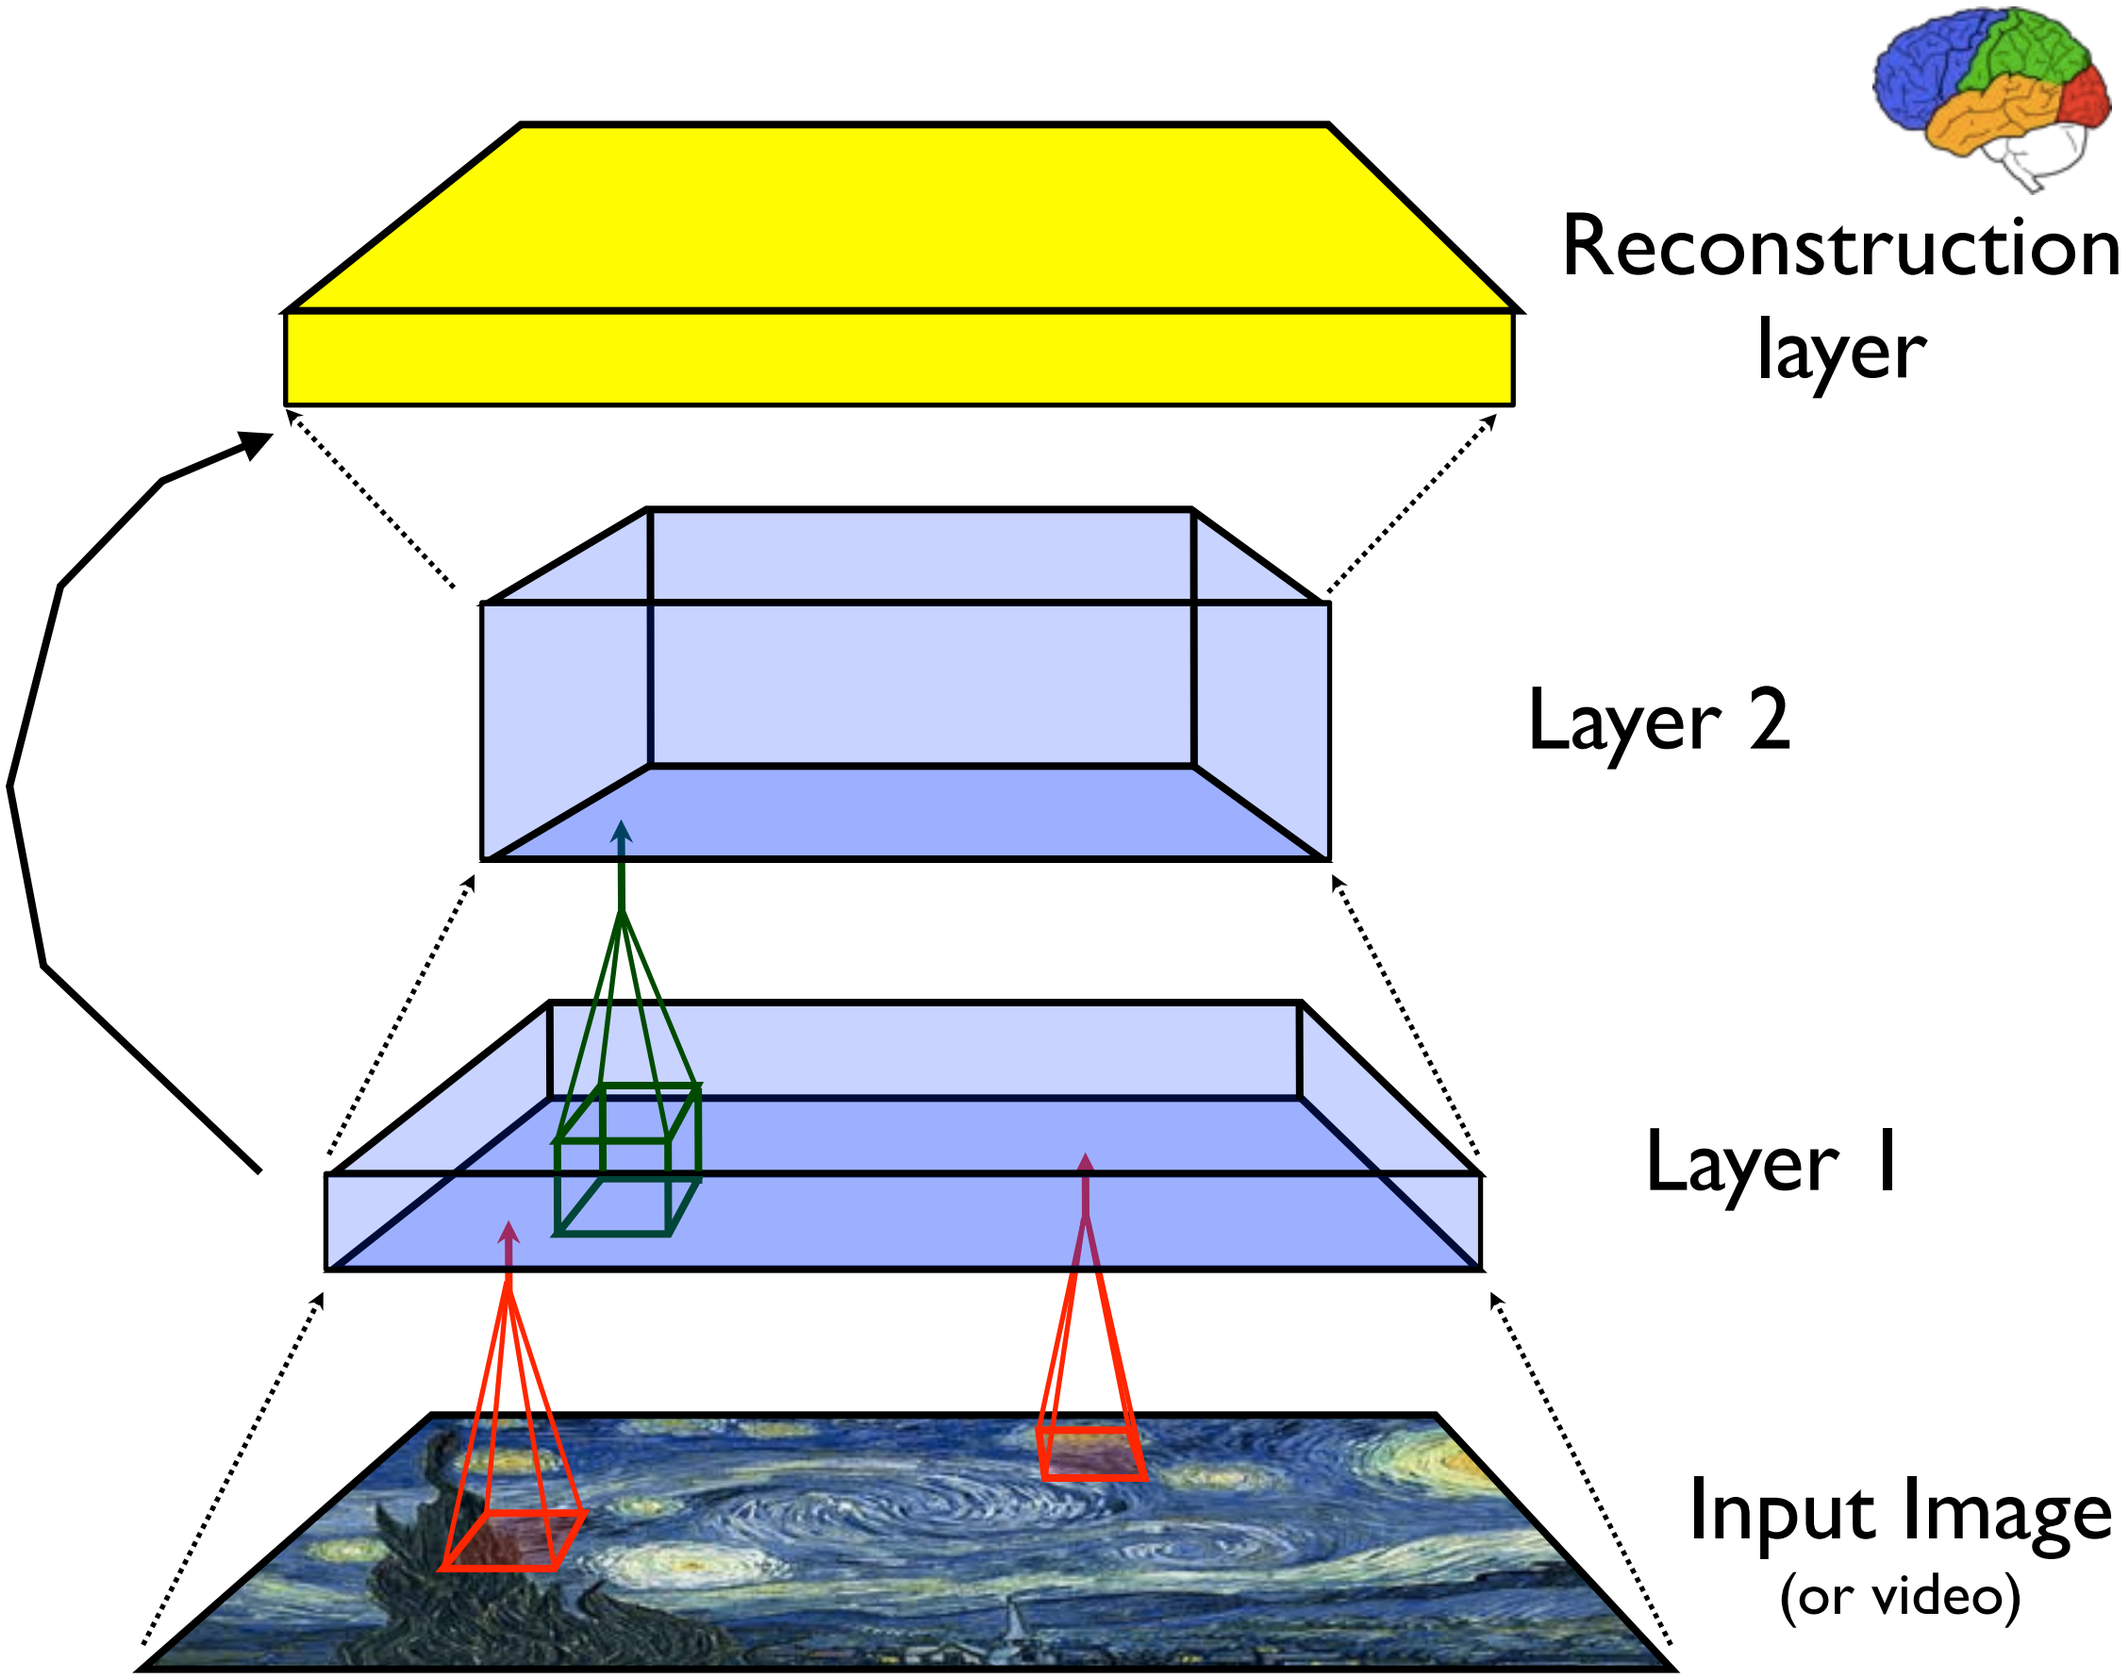
Reconstruction layer
Layer 2
Layer 1
Input Image
(or video)

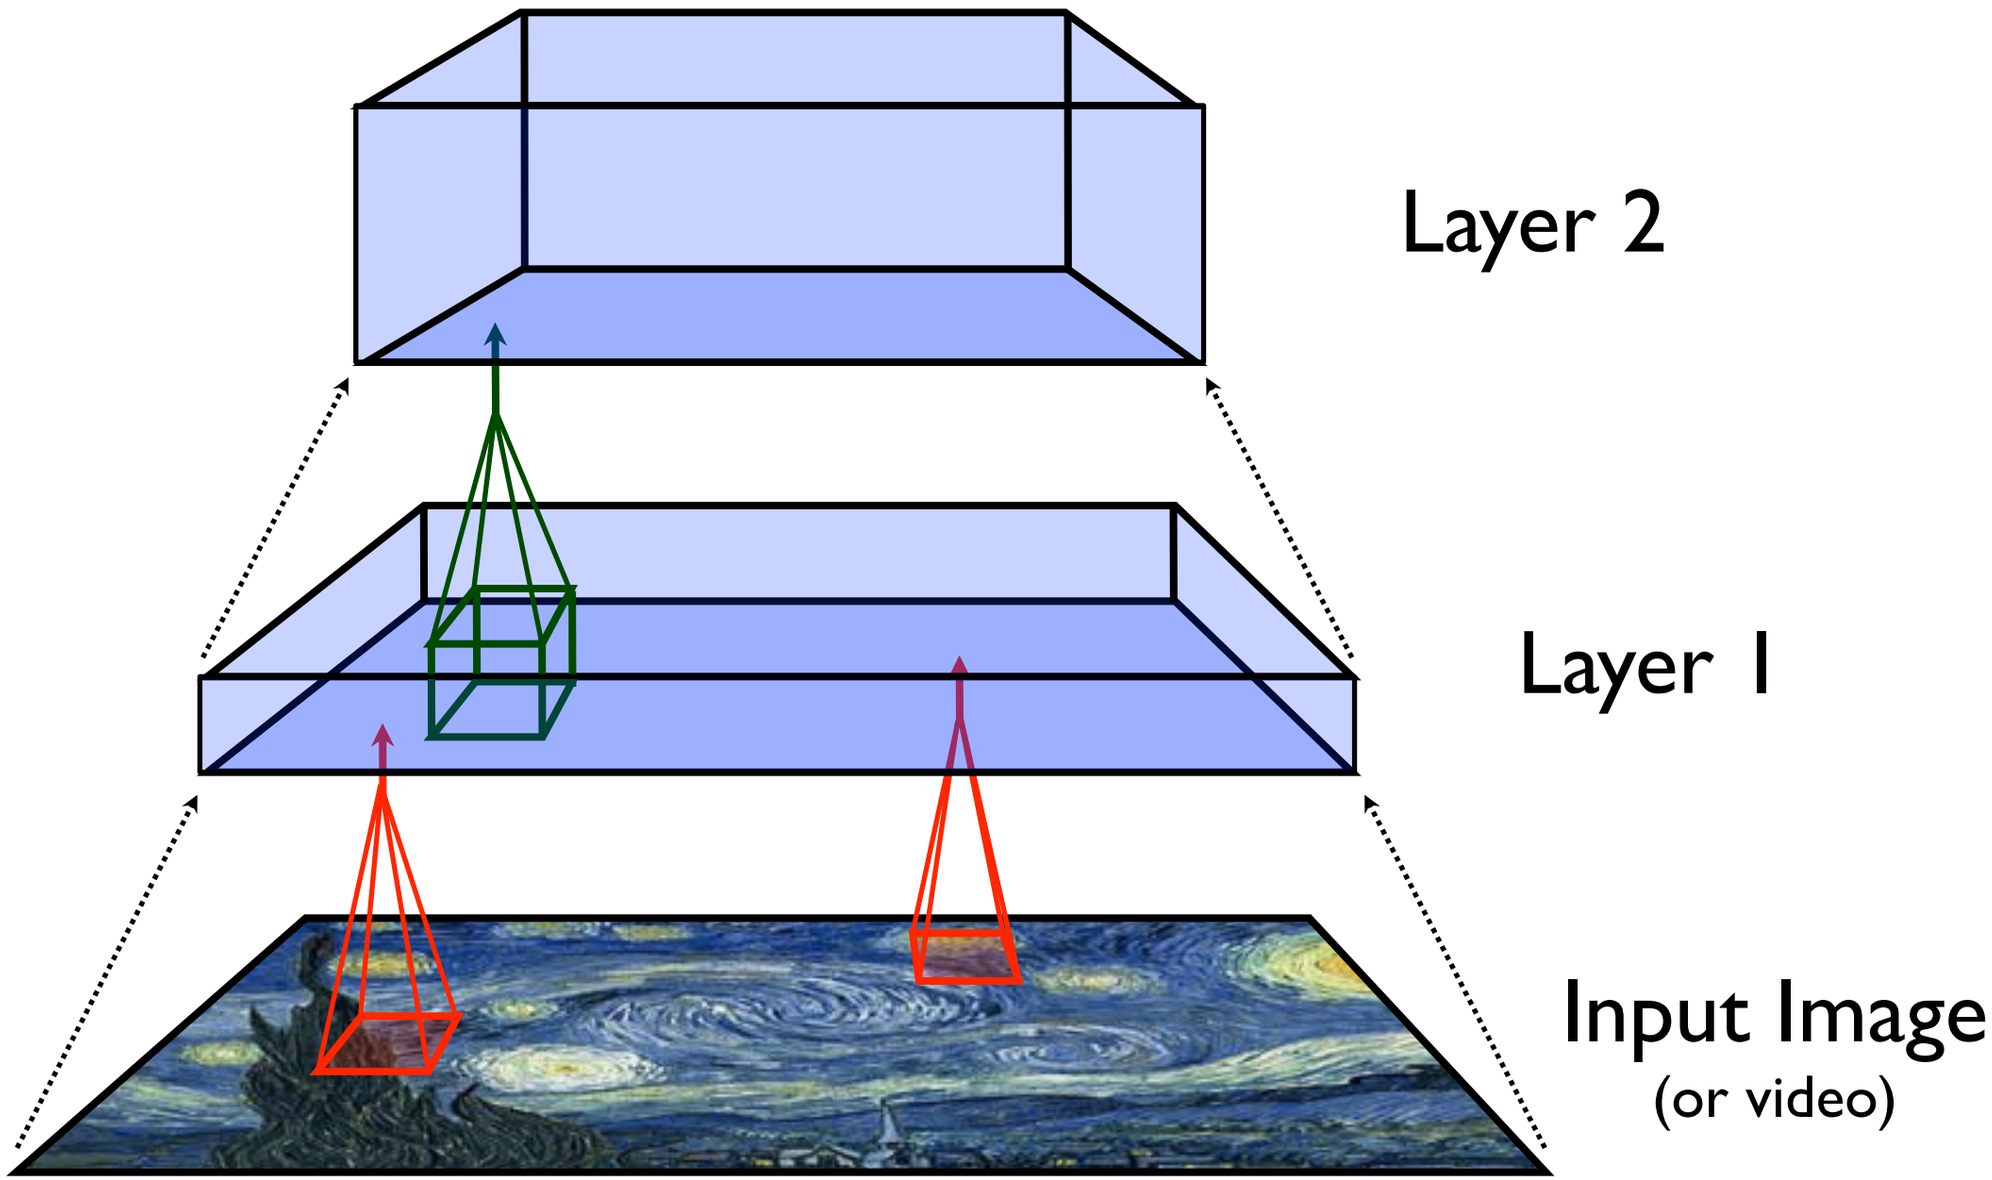
Layer 2
Layer 1
Input Image
(or video)

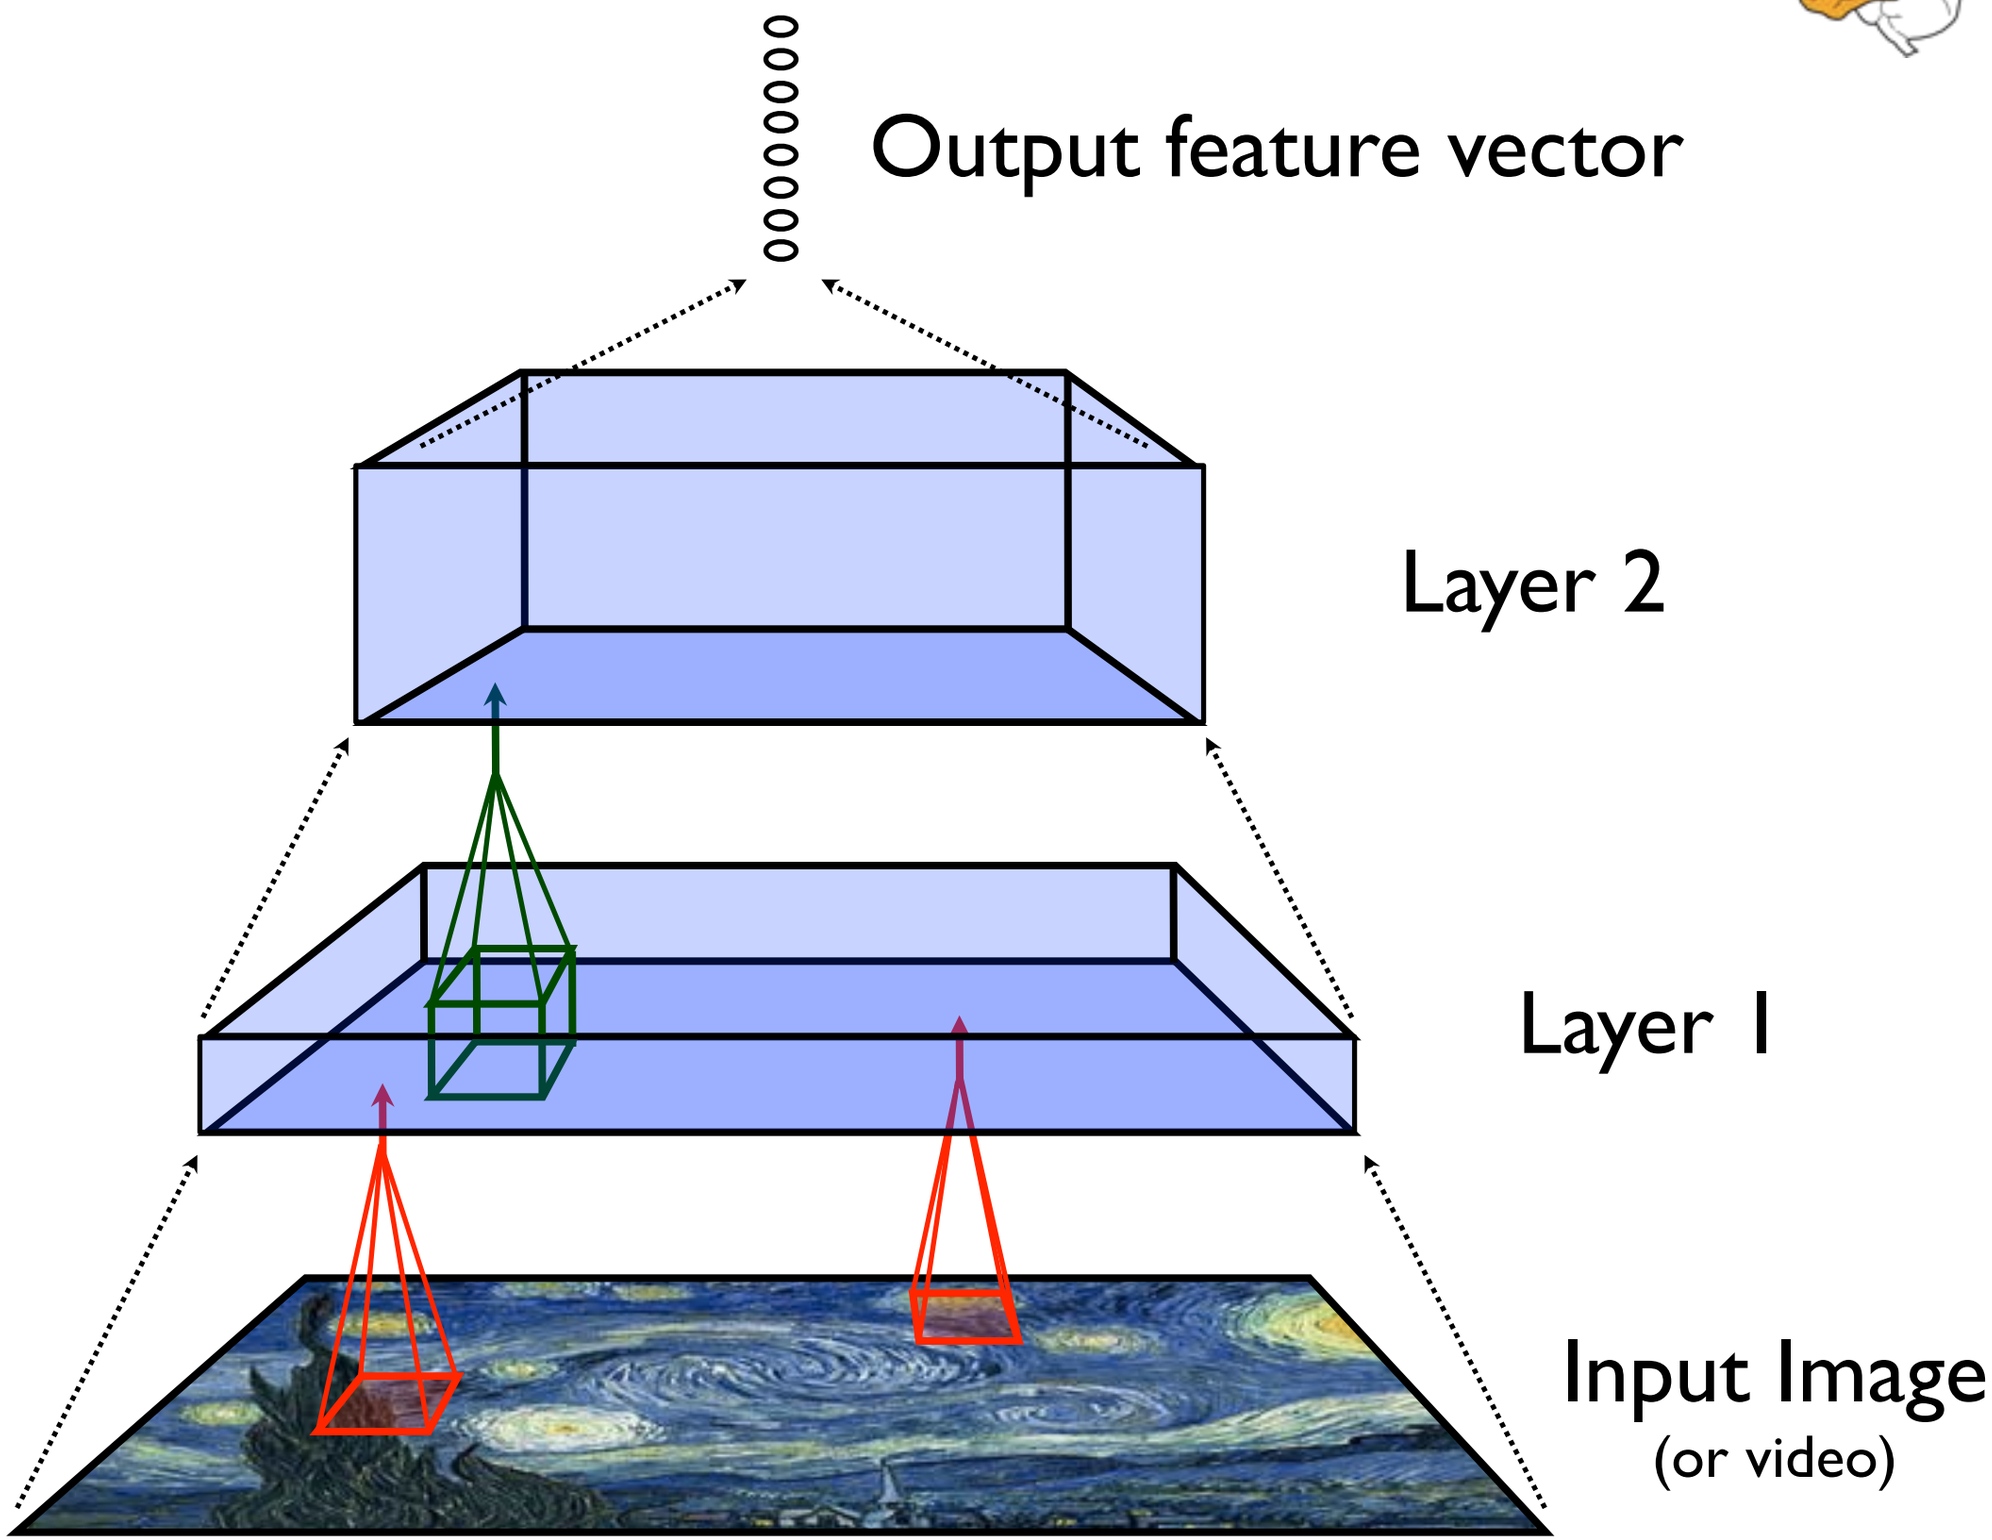
Output feature vector
Layer 2
Layer 1
Input Image
(or video)

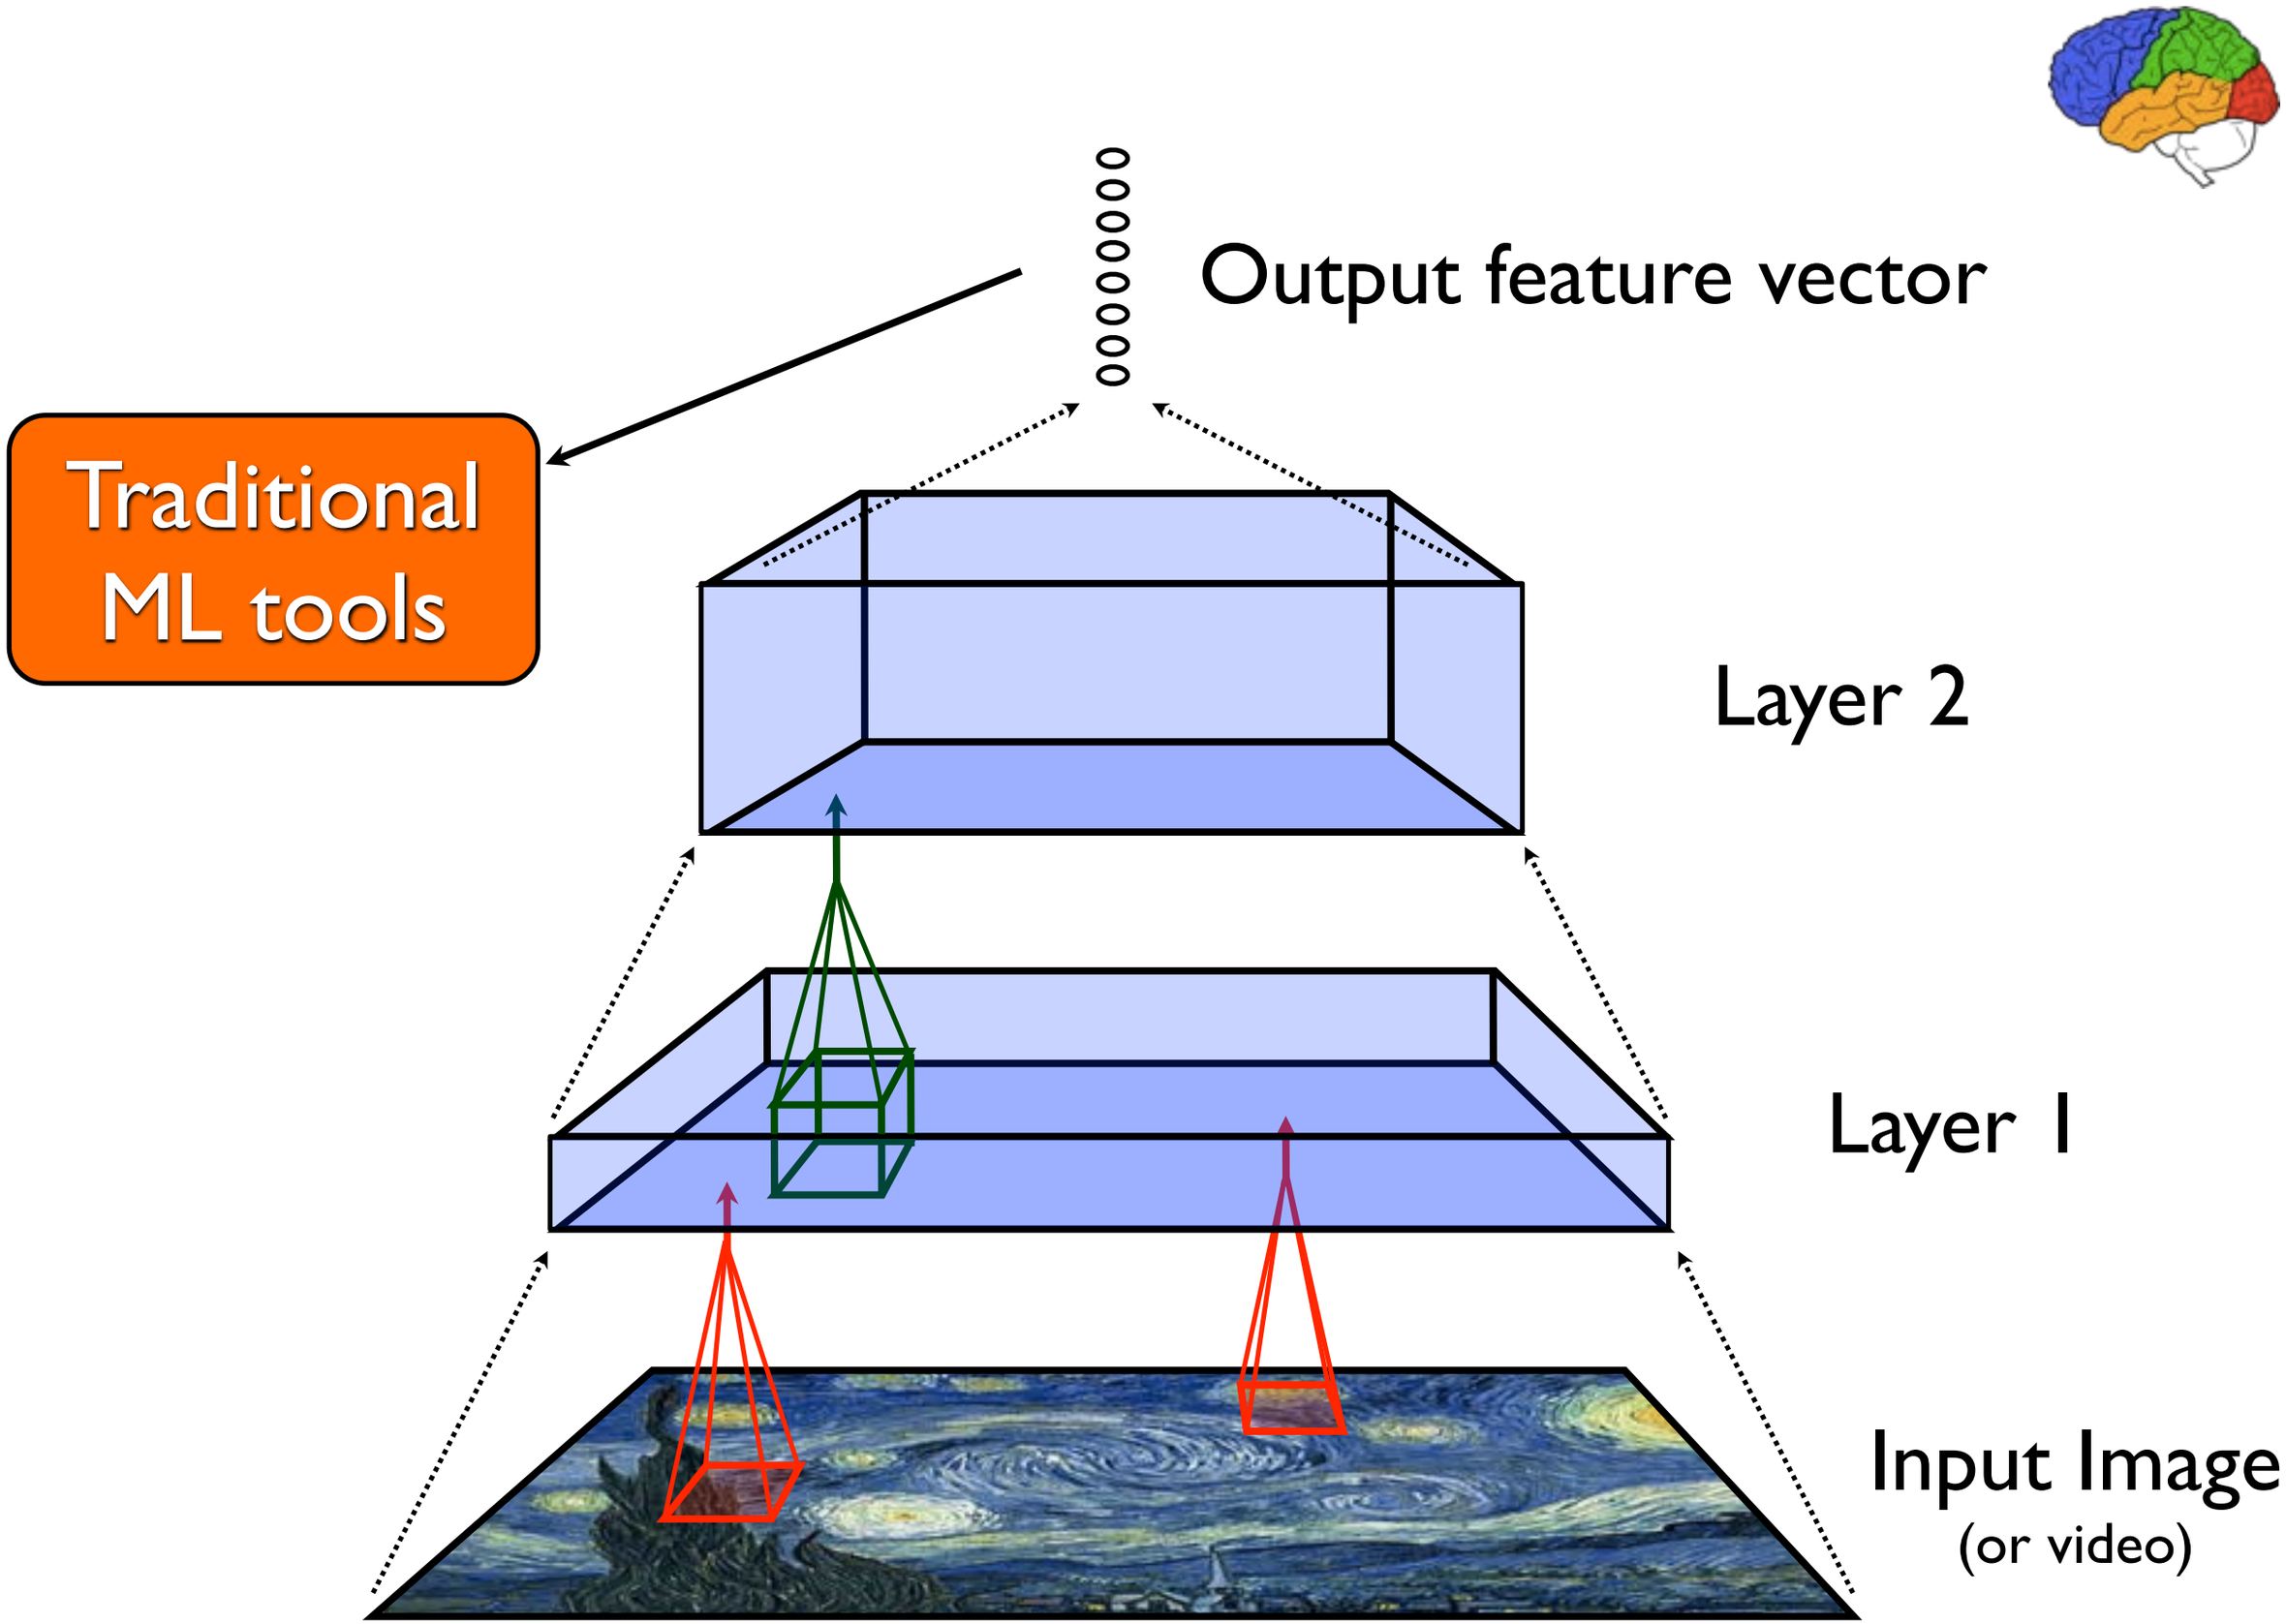
Output feature vector
Traditional ML tools
Layer 2
Layer 1
Input Image
(or video)

Supports arbitrary models, but best for large models with mostly local connections

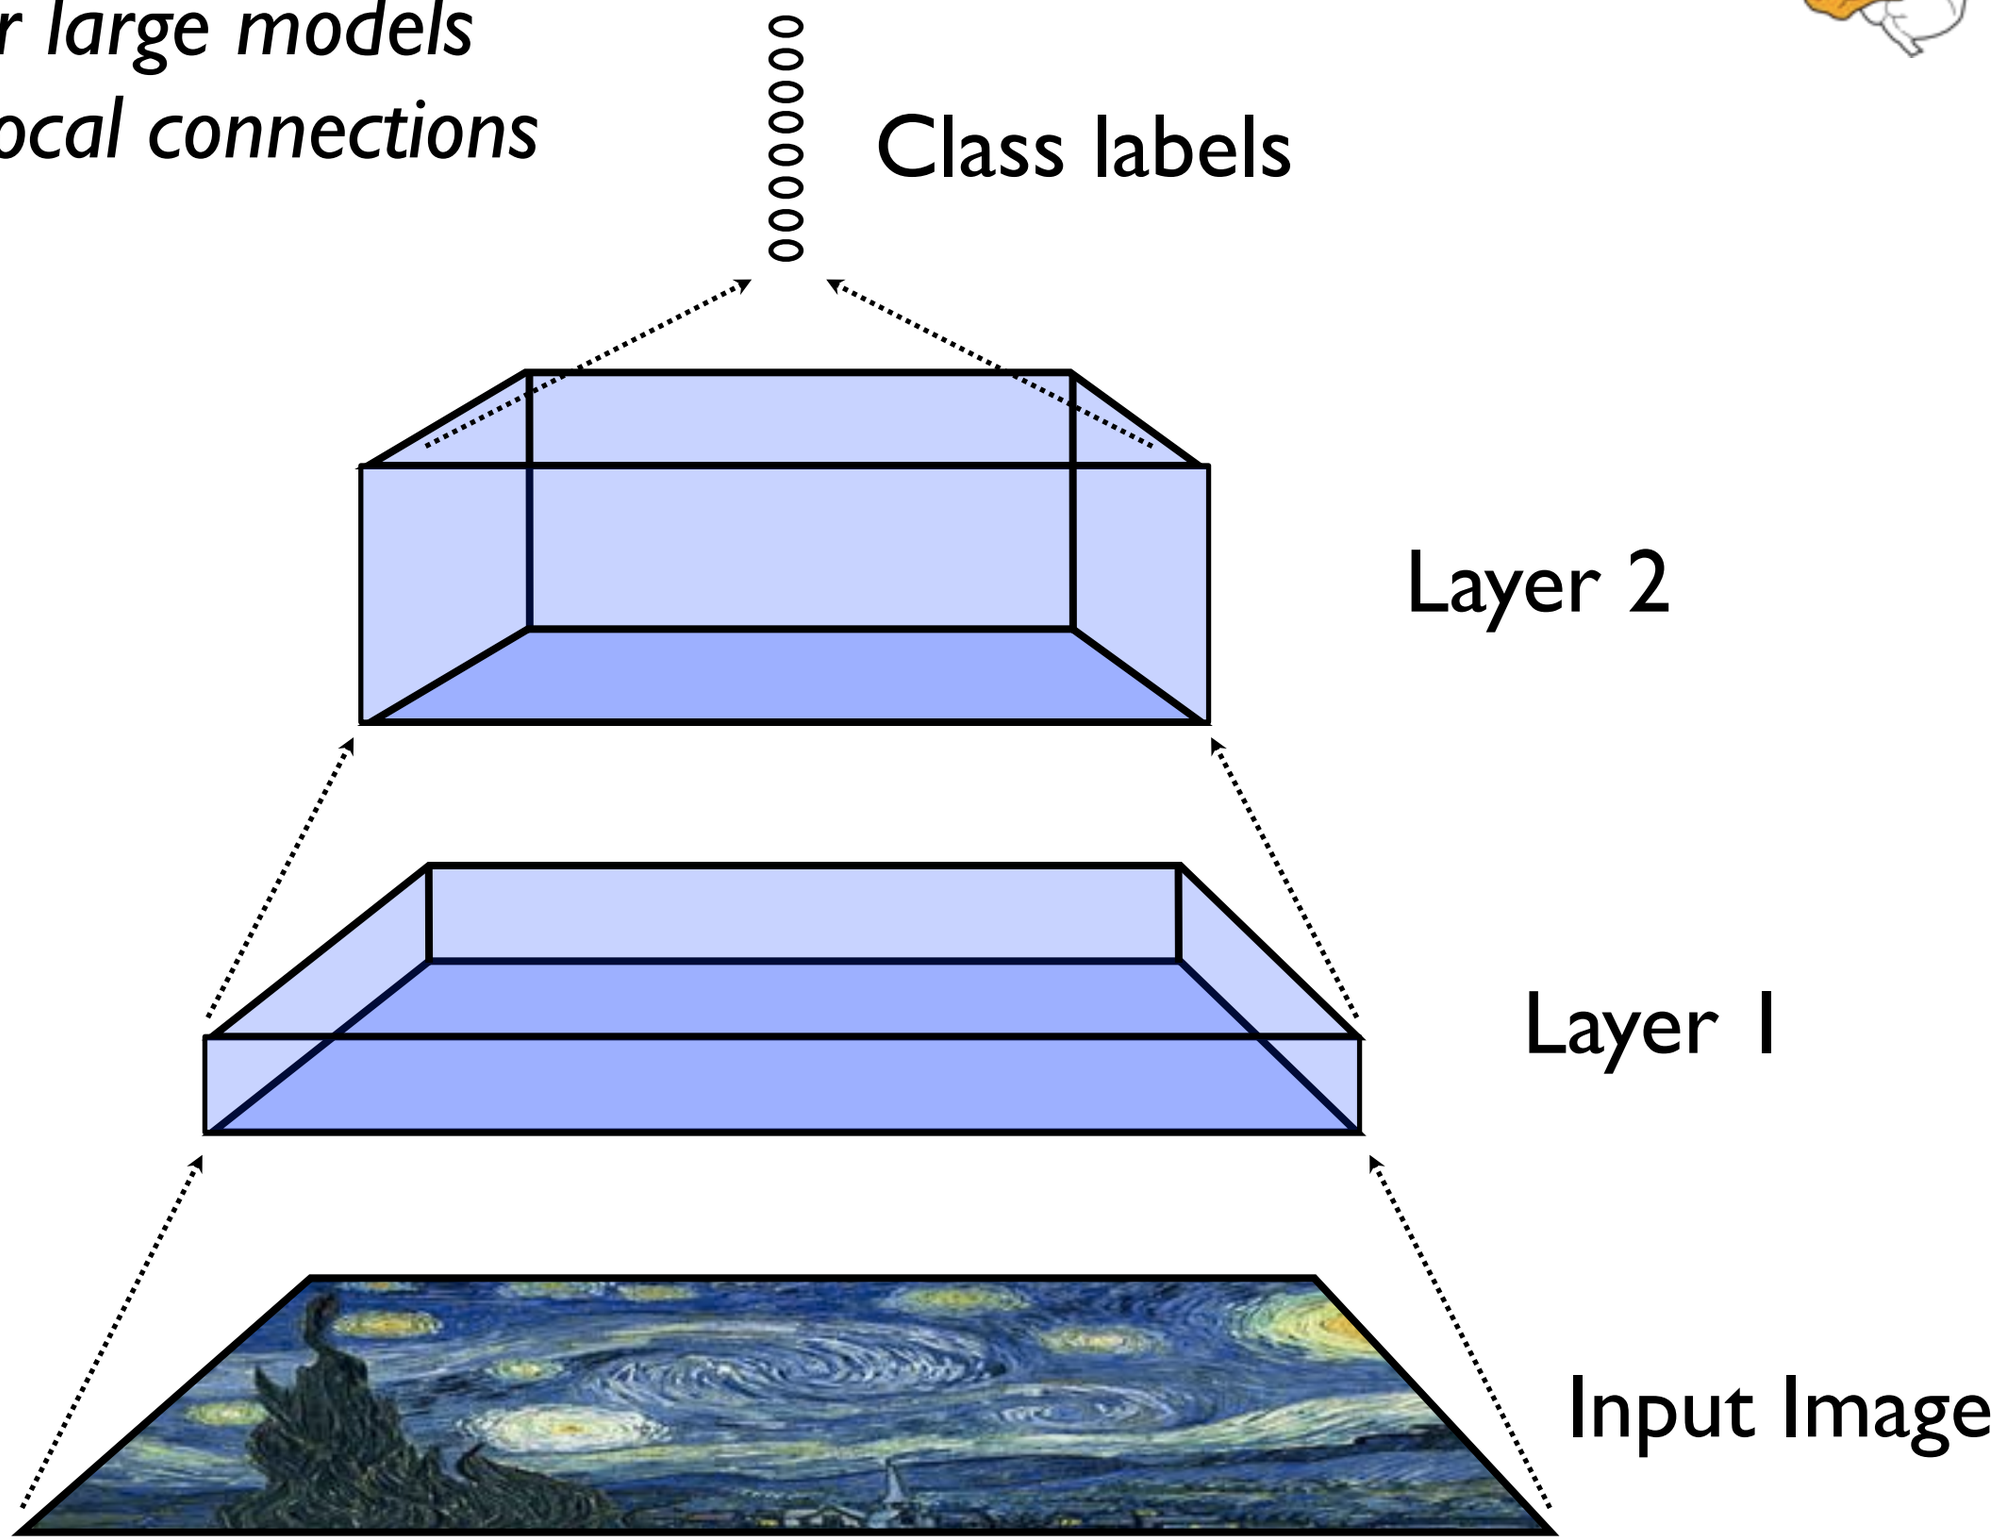

Supports arbitrary models, but best for large models with mostly local connections

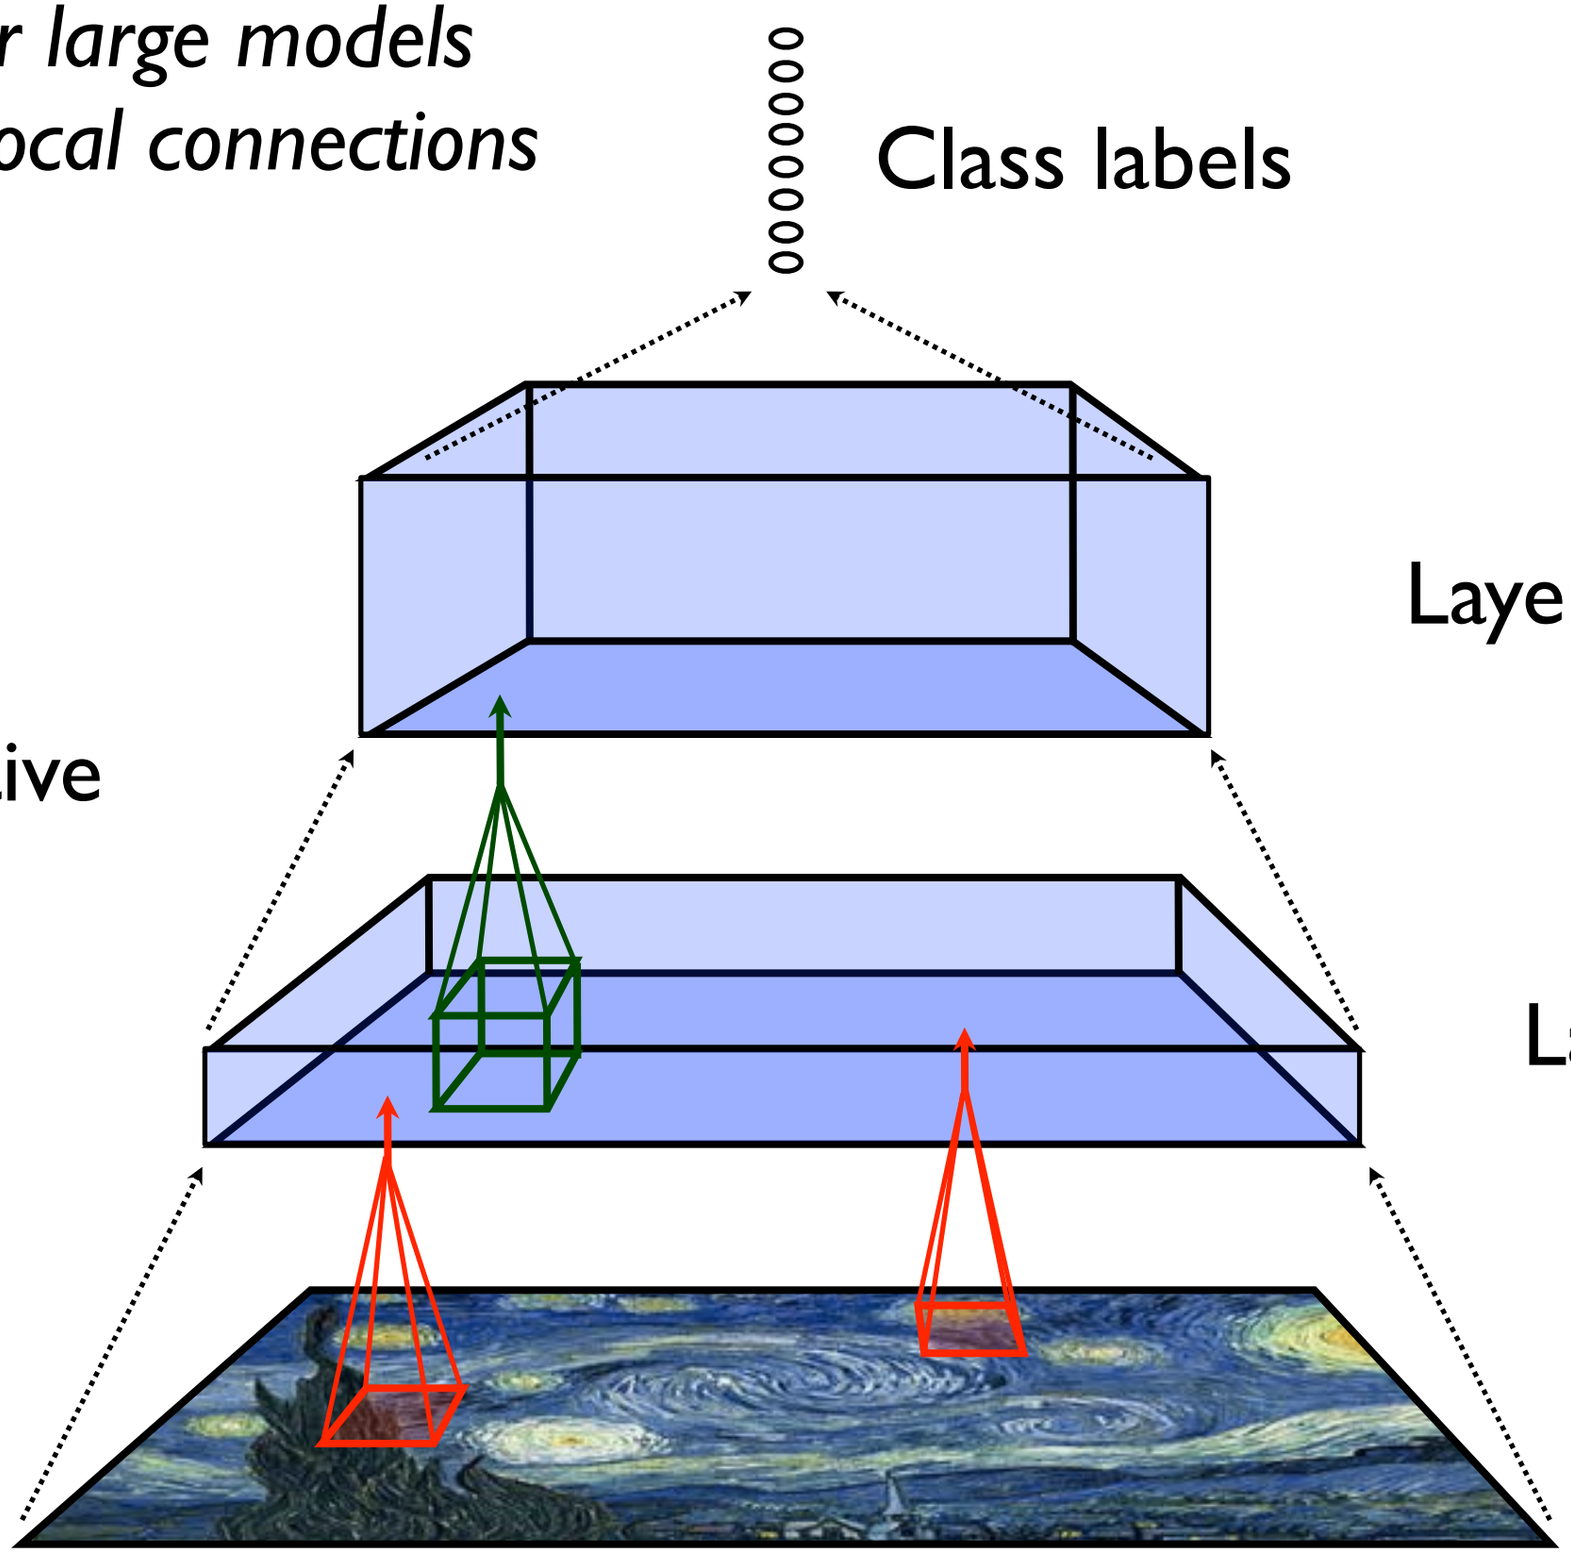

Partition model across machines

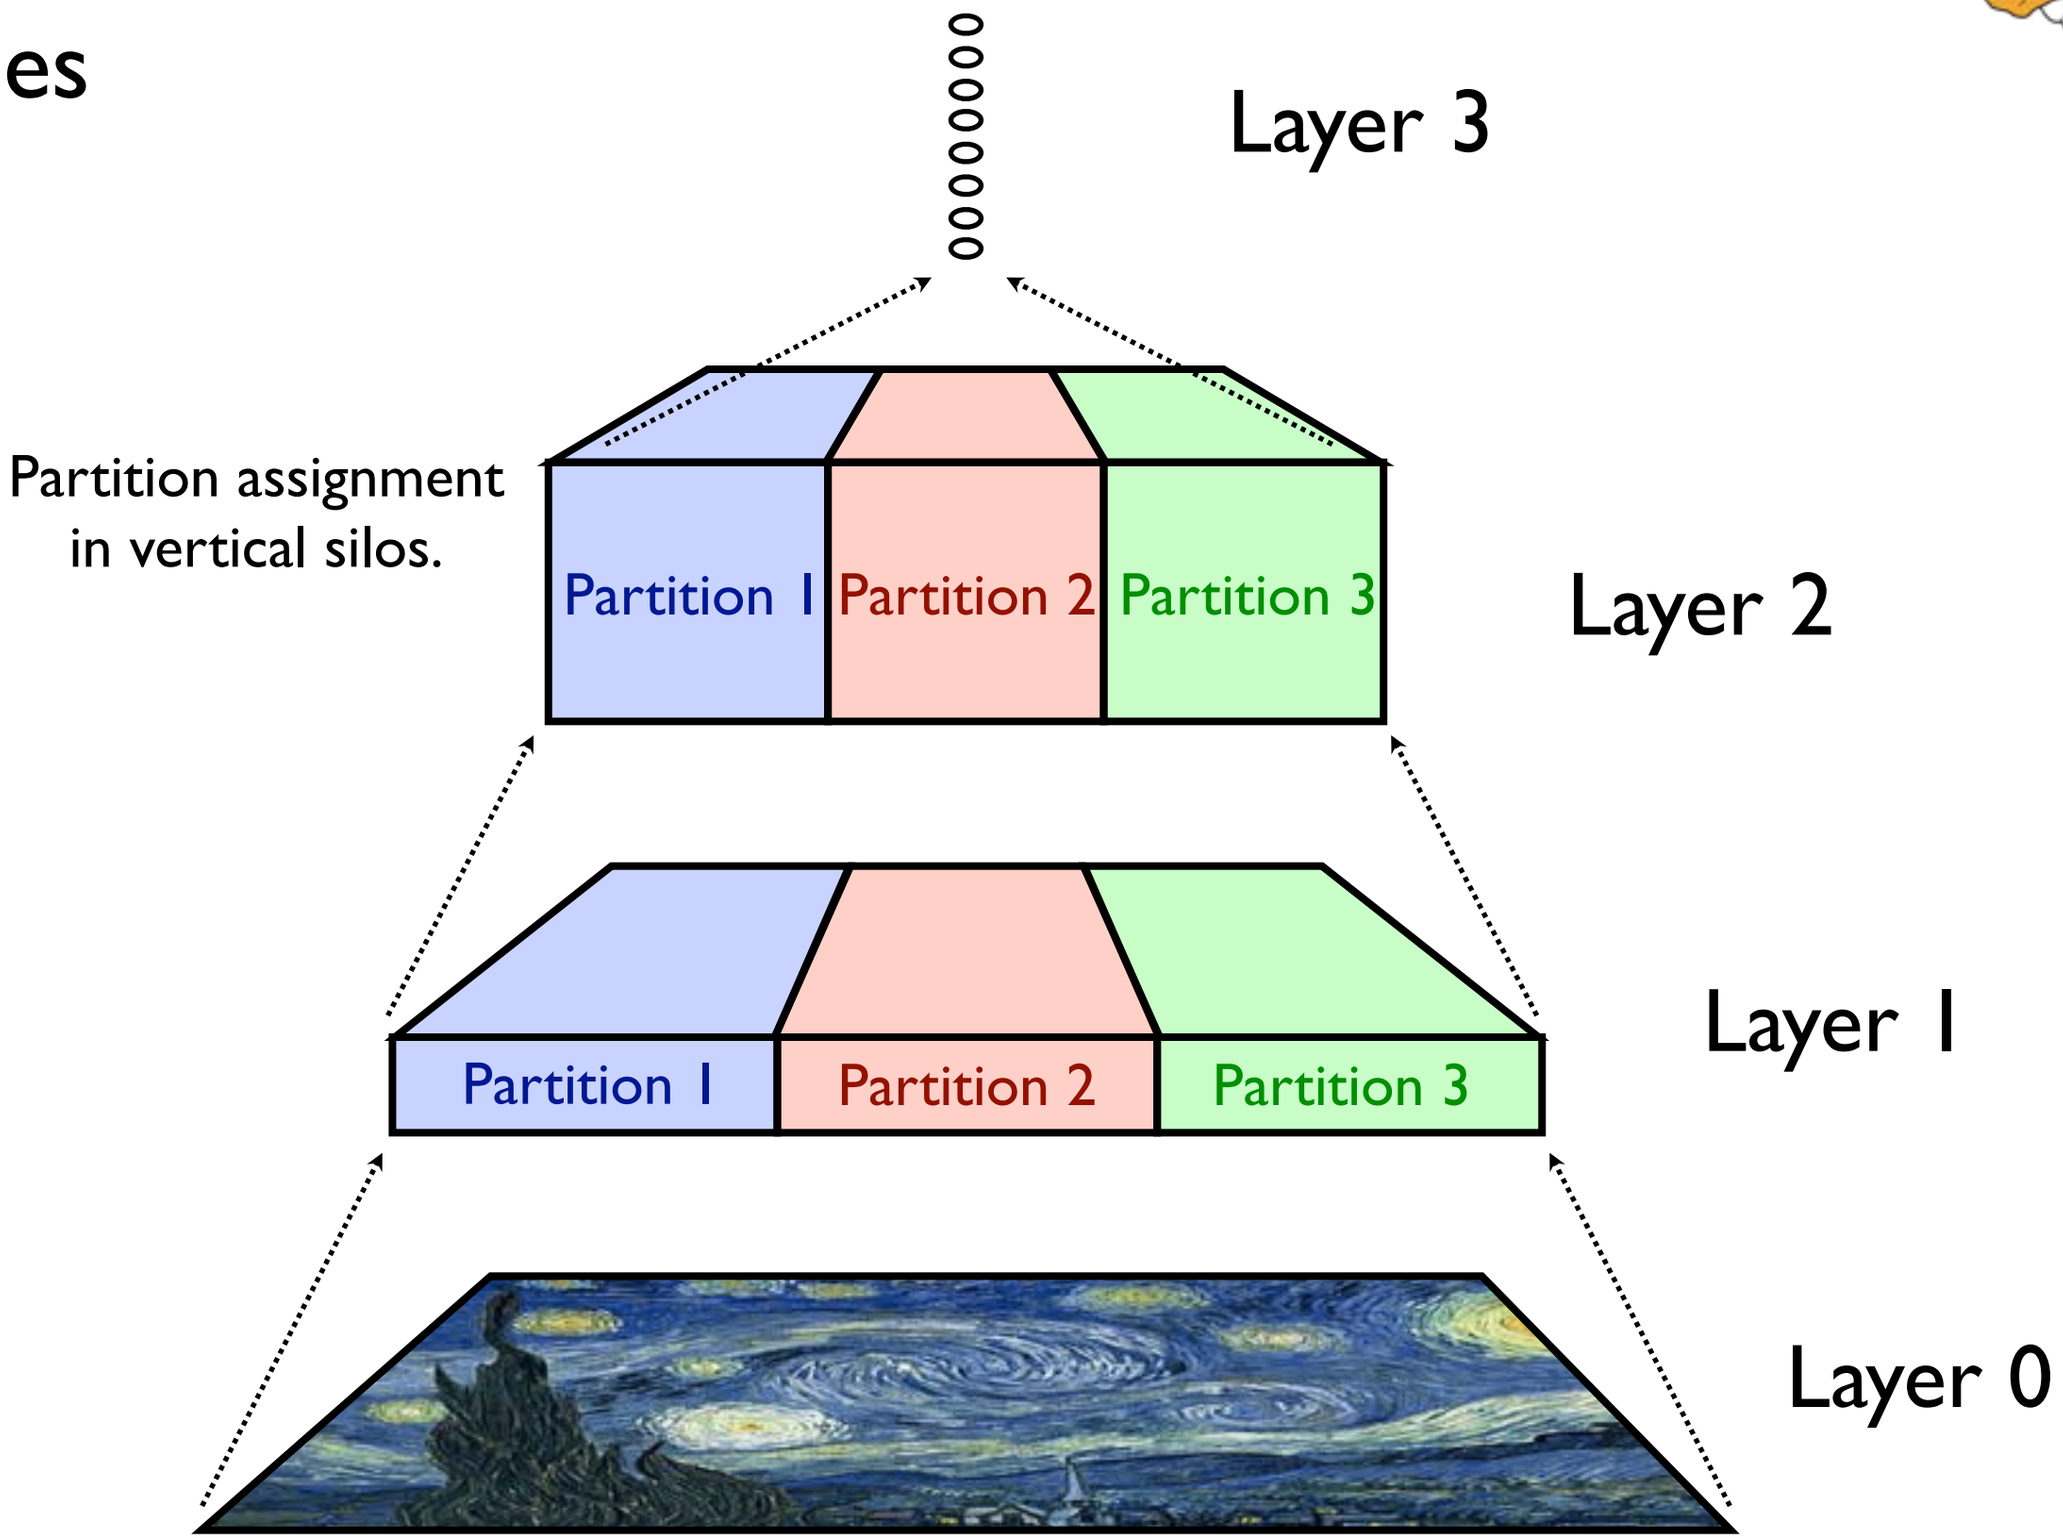

Partition model across machines

Layer 3

Partition assignment in vertical silos.

Layer 2

Minimal network traffic: The most densely connected areas are on the same partition

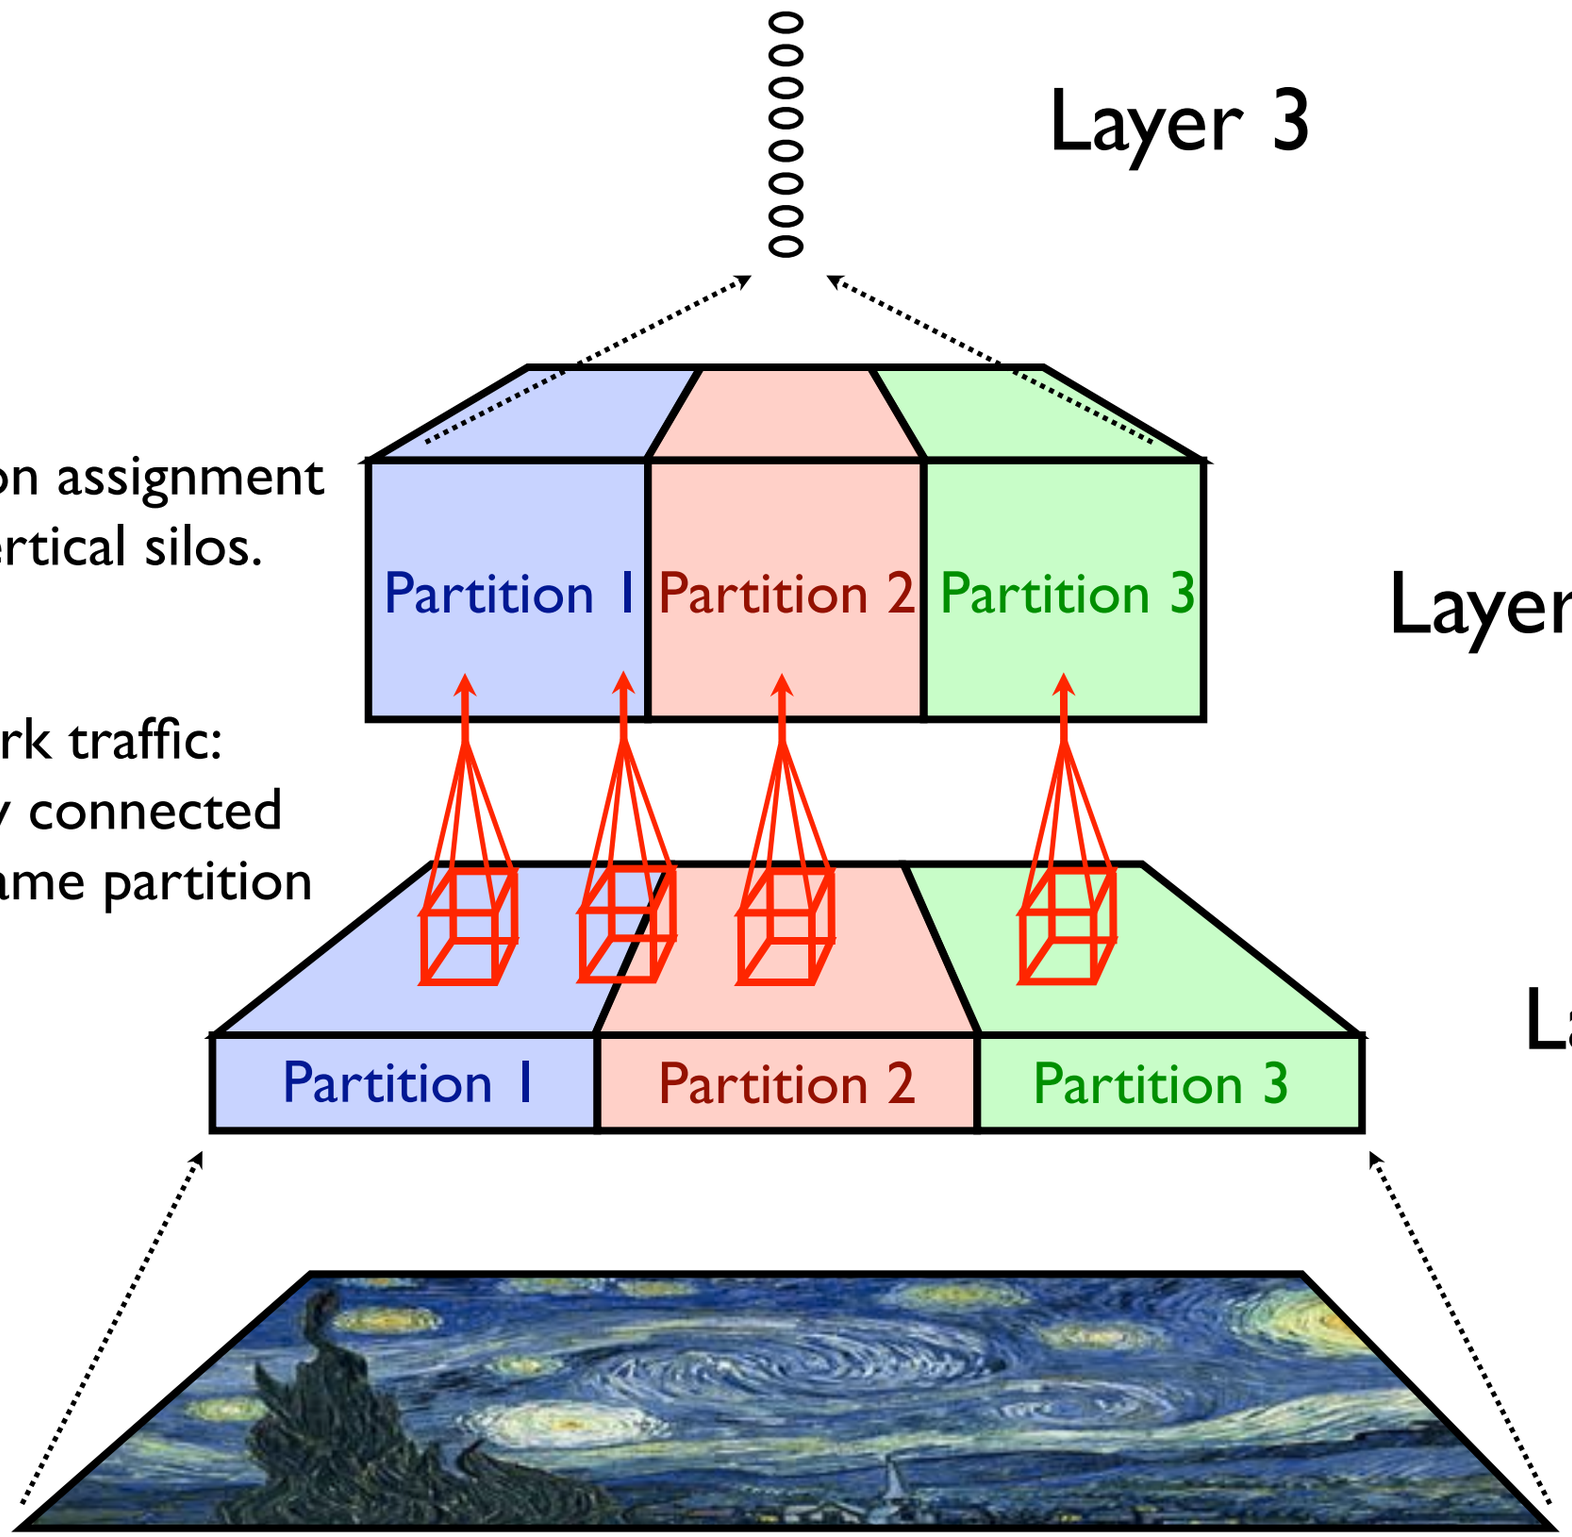

Layer 1

Layer 0

Partition model across machines

Partition assignment in vertical silos.

Minimal network traffic: The most densely connected areas are on the same partition

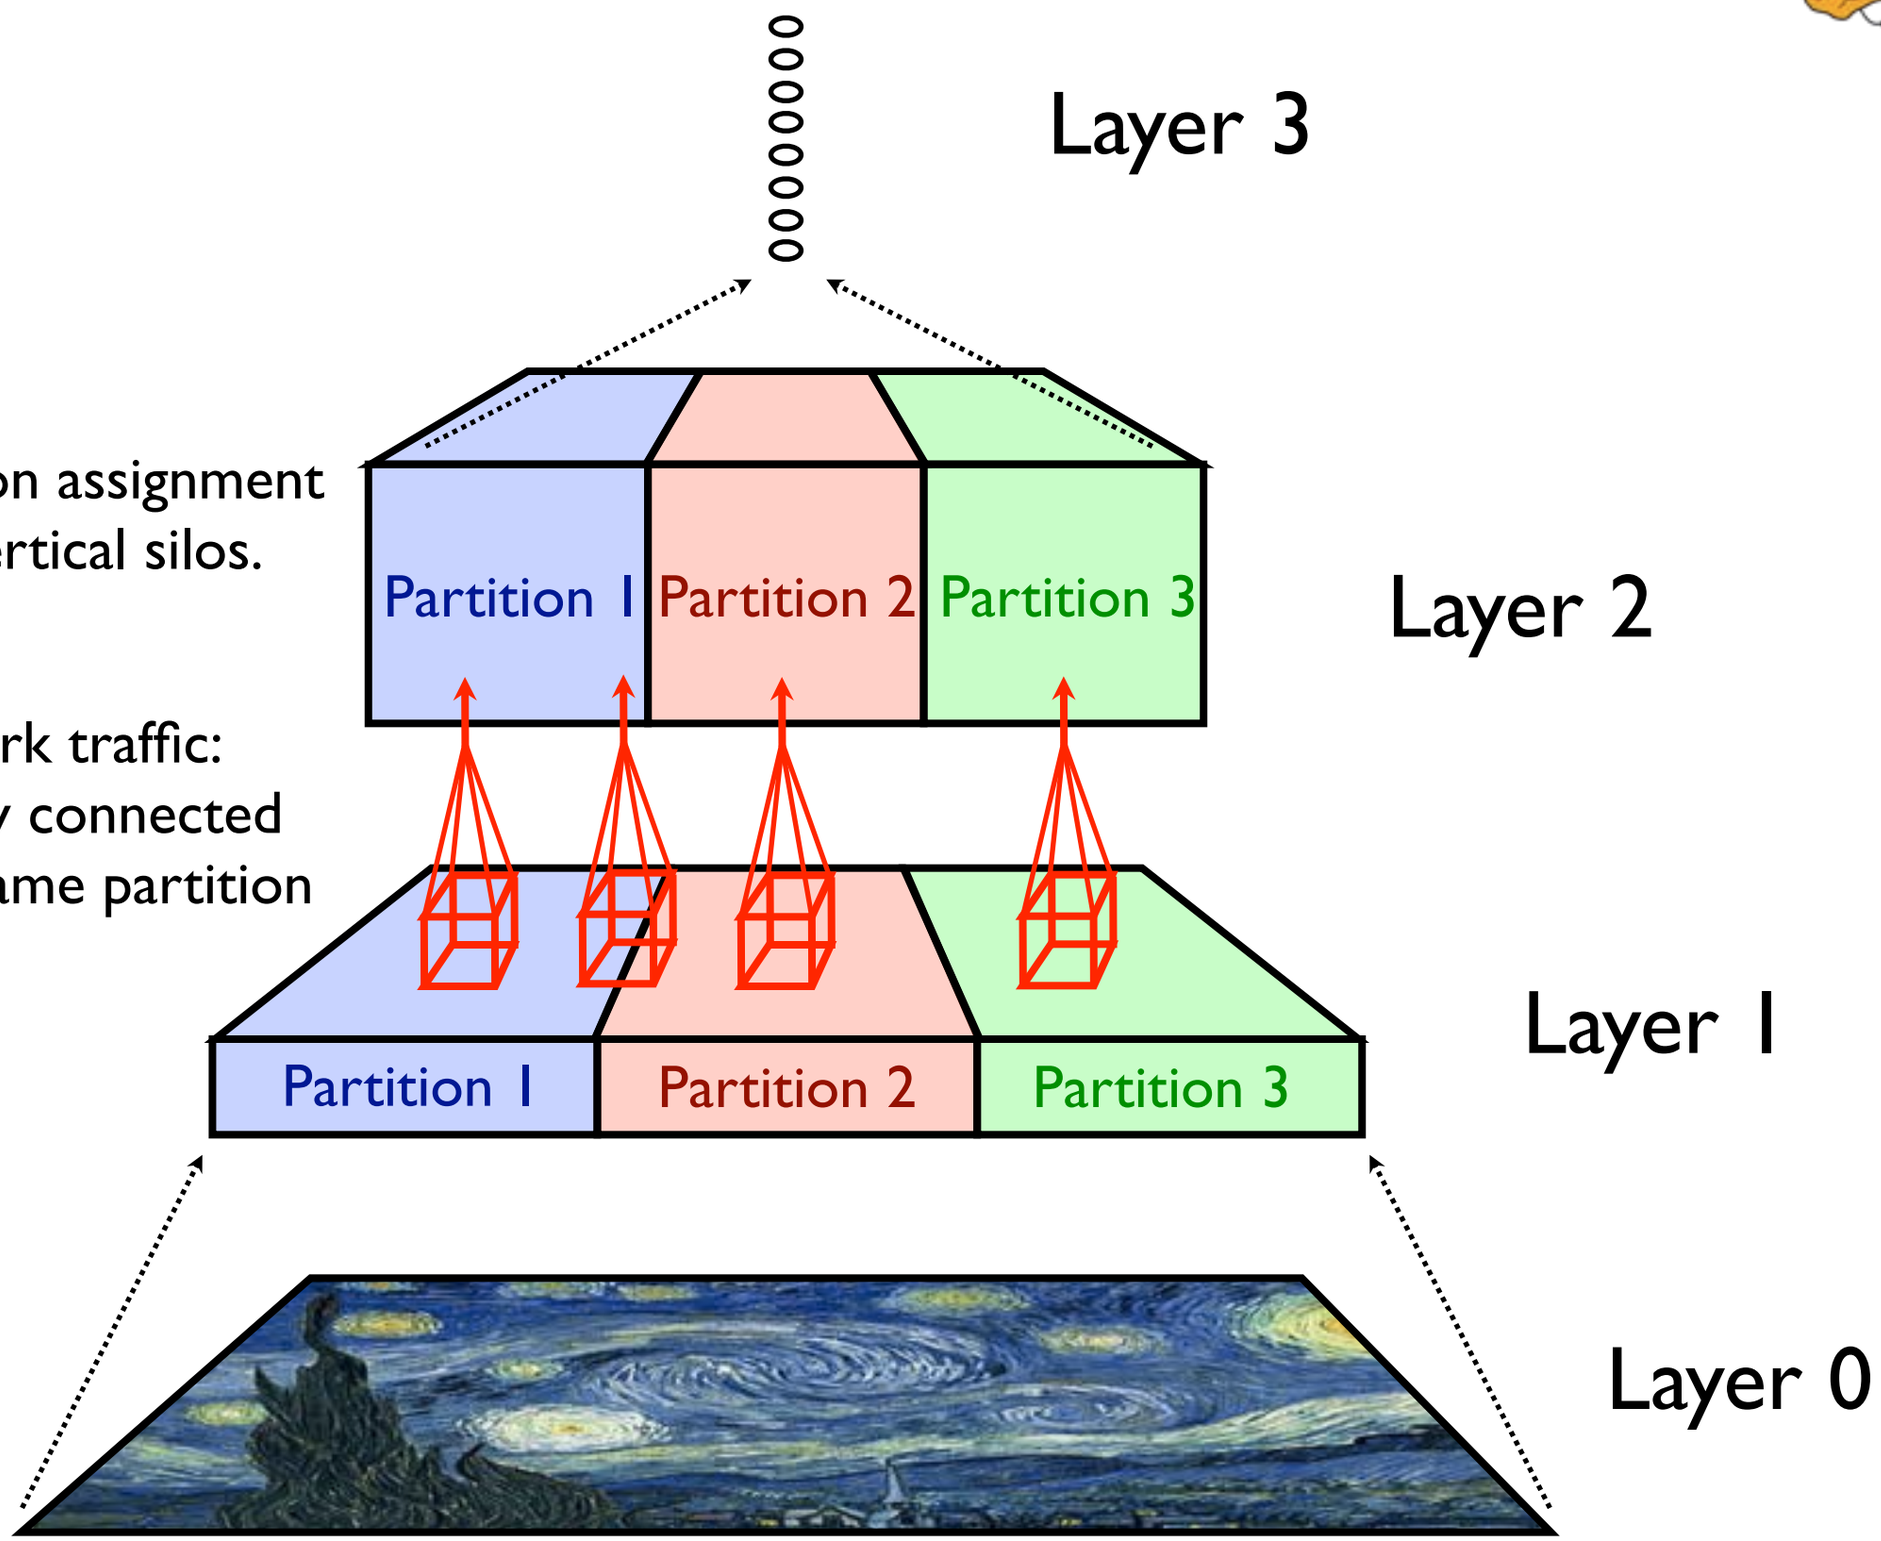

Our biggest models: 144 machines, ~2300 cores

Basic DistBelief Model Training

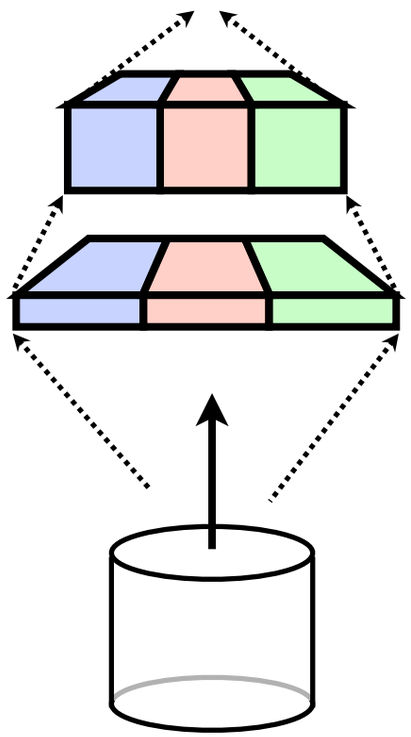

Making a single model bigger and faster was the right first step.

But training is still slow with large data sets if a model only considers tiny minibatches (10-100s of items) of data at a time.

How can we add another dimension of parallelism, and have multiple model instances train on data in parallel?

Asynchronous Distributed Stochastic Gradient Descent

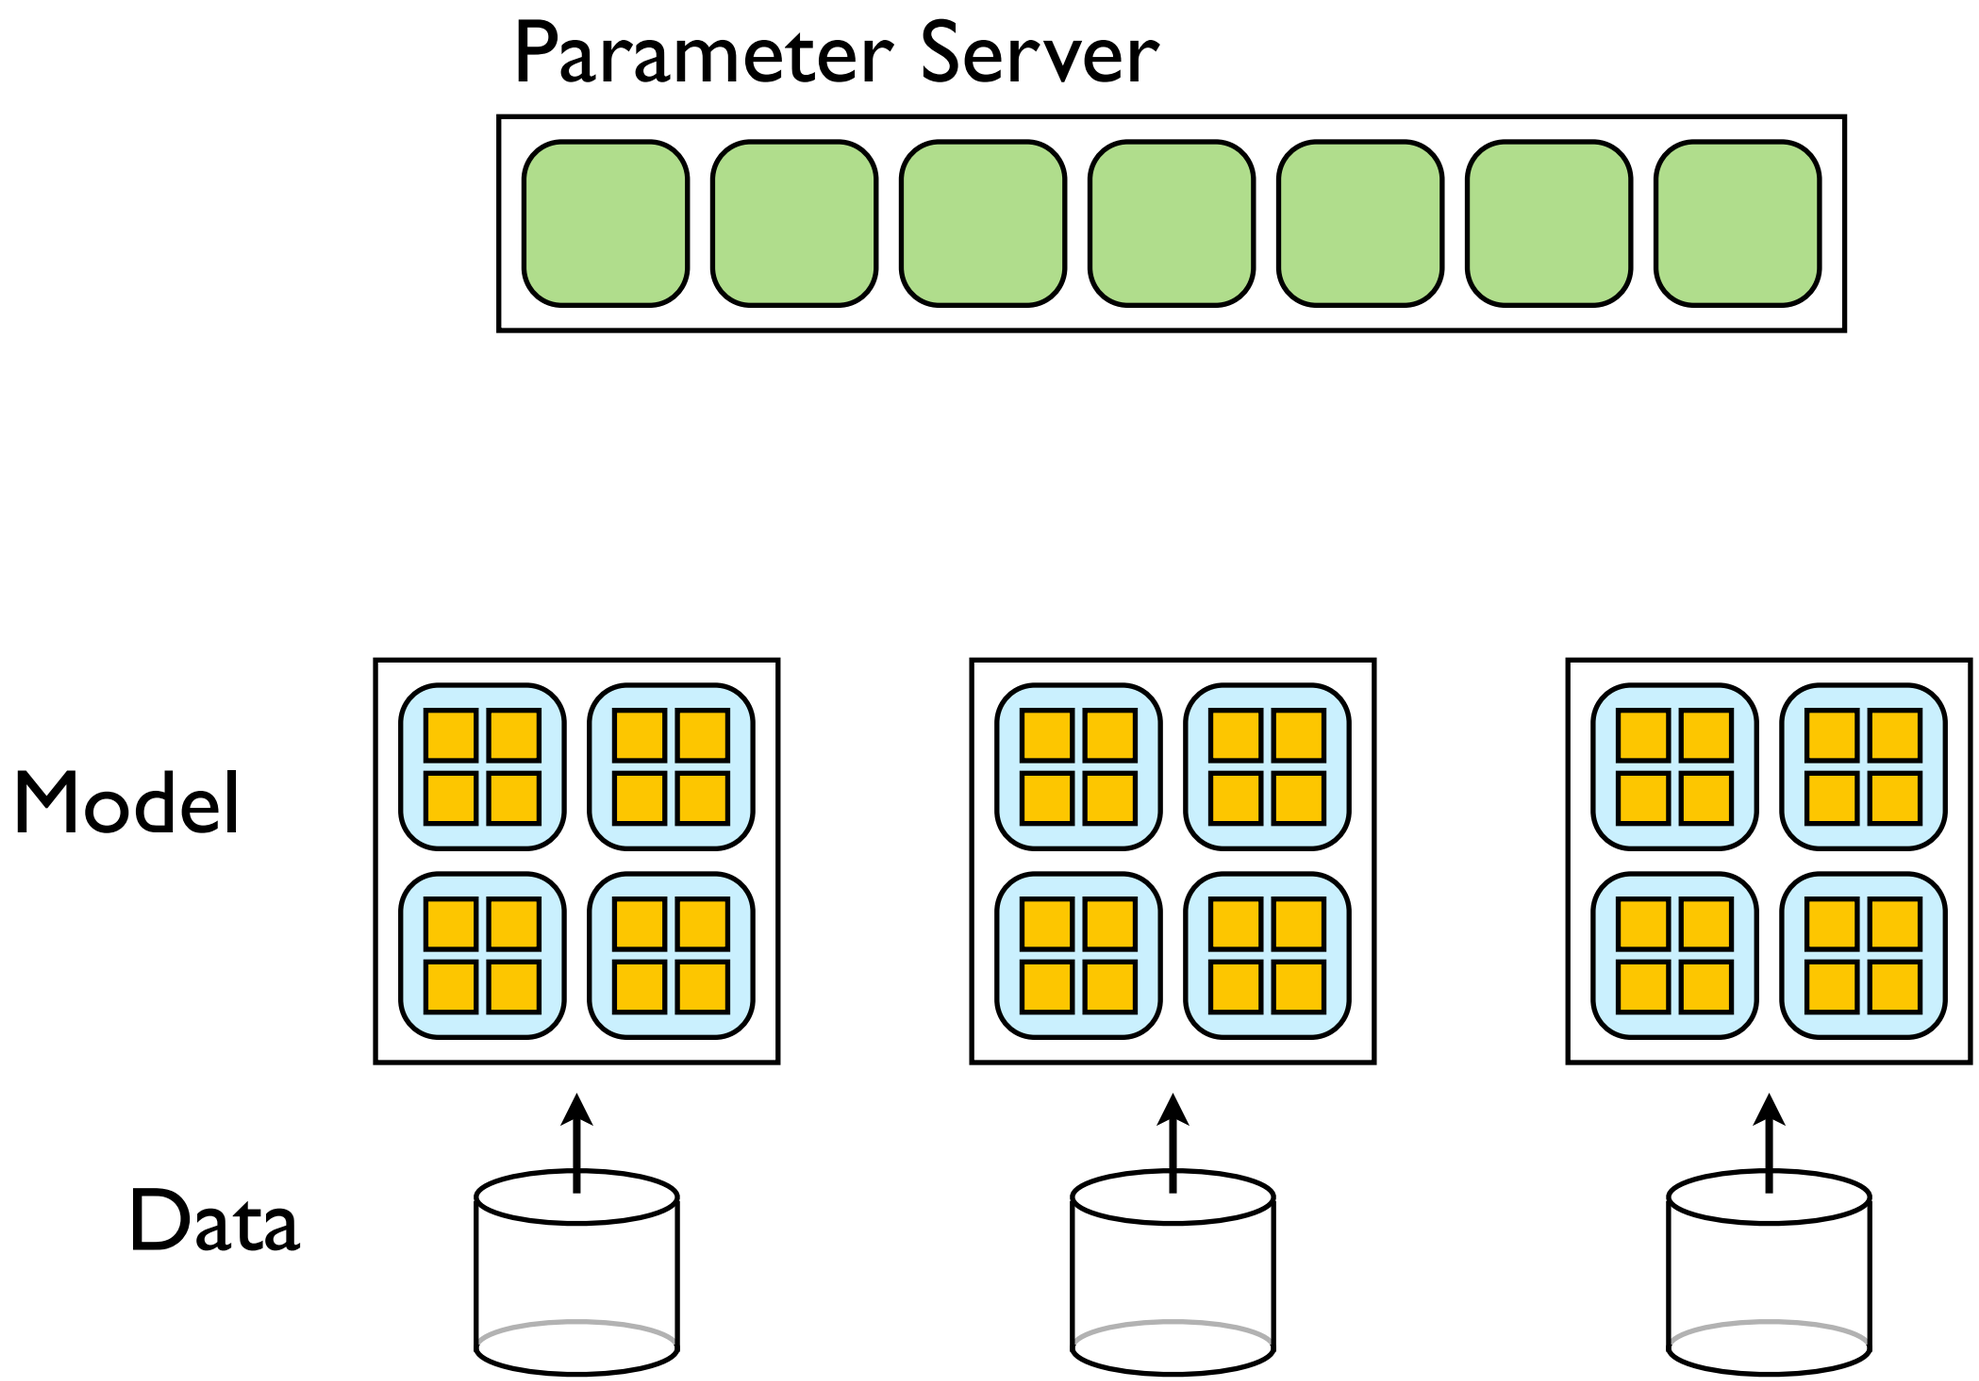

Asynchronous Distributed Stochastic Gradient Descent

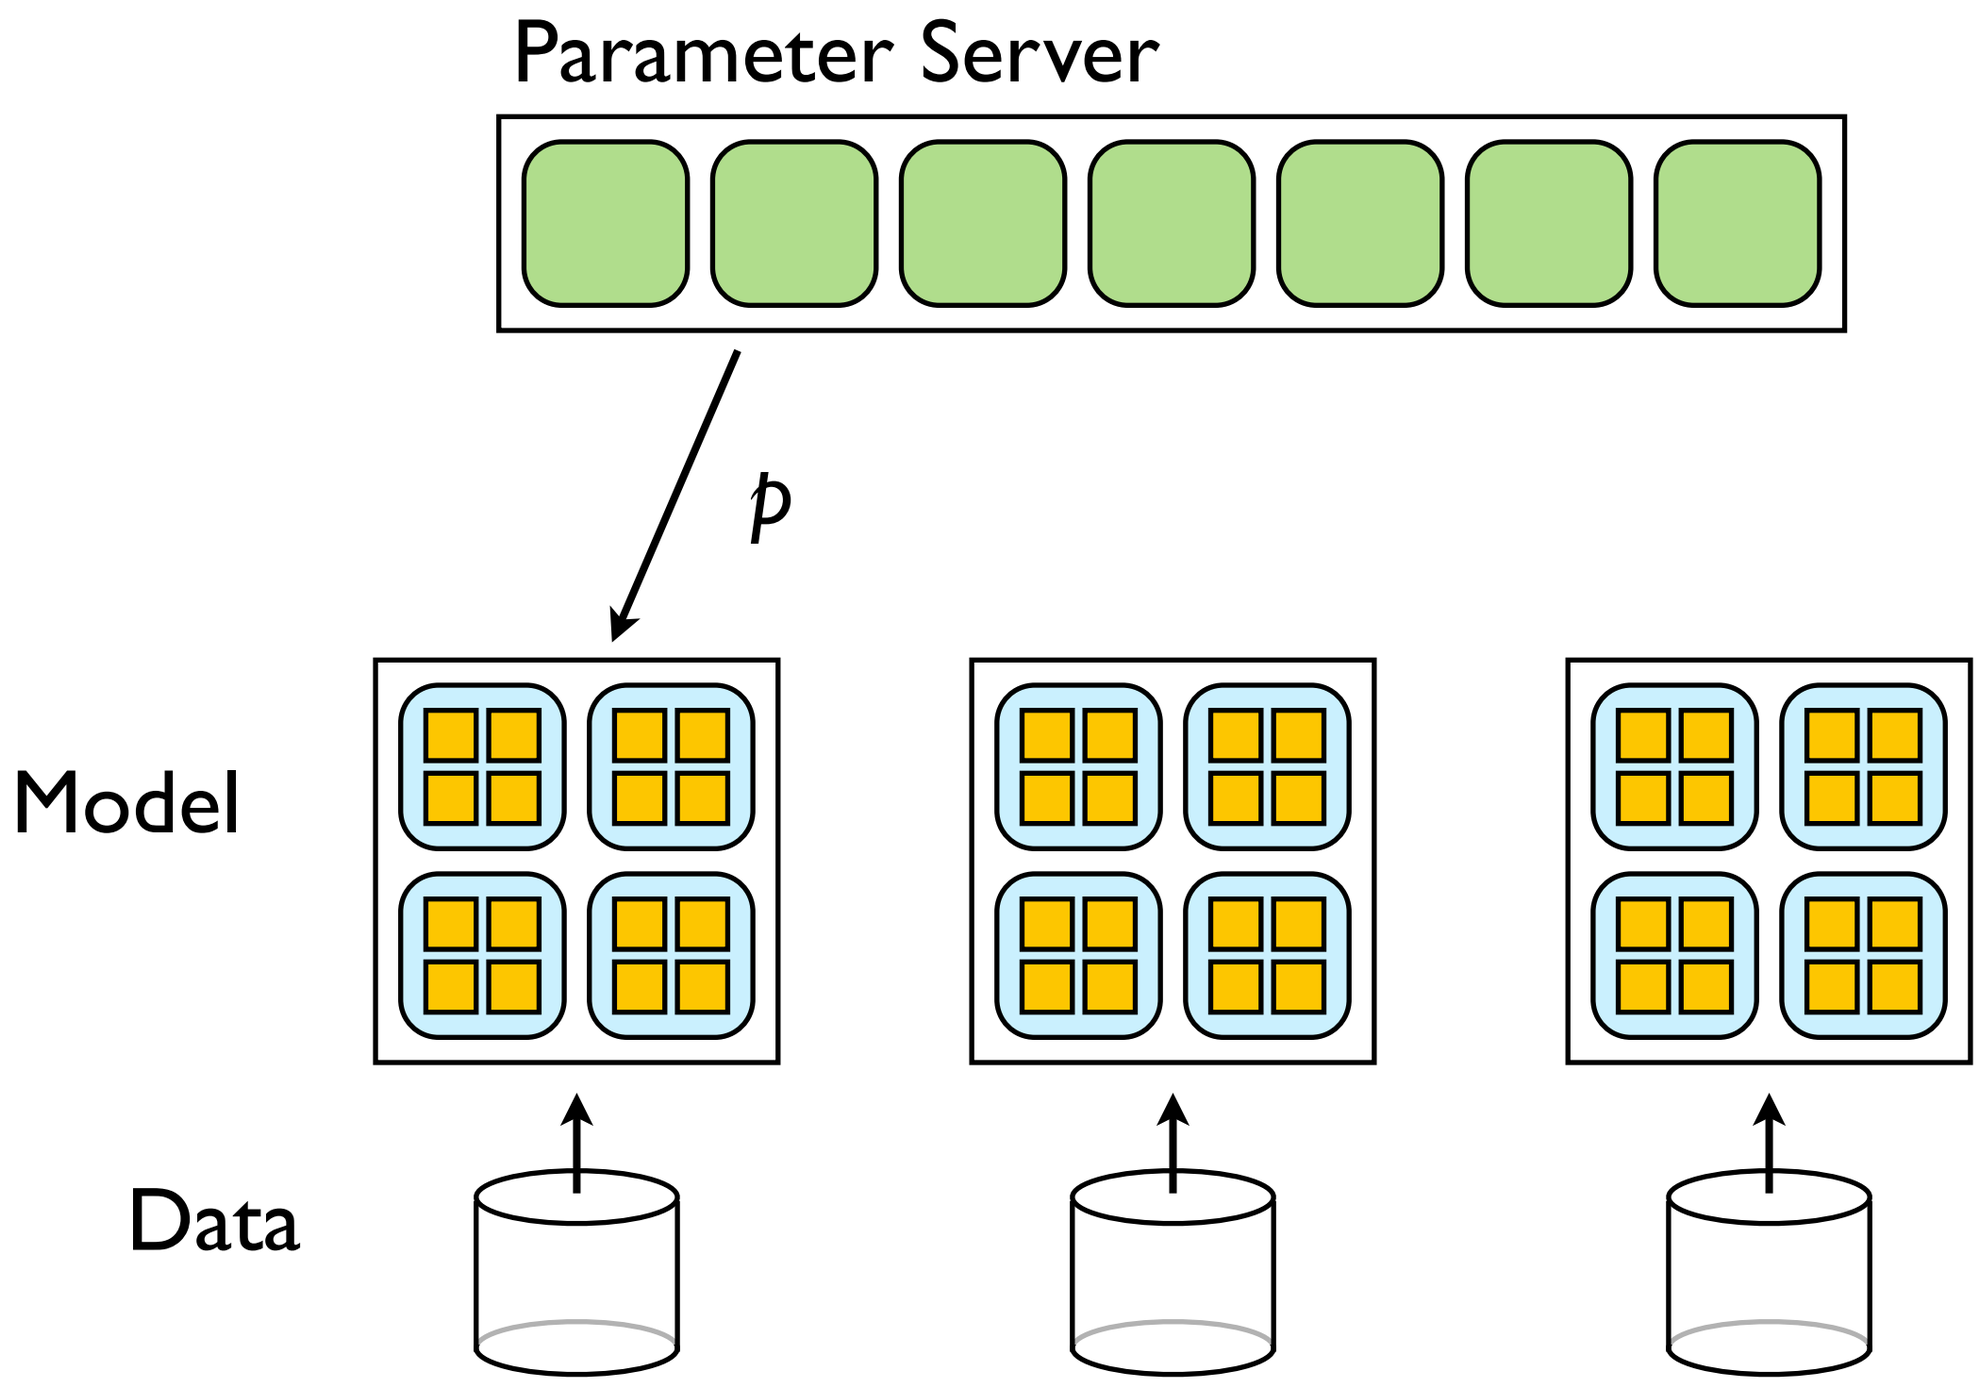

Asynchronous Distributed Stochastic Gradient Descent

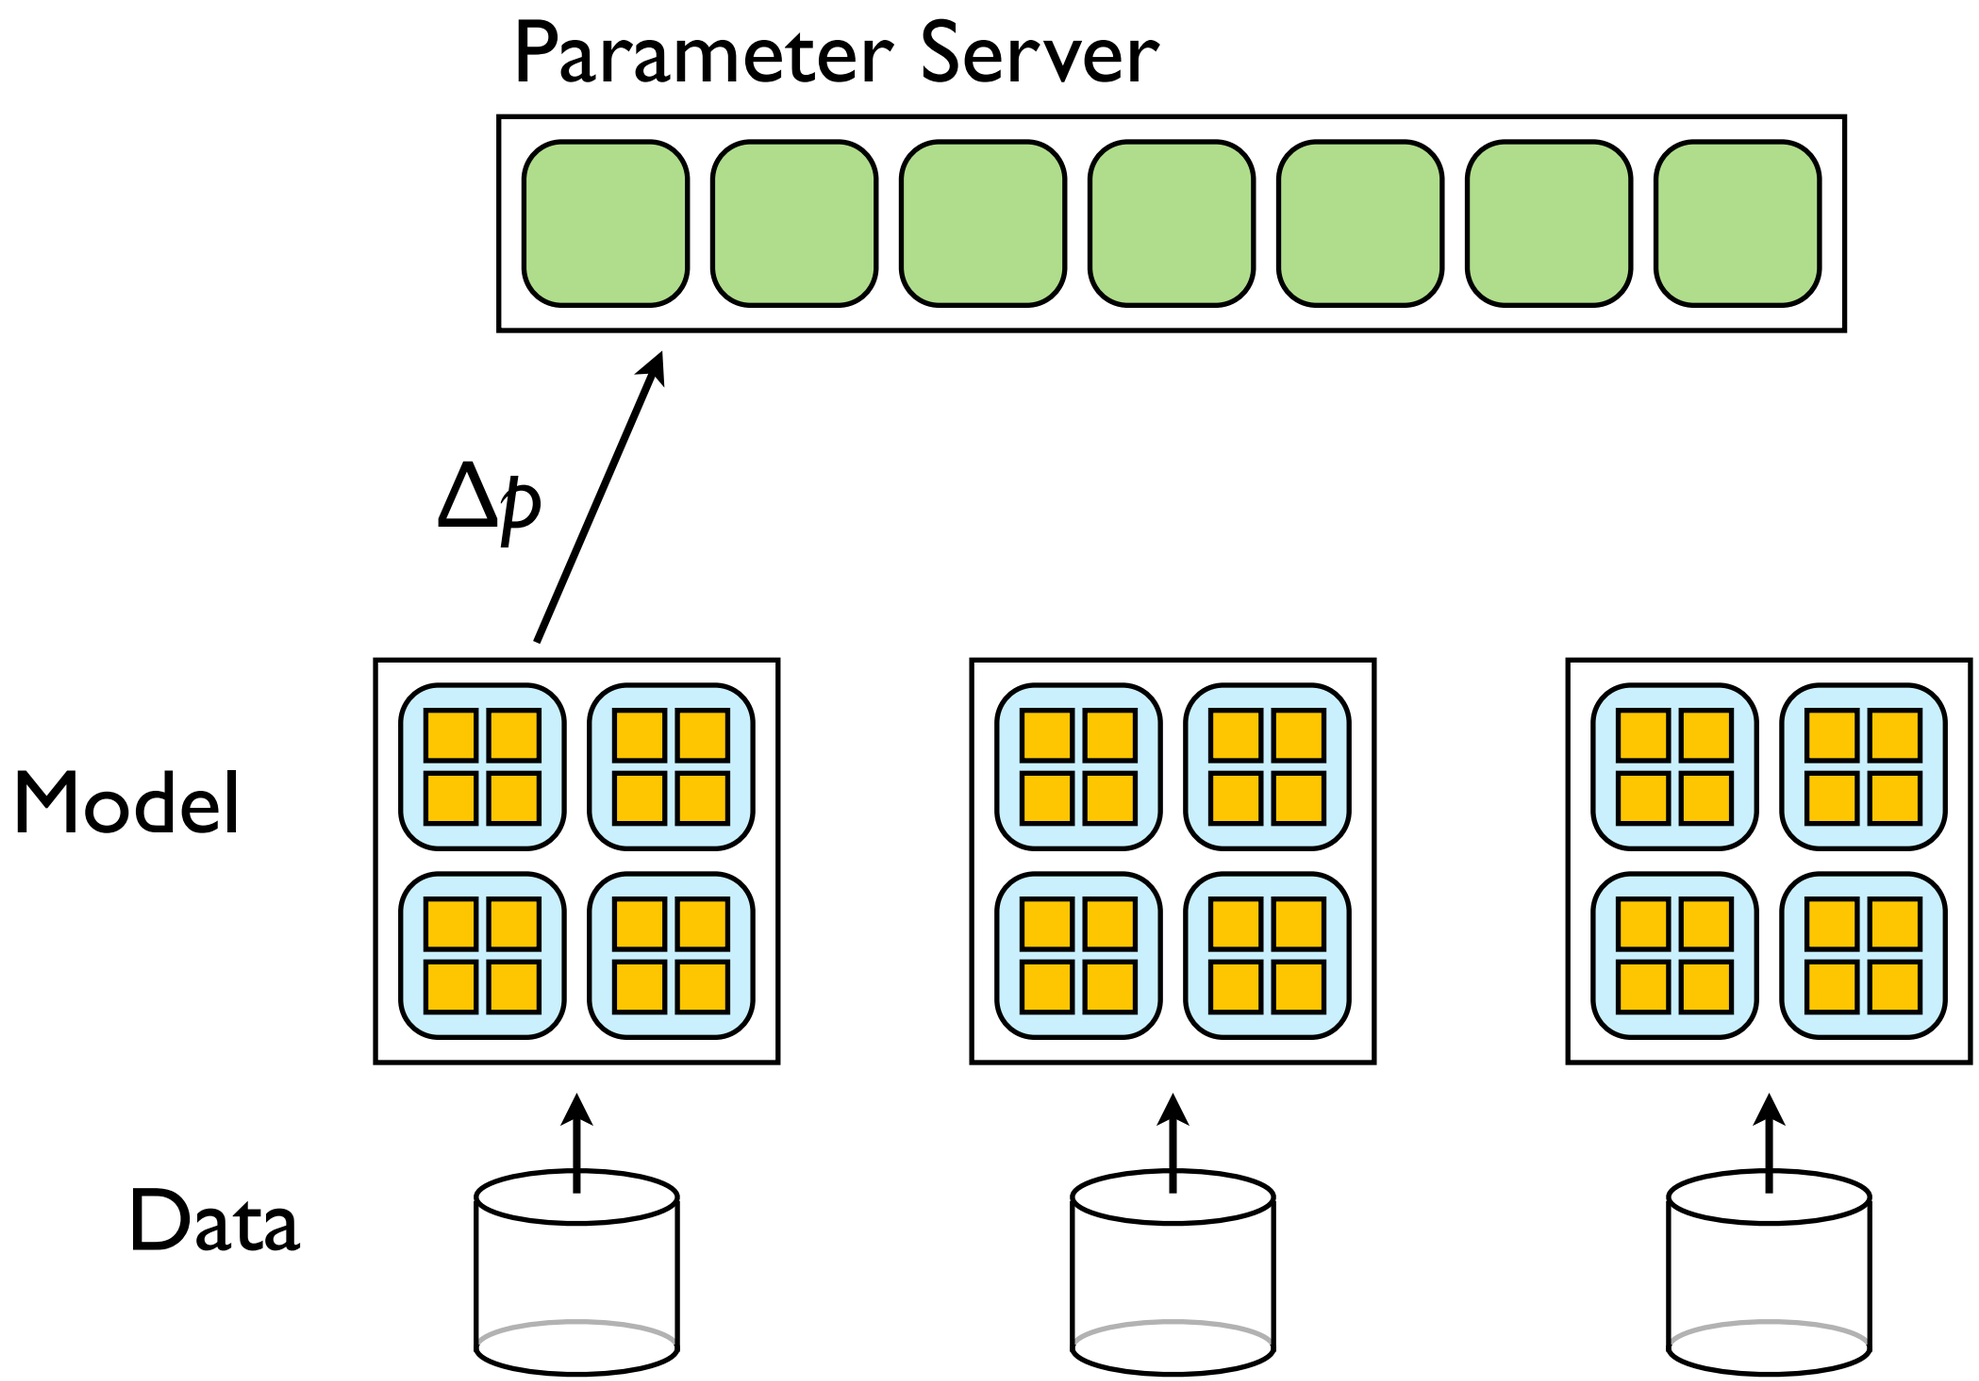

Asynchronous Distributed Stochastic Gradient Descent

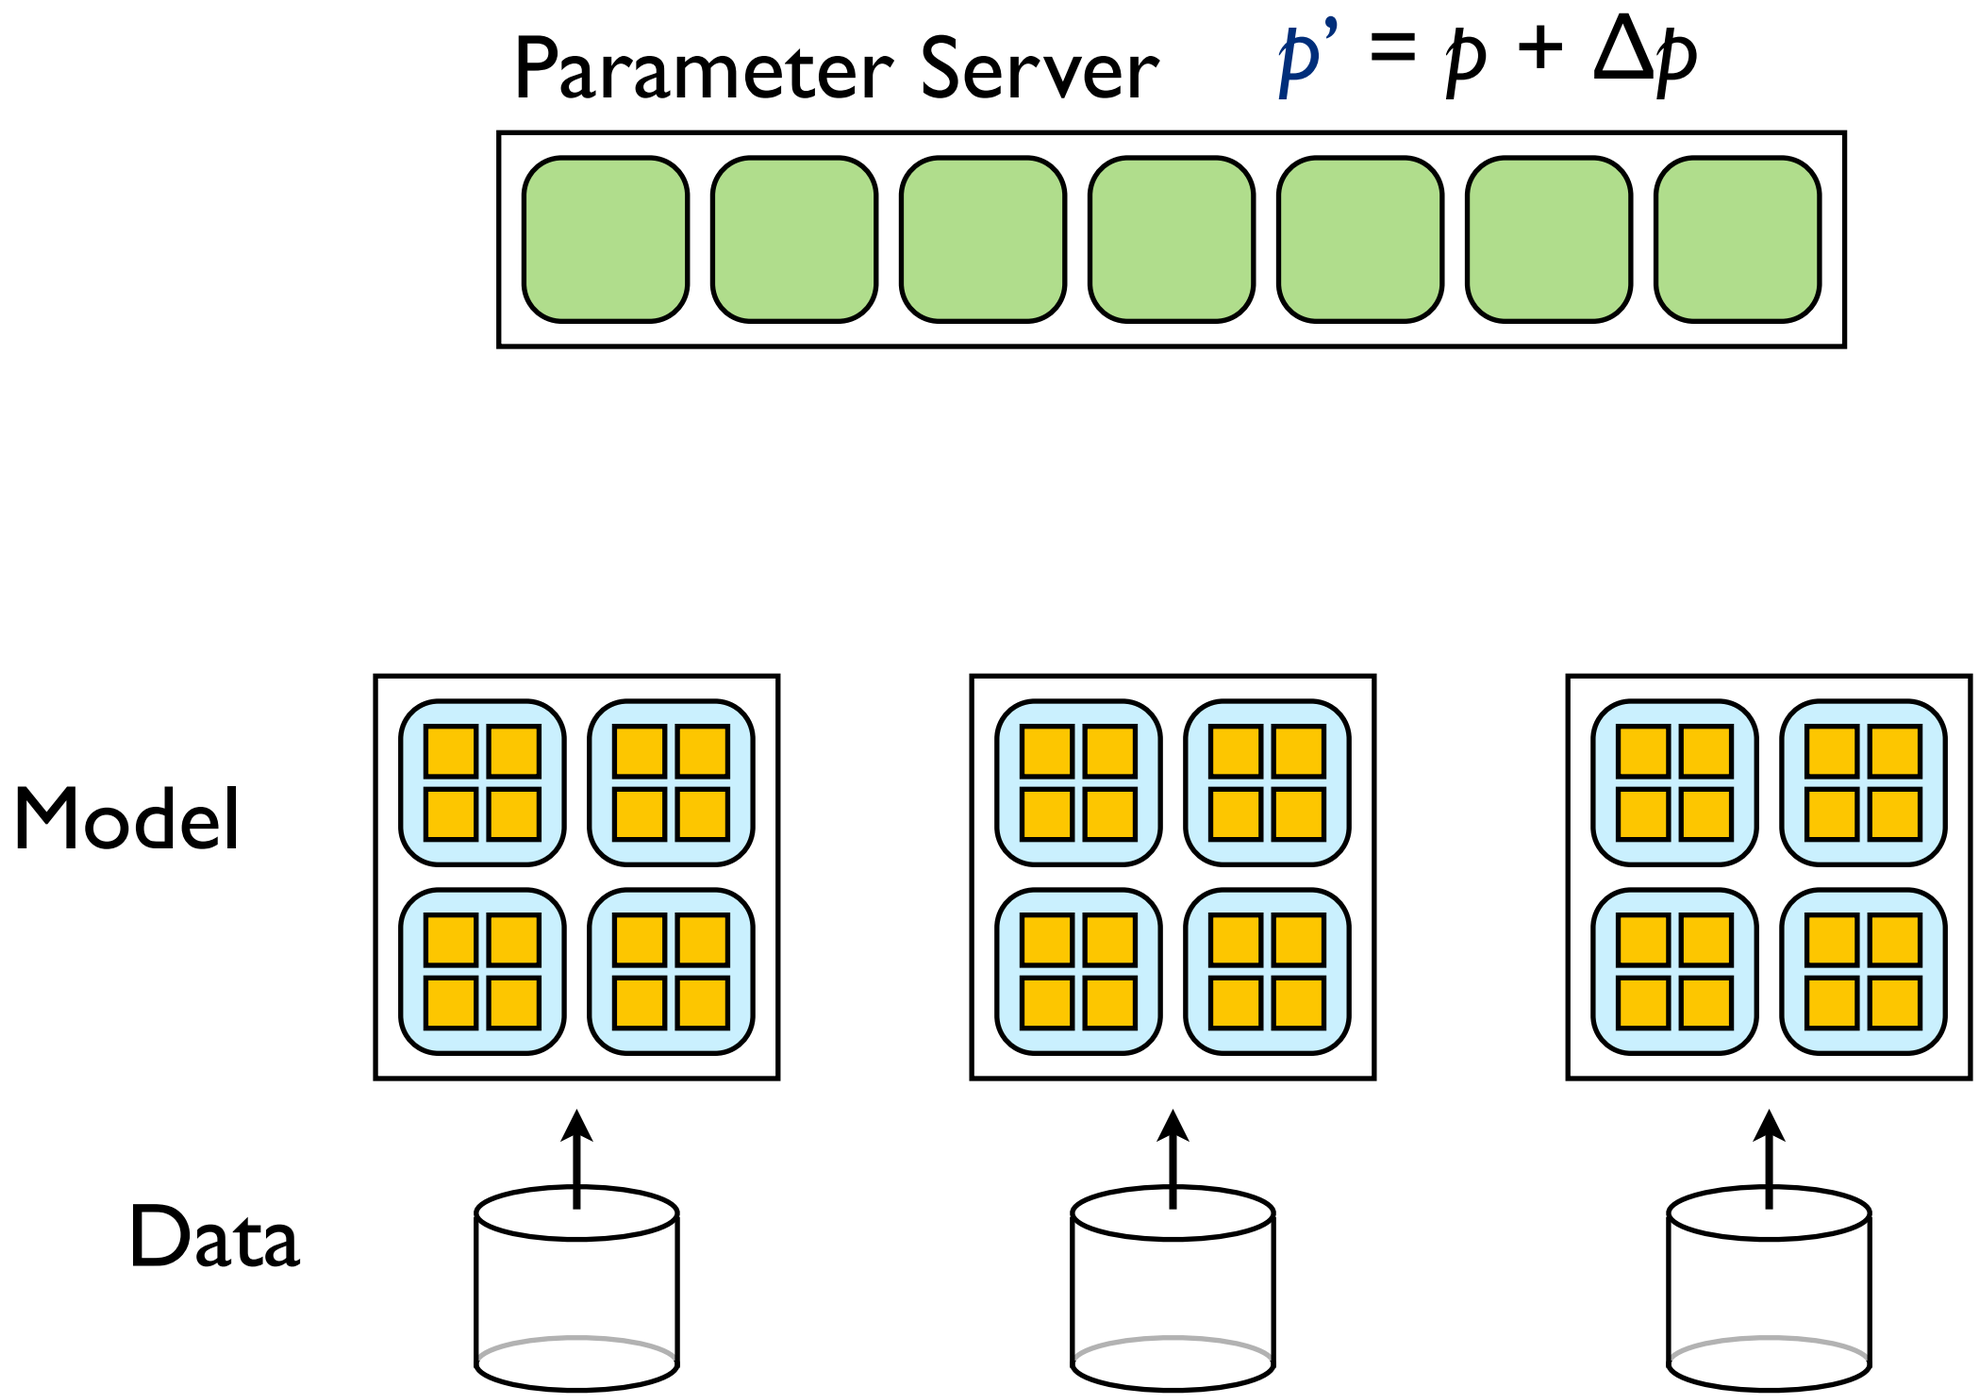

Asynchronous Distributed Stochastic Gradient Descent

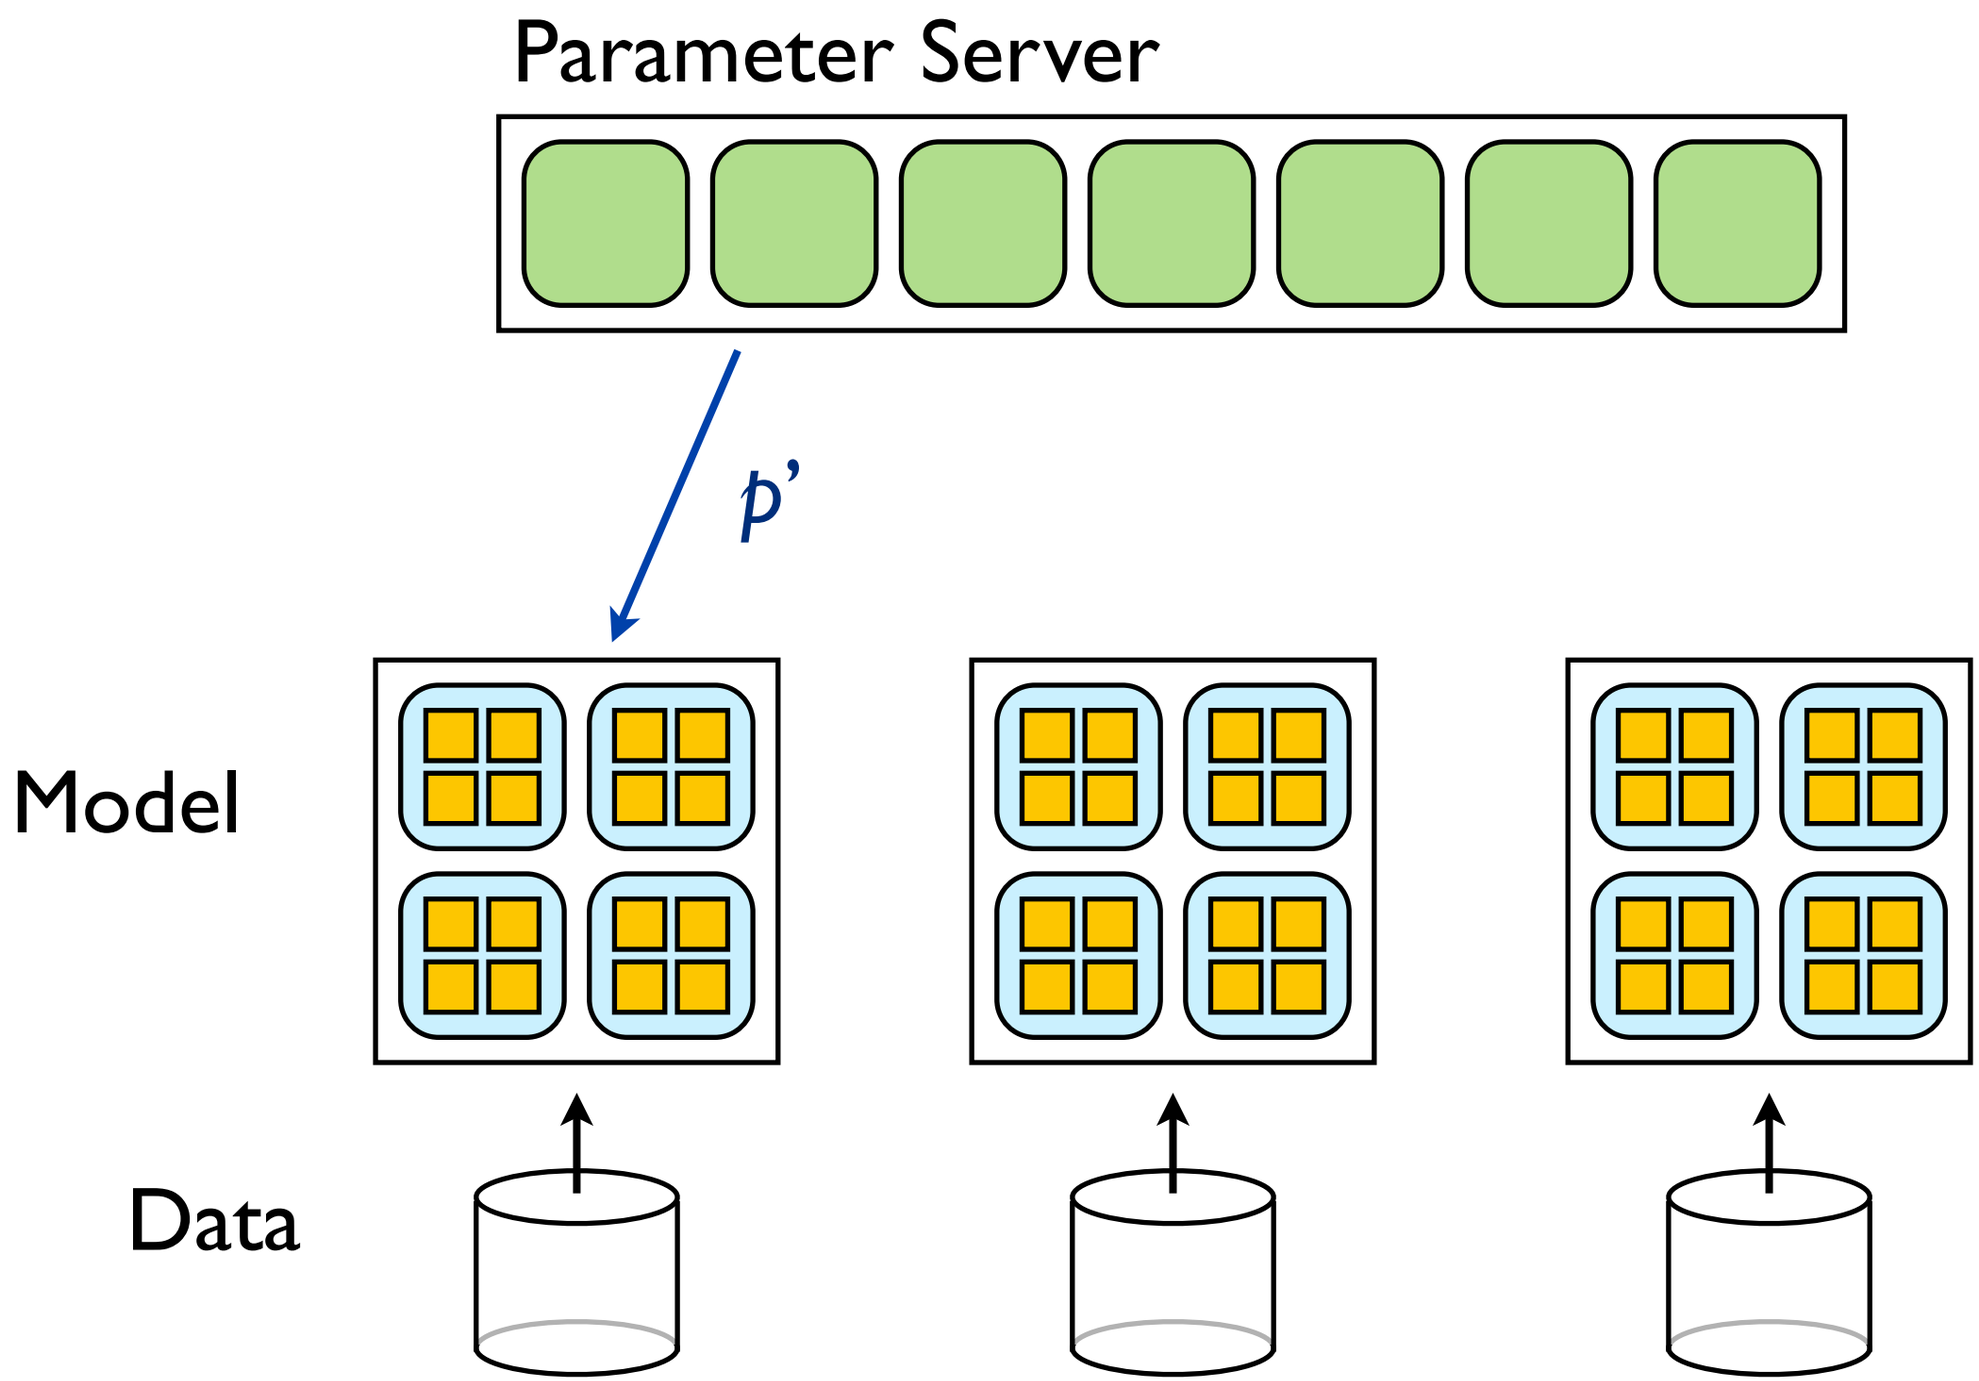

Asynchronous Distributed Stochastic Gradient Descent

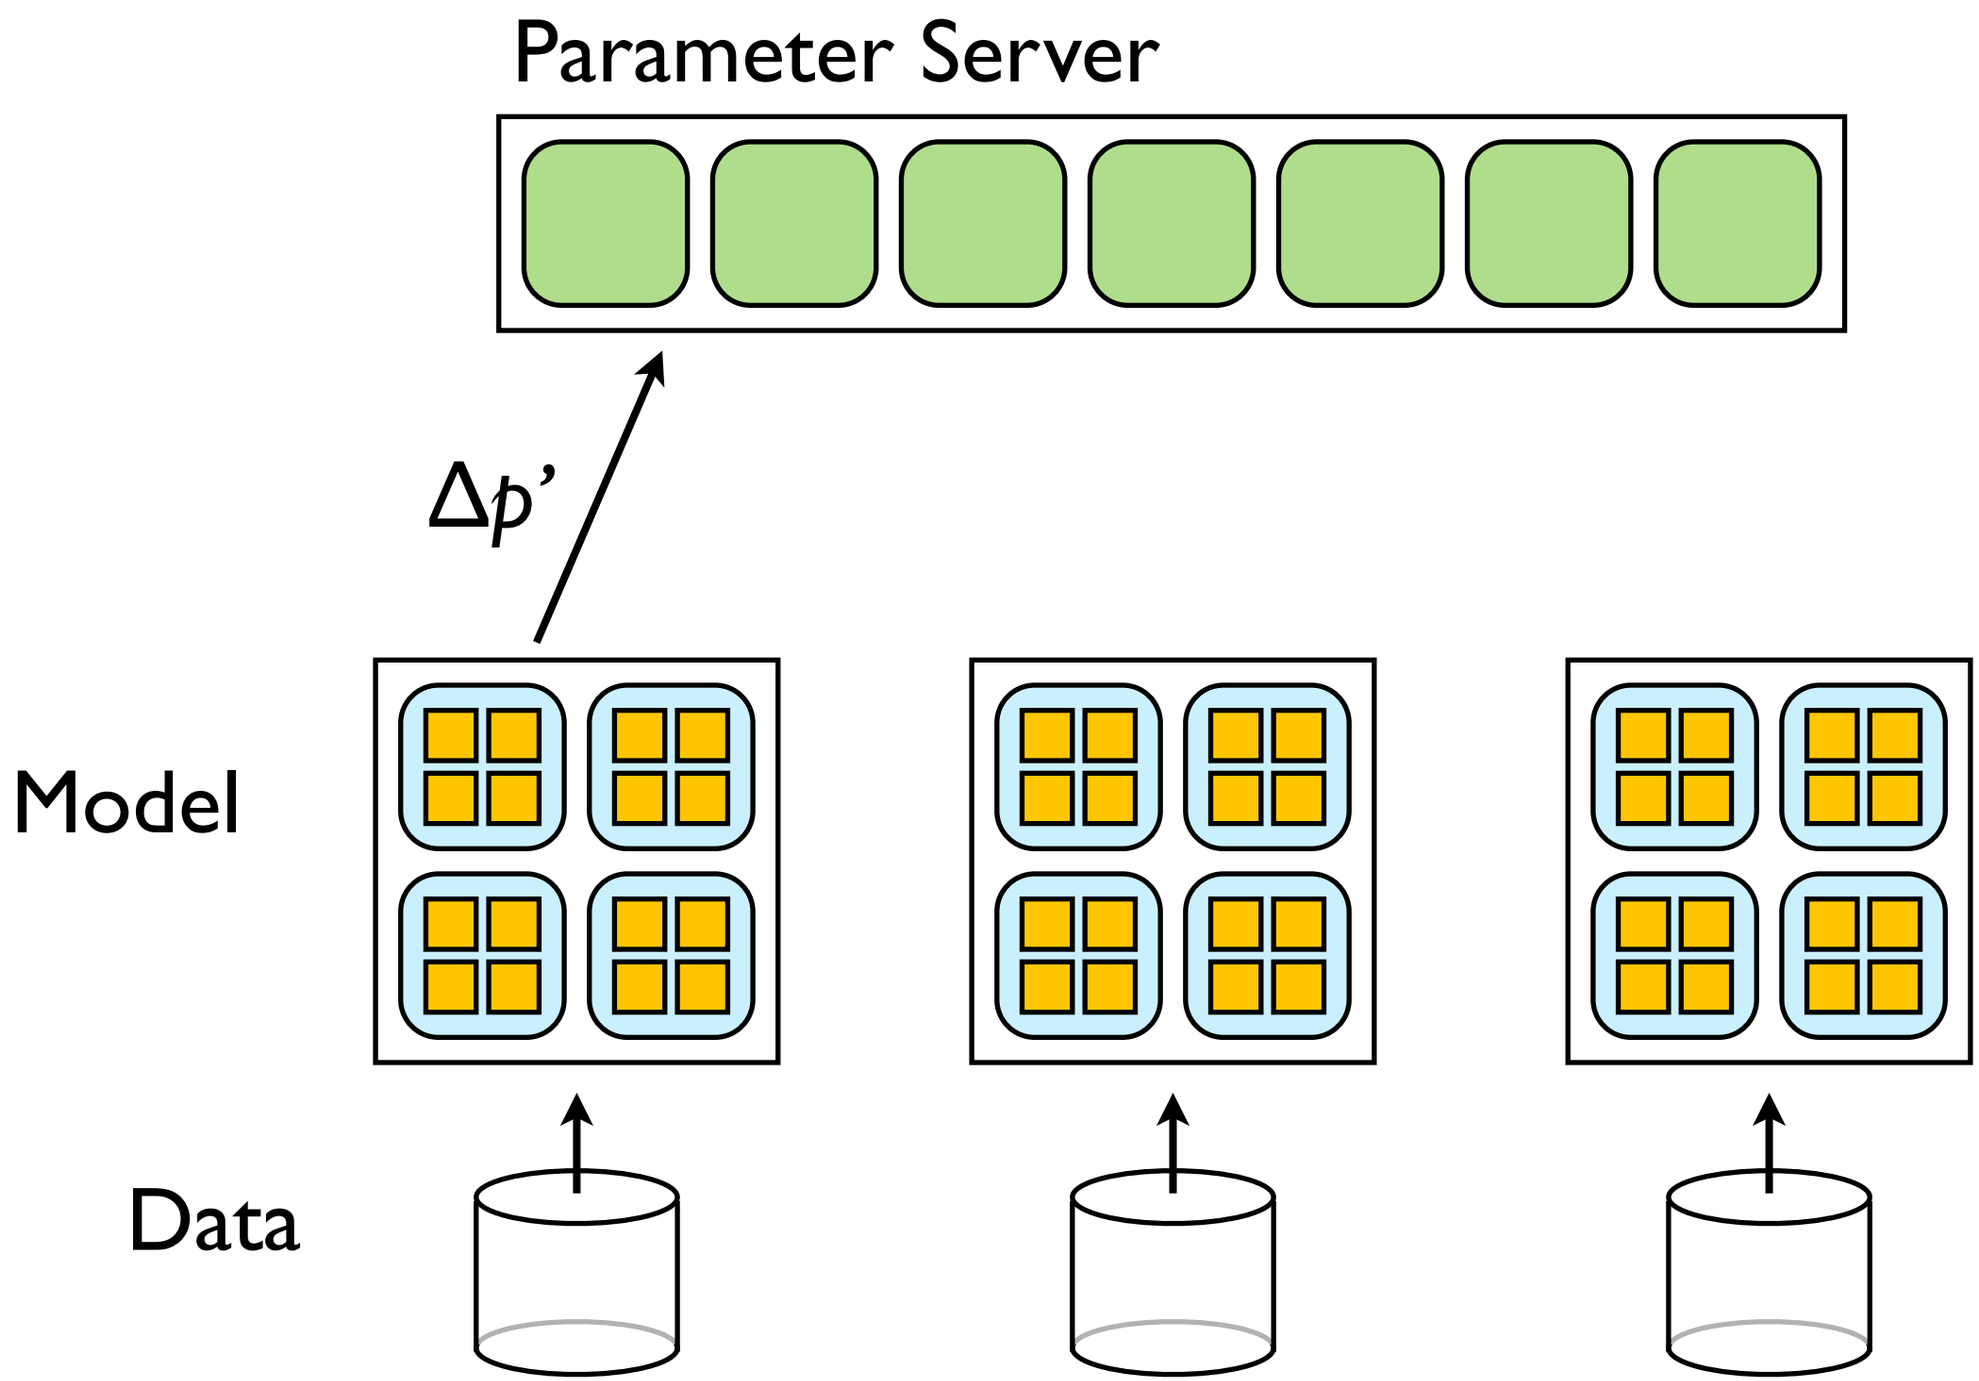

Asynchronous Distributed Stochastic Gradient Descent

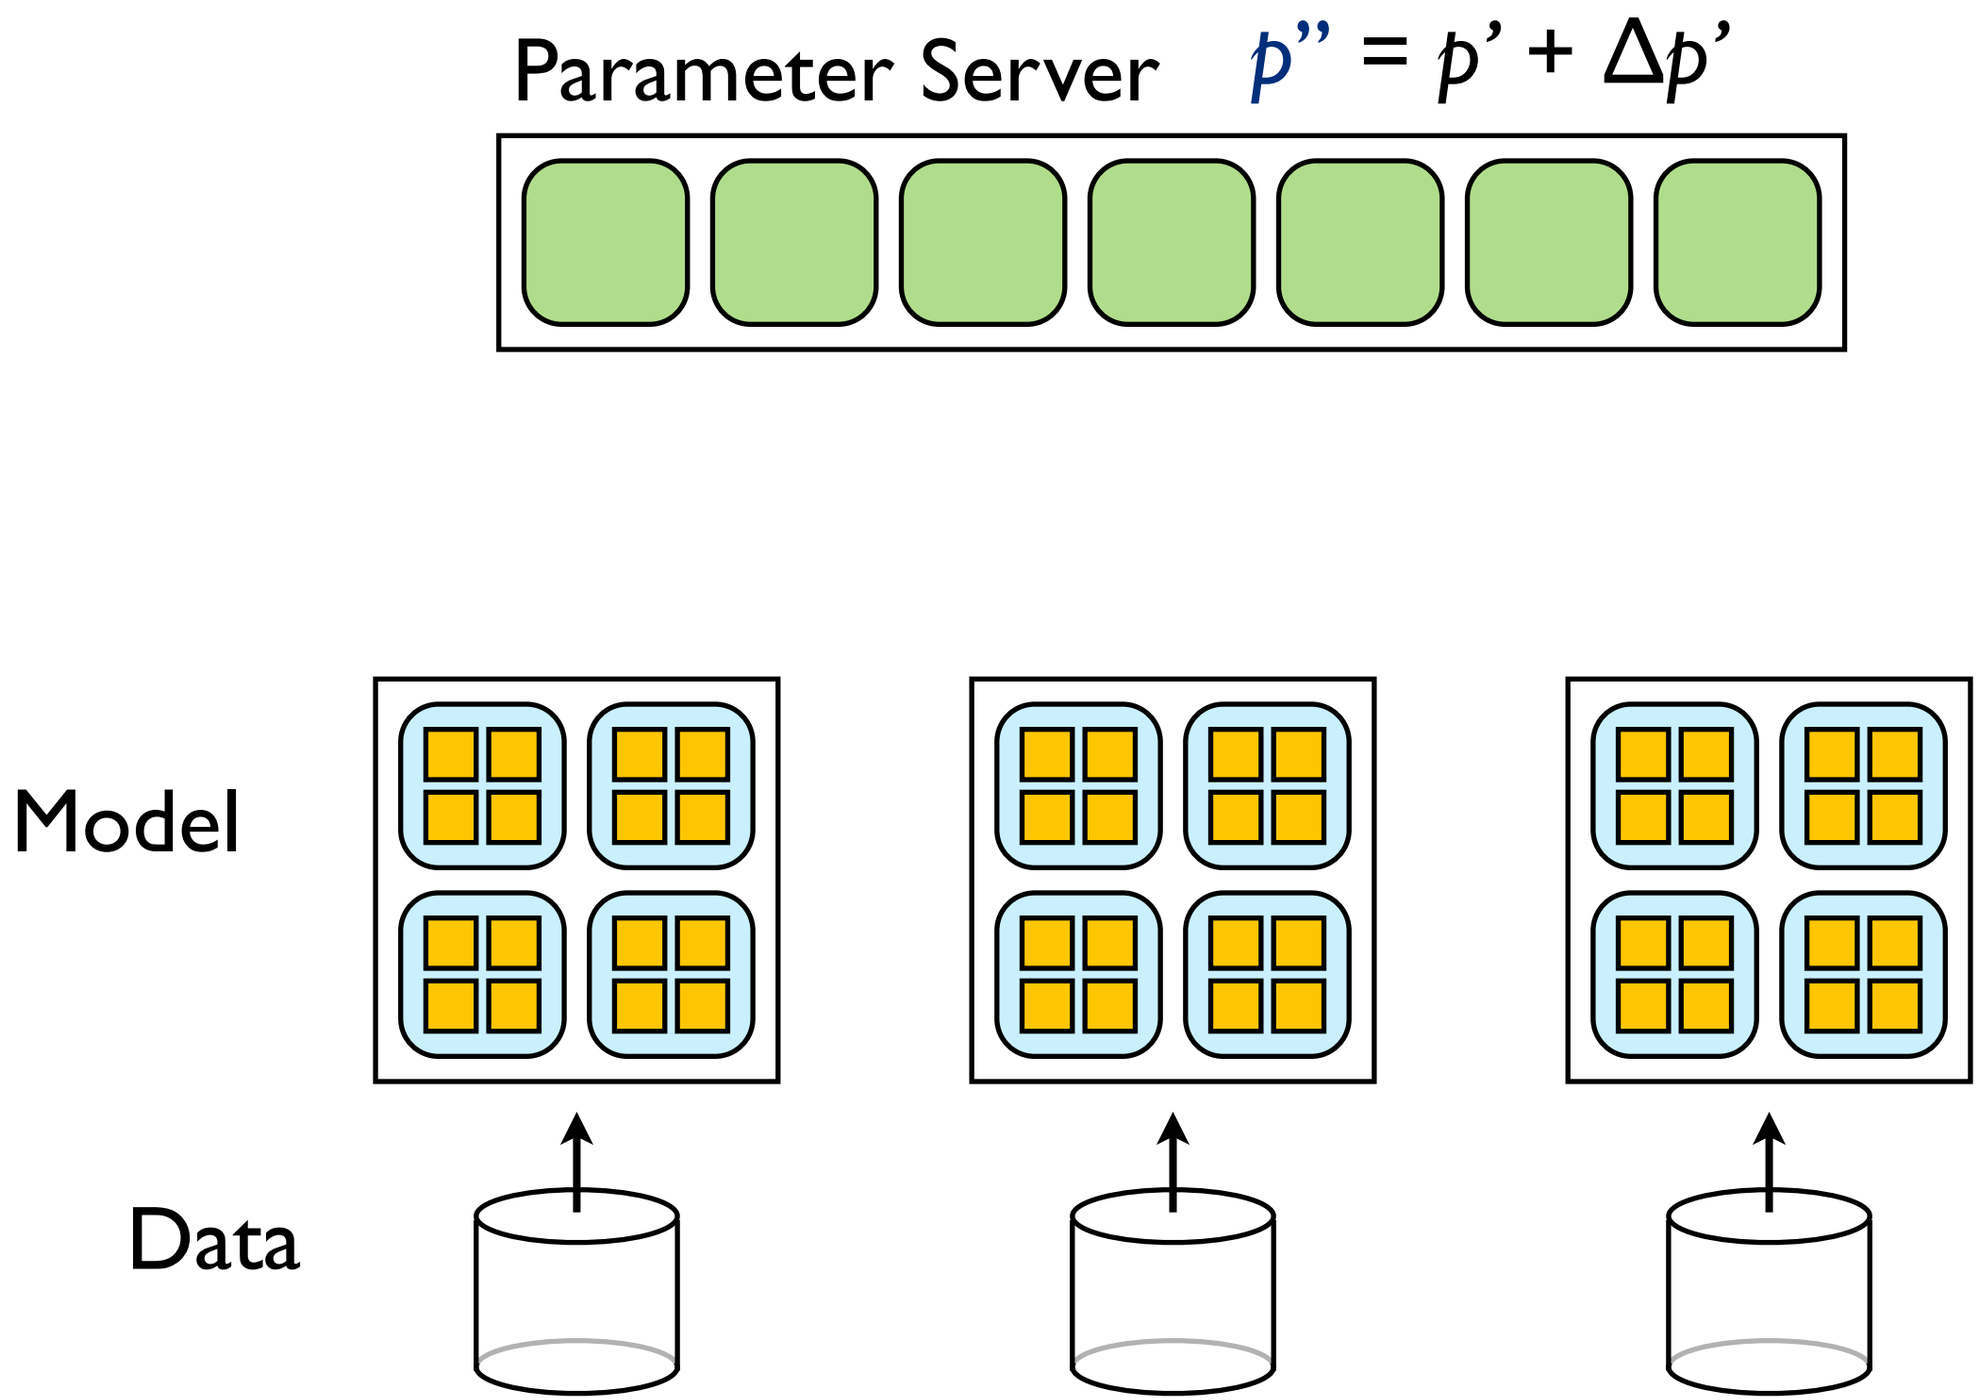

Asynchronous Distributed Stochastic Gradient Descent

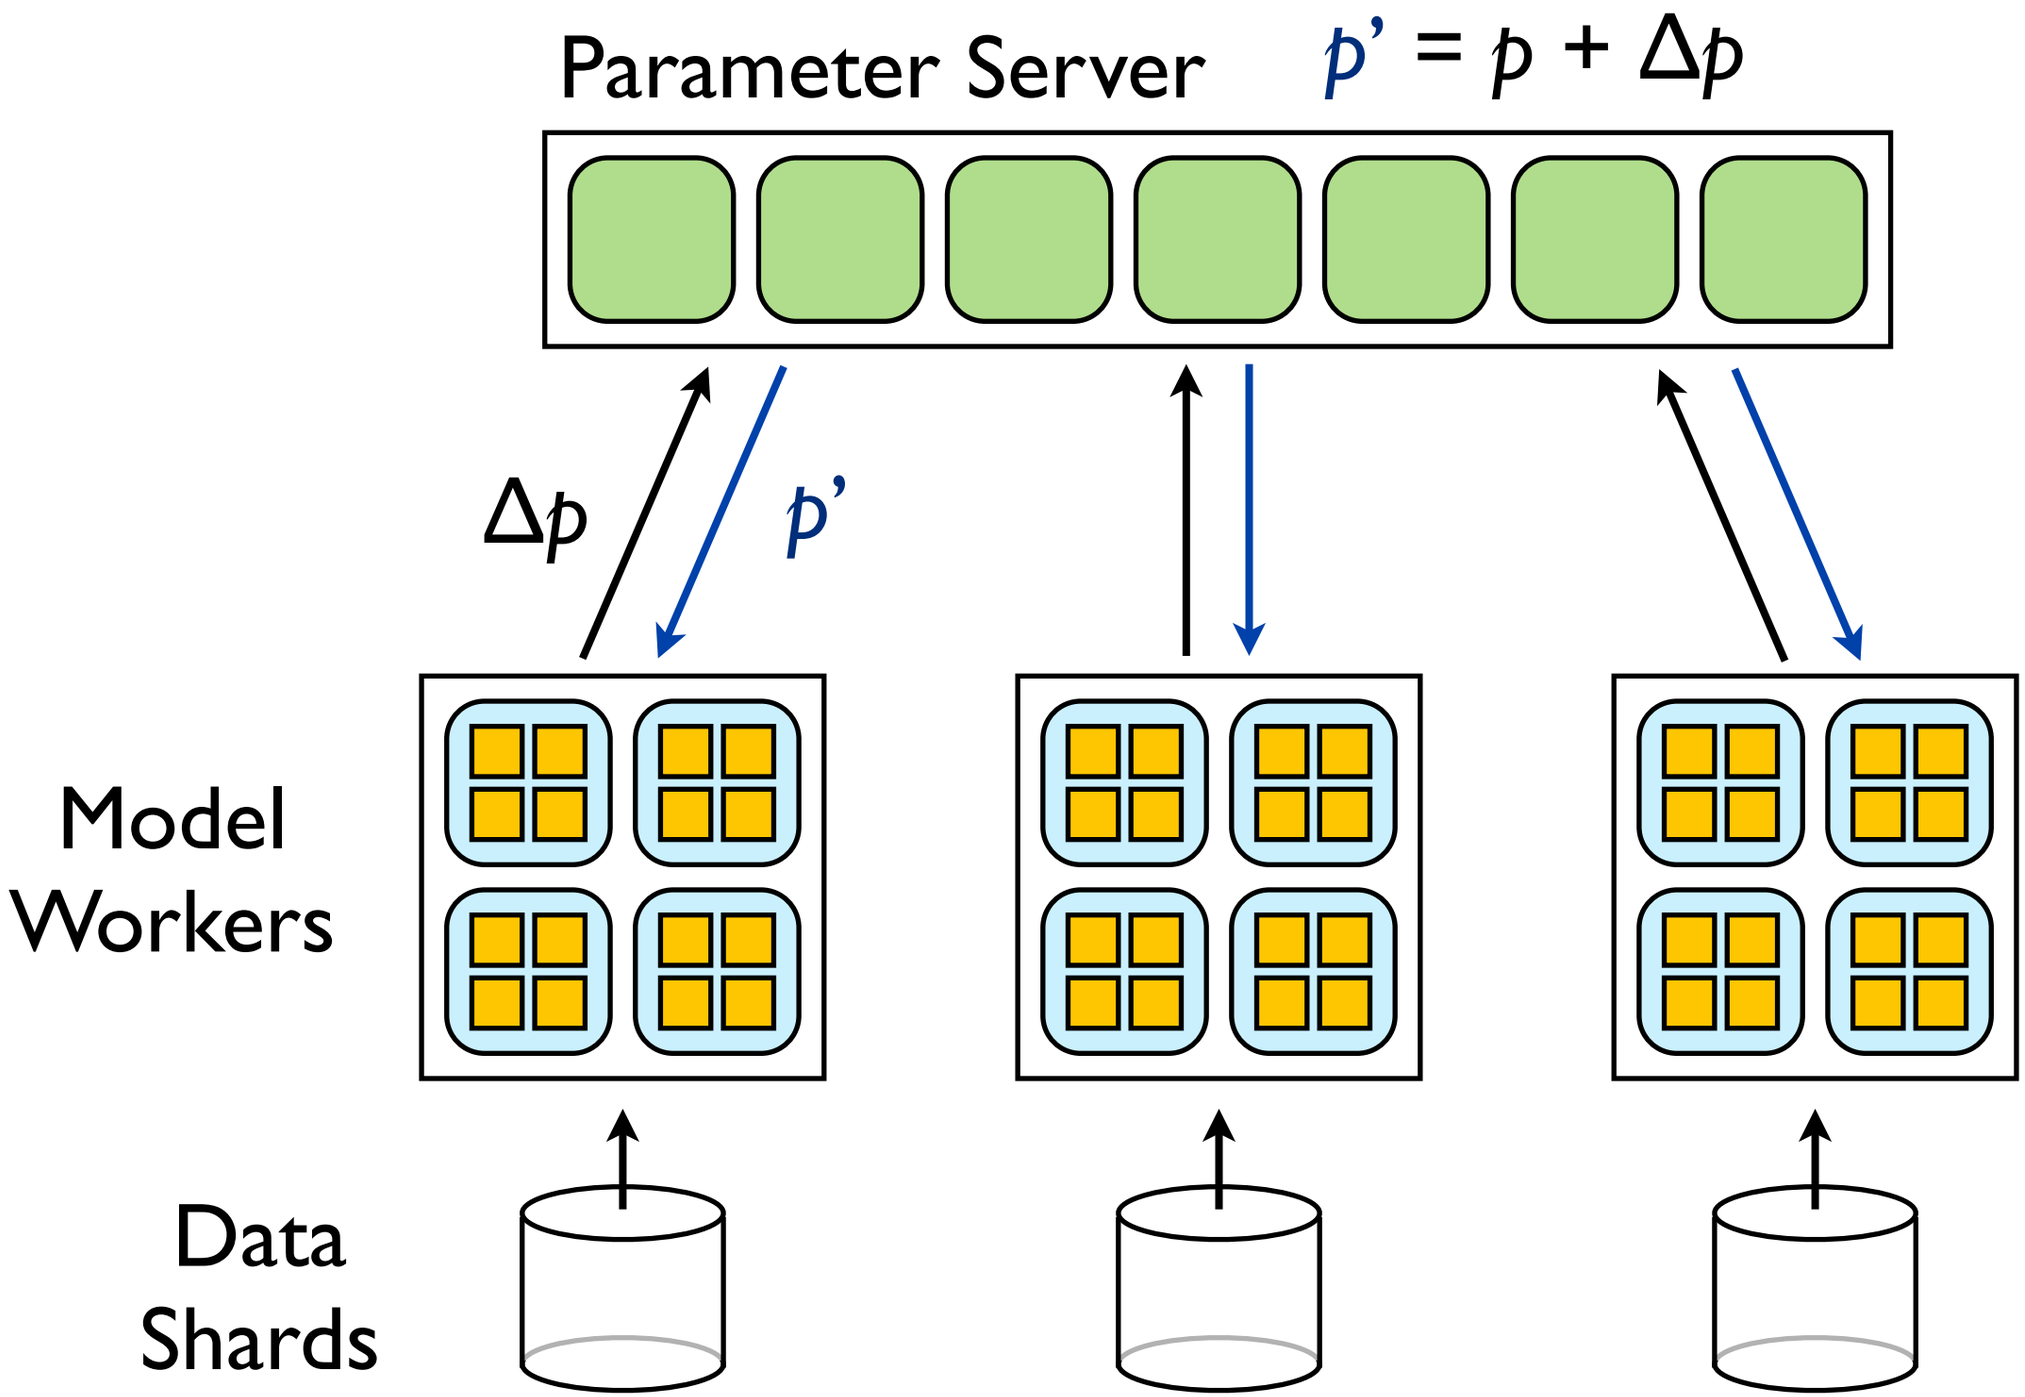

Deep Learning Design Tradeoffs

- Lots of tradeoffs can be made to improve performance. Which ones are possible without hurting learning performance too much?

- For example:
  - Use lower precision arithmetic
  - Send 1 or 2 bits instead of 32 bits across network
  - Drop results from slow partitions

- What's the right hardware for training and deploying these sorts of systems?
  - GPUs? FPGAs? Lossy computational devices?

Applications

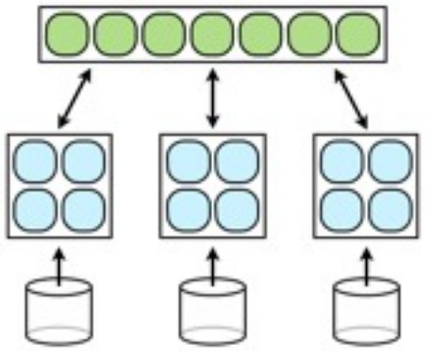

- Acoustic Models for Speech
- Unsupervised Feature Learning for Still Images
- Neural Language Models

Acoustic Modeling for Speech Recognition

8000-label Softmax

One or more hidden layers
of a few thousand nodes each.

11 Frames of 40-value Log Energy Power
Spectra and the label for central frame

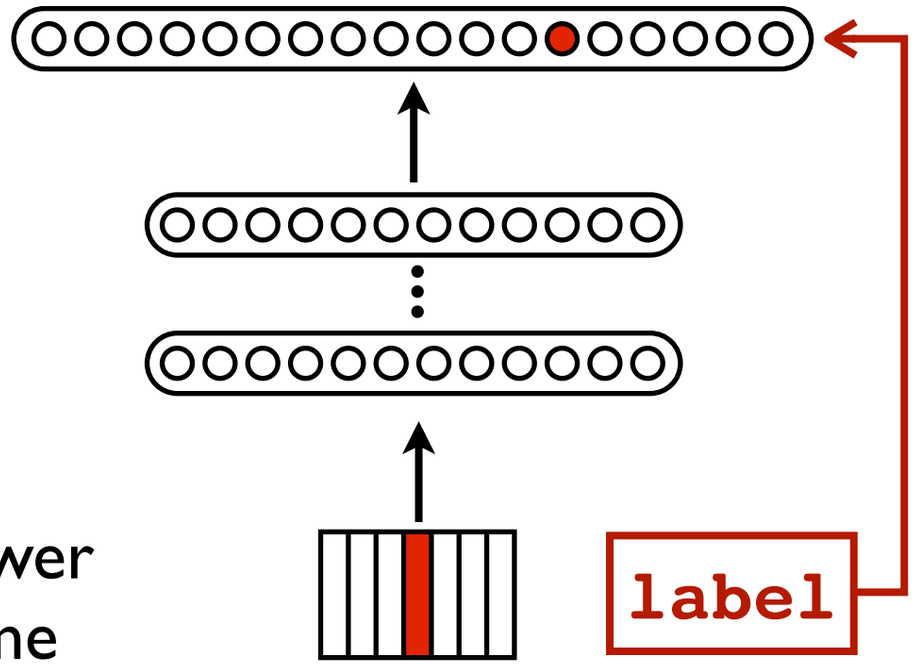

Trained in <5 days on cluster of 800 machines

Acoustic Modeling for Speech Recognition

8000-label Softmax

One or more hidden layers
of a few thousand nodes each.

11 Frames of 40-value Log Energy Power
Spectra and the label for central frame

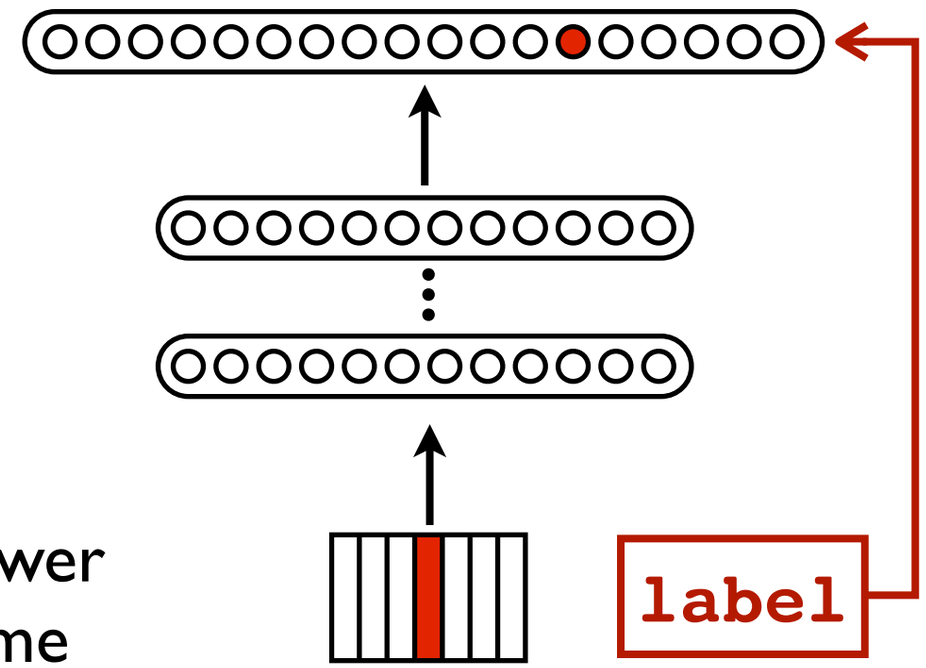

Trained in <5 days on cluster of 800 machines

Major reduction in Word Error Rate
("equivalent to 20 years of speech research")

Applications

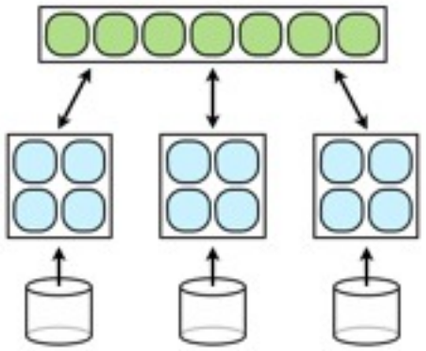

- Acoustic Models for Speech
- Unsupervised Feature Learning for Still Images
- Neural Language Models

Purely Unsupervised Feature Learning in Images

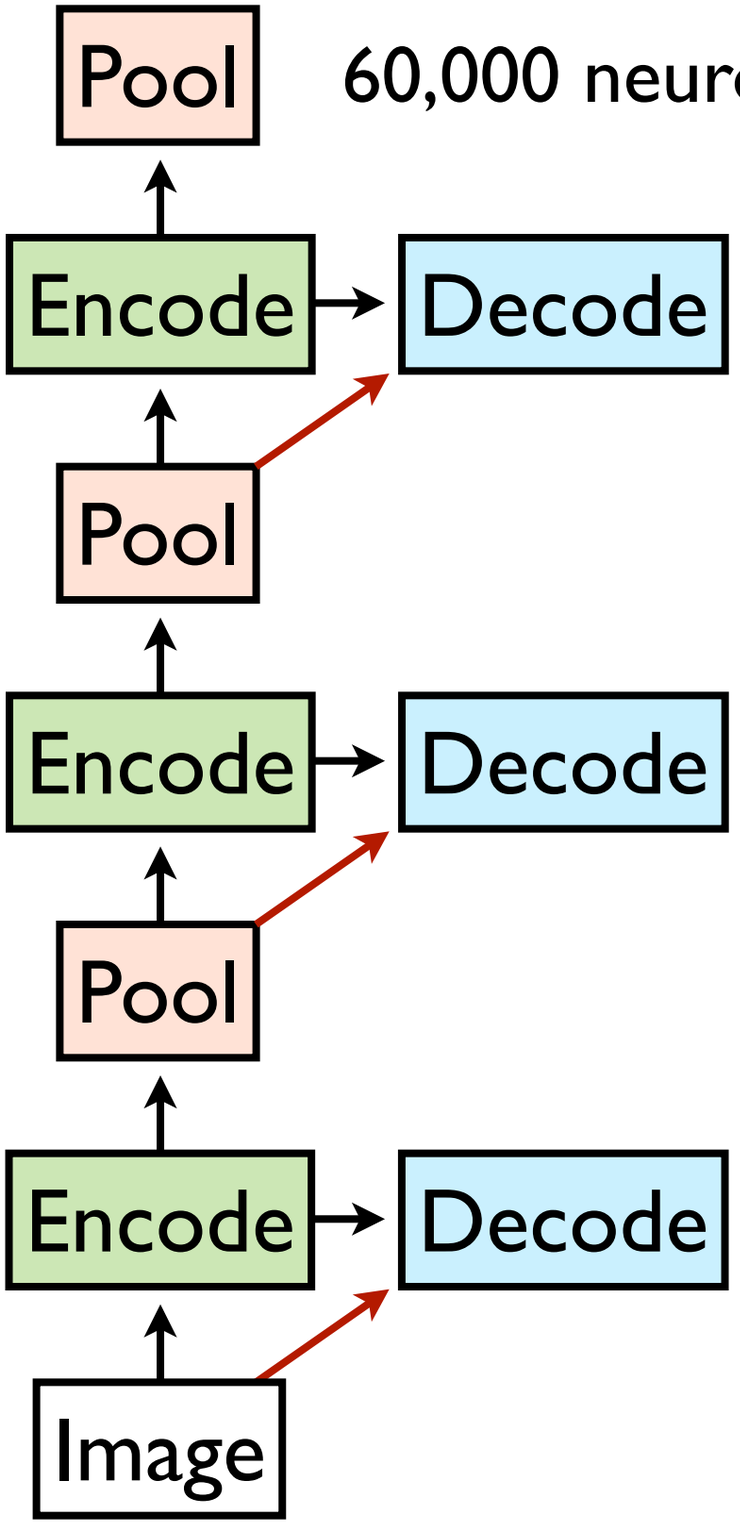

60,000 neurons at top level

- 1.15 billion parameters (50x larger than largest deep network in the literature)
- Trained on 16k cores for 1 week using Async-SGD
- Do **unsupervised** training on one frame from each of 10 million YouTube videos (200x200 pixels)
- No labels!

Details in our ICML paper [Le et al. 2012]

Purely Unsupervised Feature Learning in Images

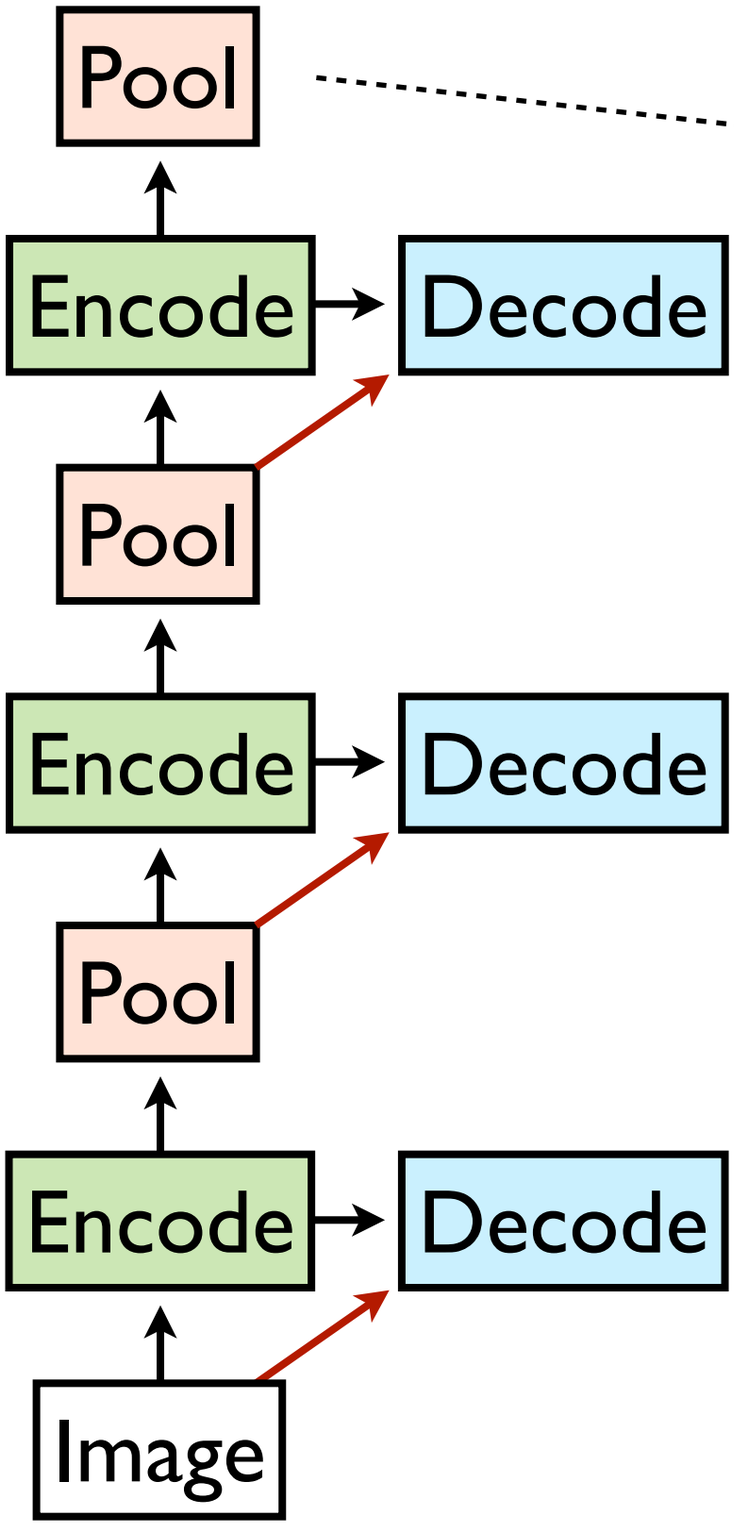

Top level neurons seem to discover high-level concepts. For example, one neuron is a decent face detector:

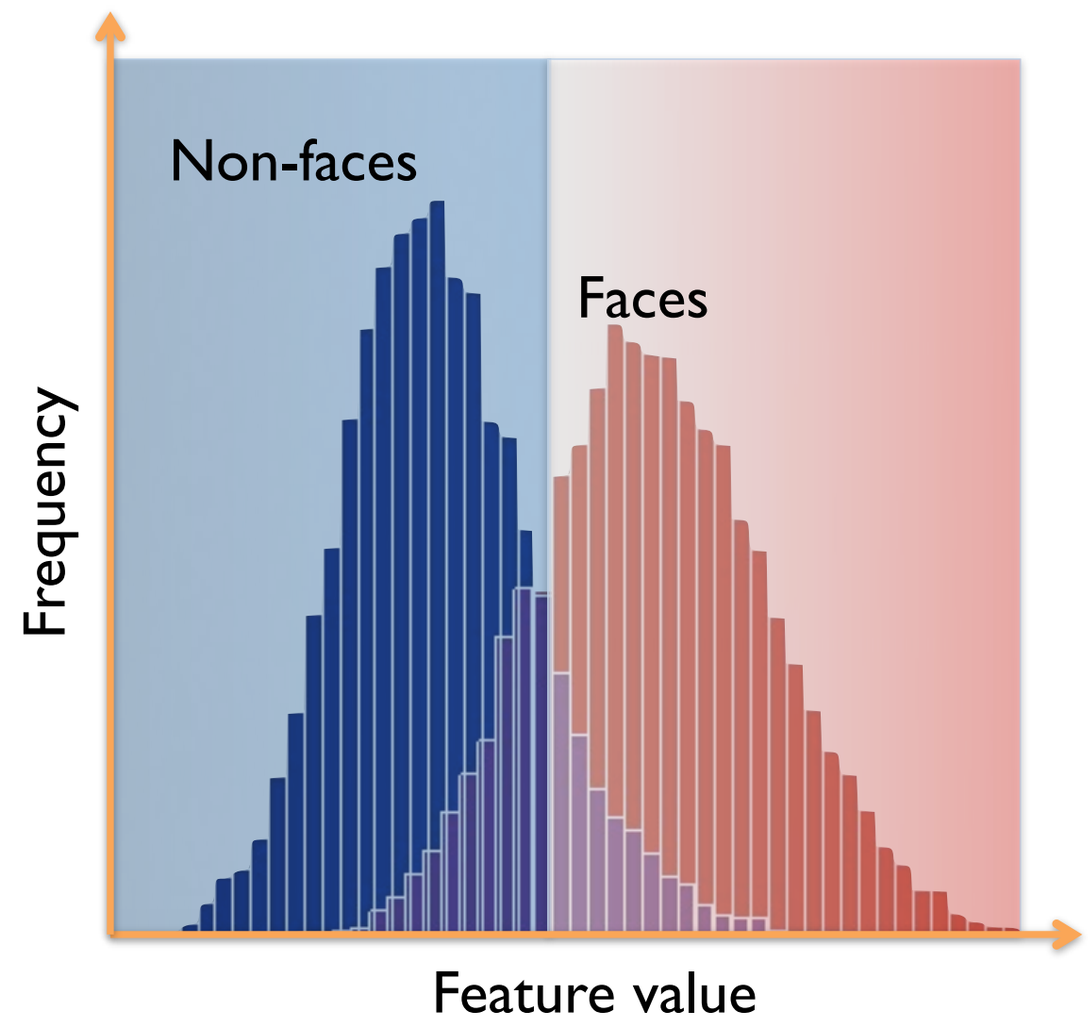

Purely Unsupervised Feature Learning in Images

Most face-selective neuron

Top 48 stimuli from the test set

Purely Unsupervised Feature Learning in Images

Most face-selective neuron

Top 48 stimuli from the test set

Optimal stimulus
by numerical optimization

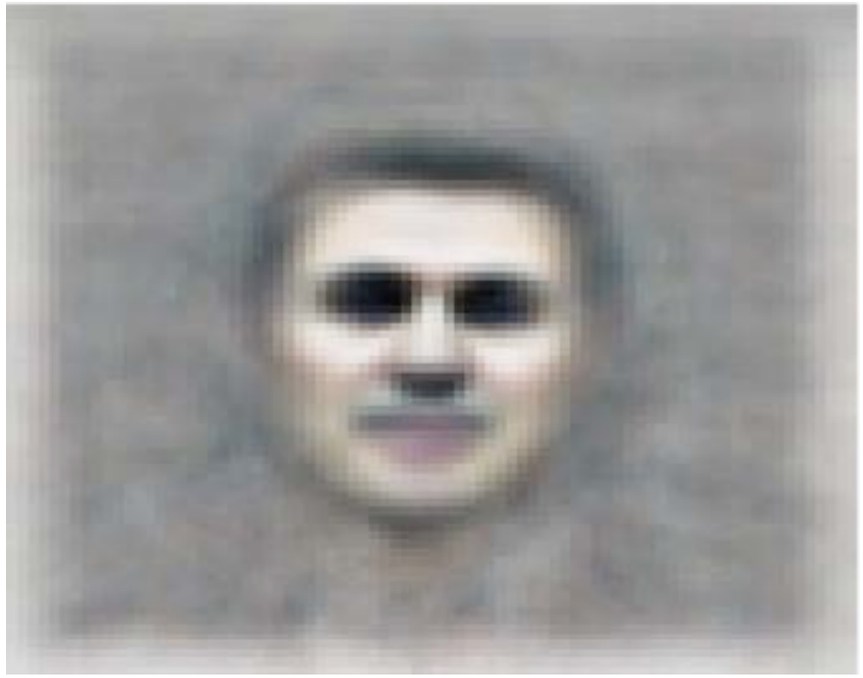

Purely Unsupervised Feature Learning in Images

It is YouTube... We also have a cat neuron!

Top stimuli from the test set

Purely Unsupervised Feature Learning in Images

It is YouTube... We also have a cat neuron!

Top stimuli from the test set

Optimal stimulus

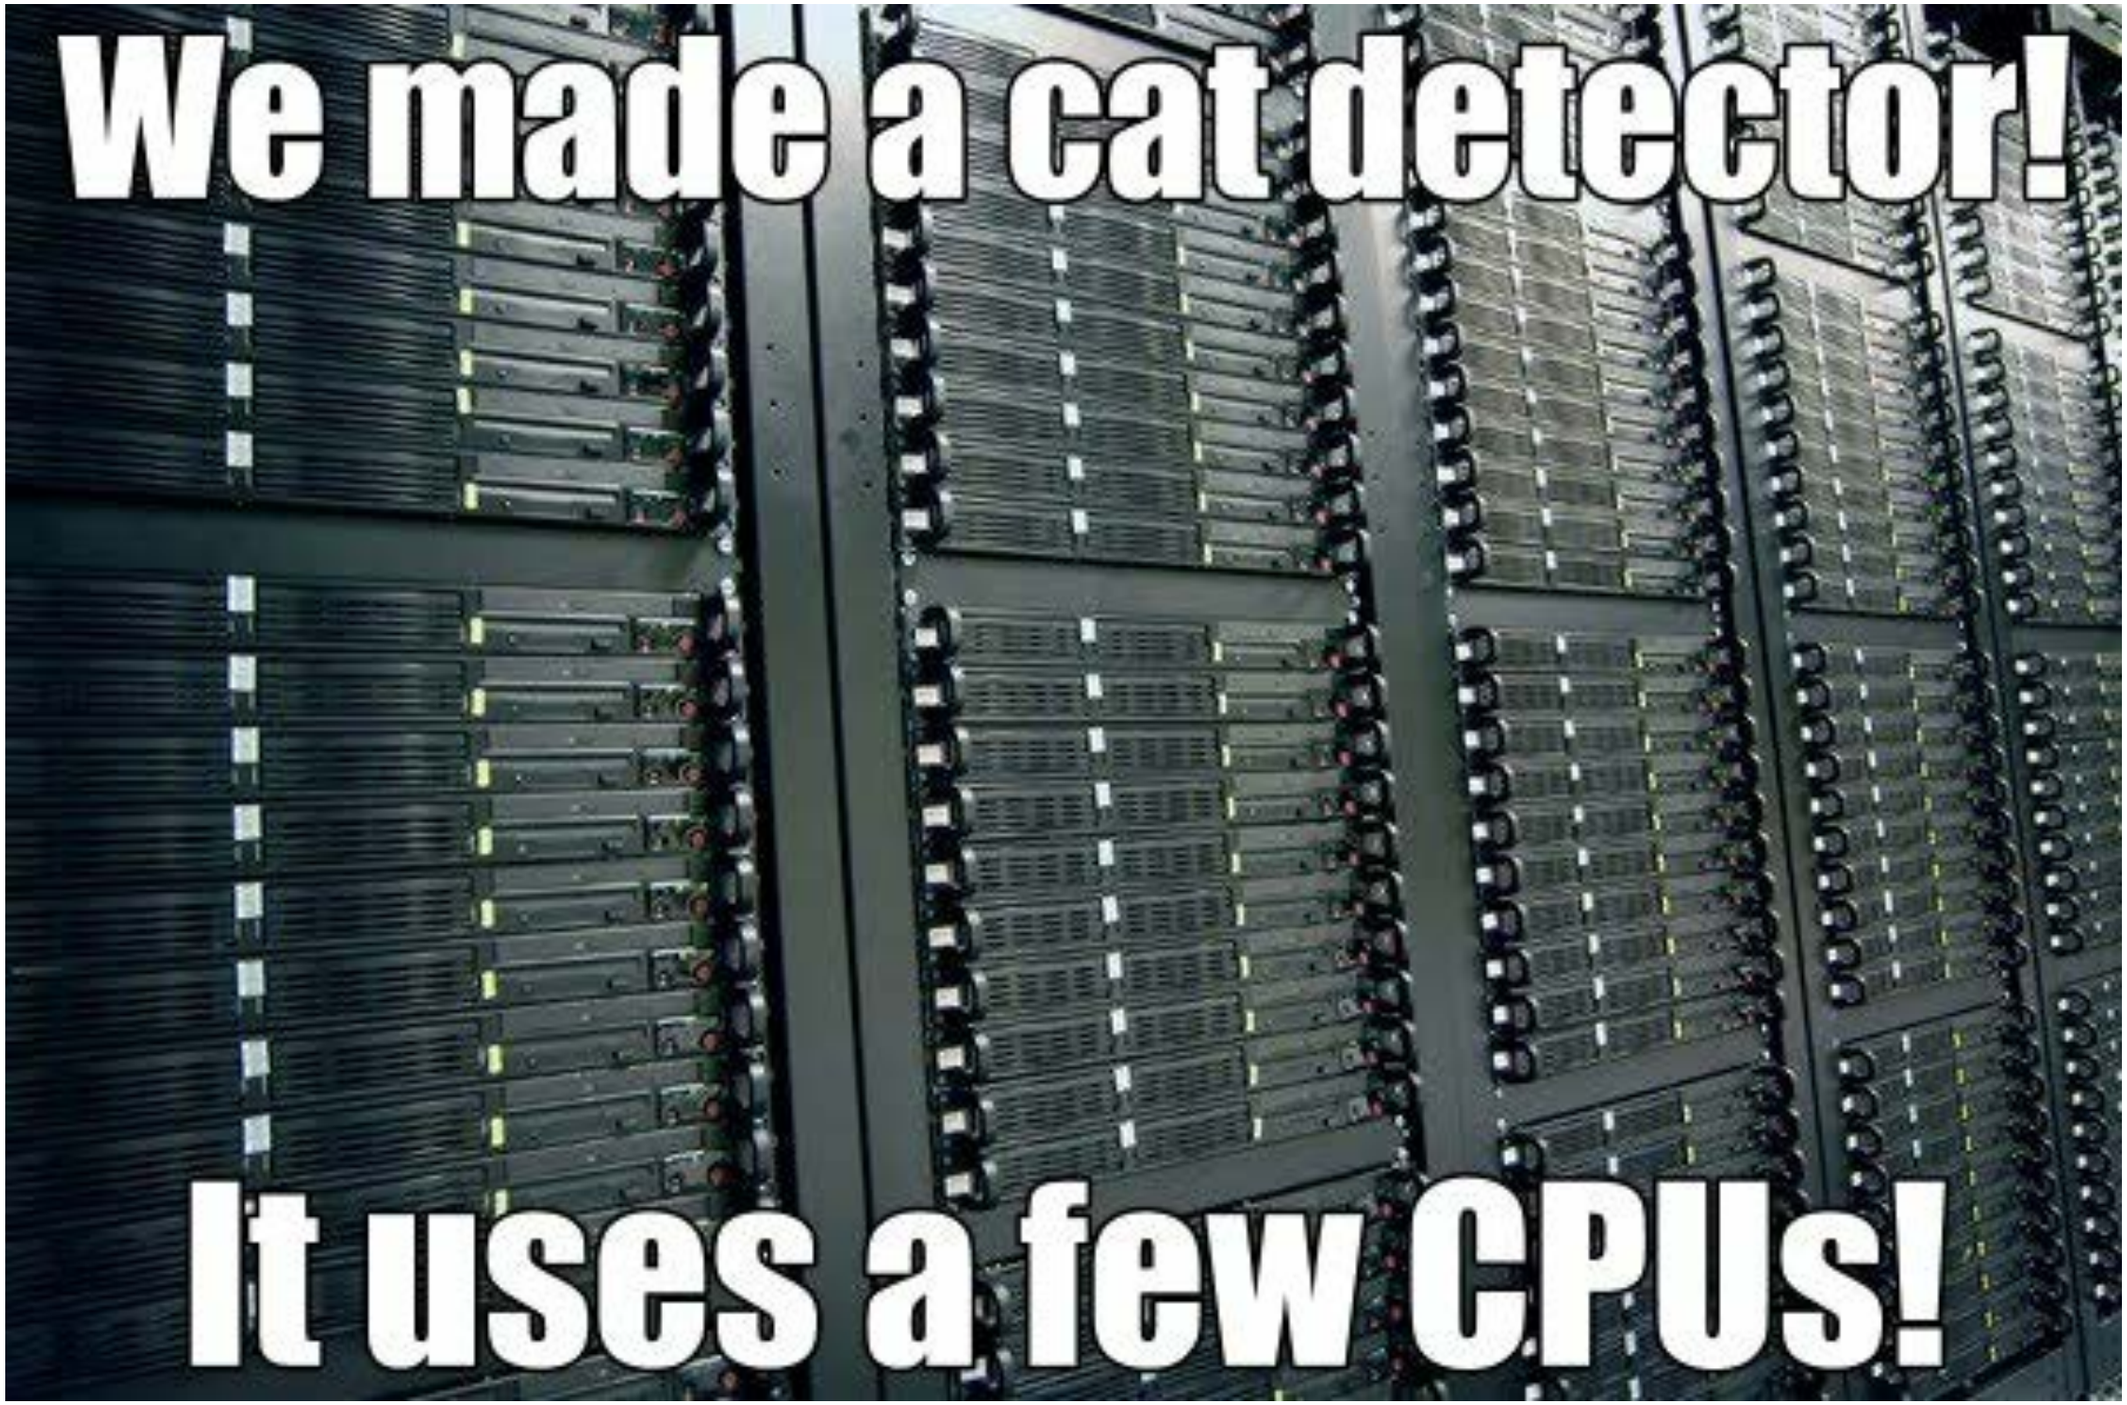
We made a cat detector!
It uses a few CPUs!

Semi-supervised Feature Learning in Images

Are the higher-level representations learned by unsupervised training a useful starting point for supervised training?

We do have *some* labeled data, so let's fine tune this same network for a challenging image classification task.

Semi-supervised Feature Learning in Images

Are the higher-level representations learned by unsupervised training a useful starting point for supervised training?

We do have *some* labeled data, so let's fine tune this same network for a challenging image classification task.

ImageNet:

- 16 million images
- ~21,000 categories
- Recurring academic competitions

Aside: 20,000 is a lot of categories....

```
01496331 electric ray, crampfish, numbfish, torpedo
01497118 sawfish
01497413 smalltooth sawfish, Pristis pectinatus
01497738 guitarfish
01498041 stingray
01498406 roughtail stingray, Dasyatis centroura
01498699 butterfly ray
01498989 eagle ray
01499396 spotted eagle ray, spotted ray, Aetobatus narinari
01499732 cownose ray, cow-nosed ray, Rhinoptera bonasus
01500091 manta, manta ray, devilfish
01500476 Atlantic manta, Manta birostris
01500854 devil ray, Mobula hypostoma
01501641 grey skate, gray skate, Raja batis
01501777 little skate, Raja erinacea
01501948 thorny skate, Raja radiata
01502101 barndoor skate, Raja laevis
01503976 dickeybird, dickey-bird, dickybird, dicky-bird
01504179 fledgling, fledgeling
01504344 nestling, baby bird
```

Aside: 20,000 is a lot of categories....

0149 **roughtail stingray**, numbfish, torpedo

0149 us
0149
0149
0149 ra
0149
0149
0149
0149 batus narinari
01499732 cownose ray, cow-nosed ray, Rhinoptera bonasus
01500091 manta, ma fish
01500476 Atlantic rostris
01500854 devil ray
01501641 grey skat
01501777 little sk
01501948 thorny sk
01502101 barndoor
01503976 dickeybir
01504179 fledgling
01504344 nestling, baby bird

manta ray

Semi-supervised Feature Learning in Images

Example top stimuli after fine tuning on ImageNet:

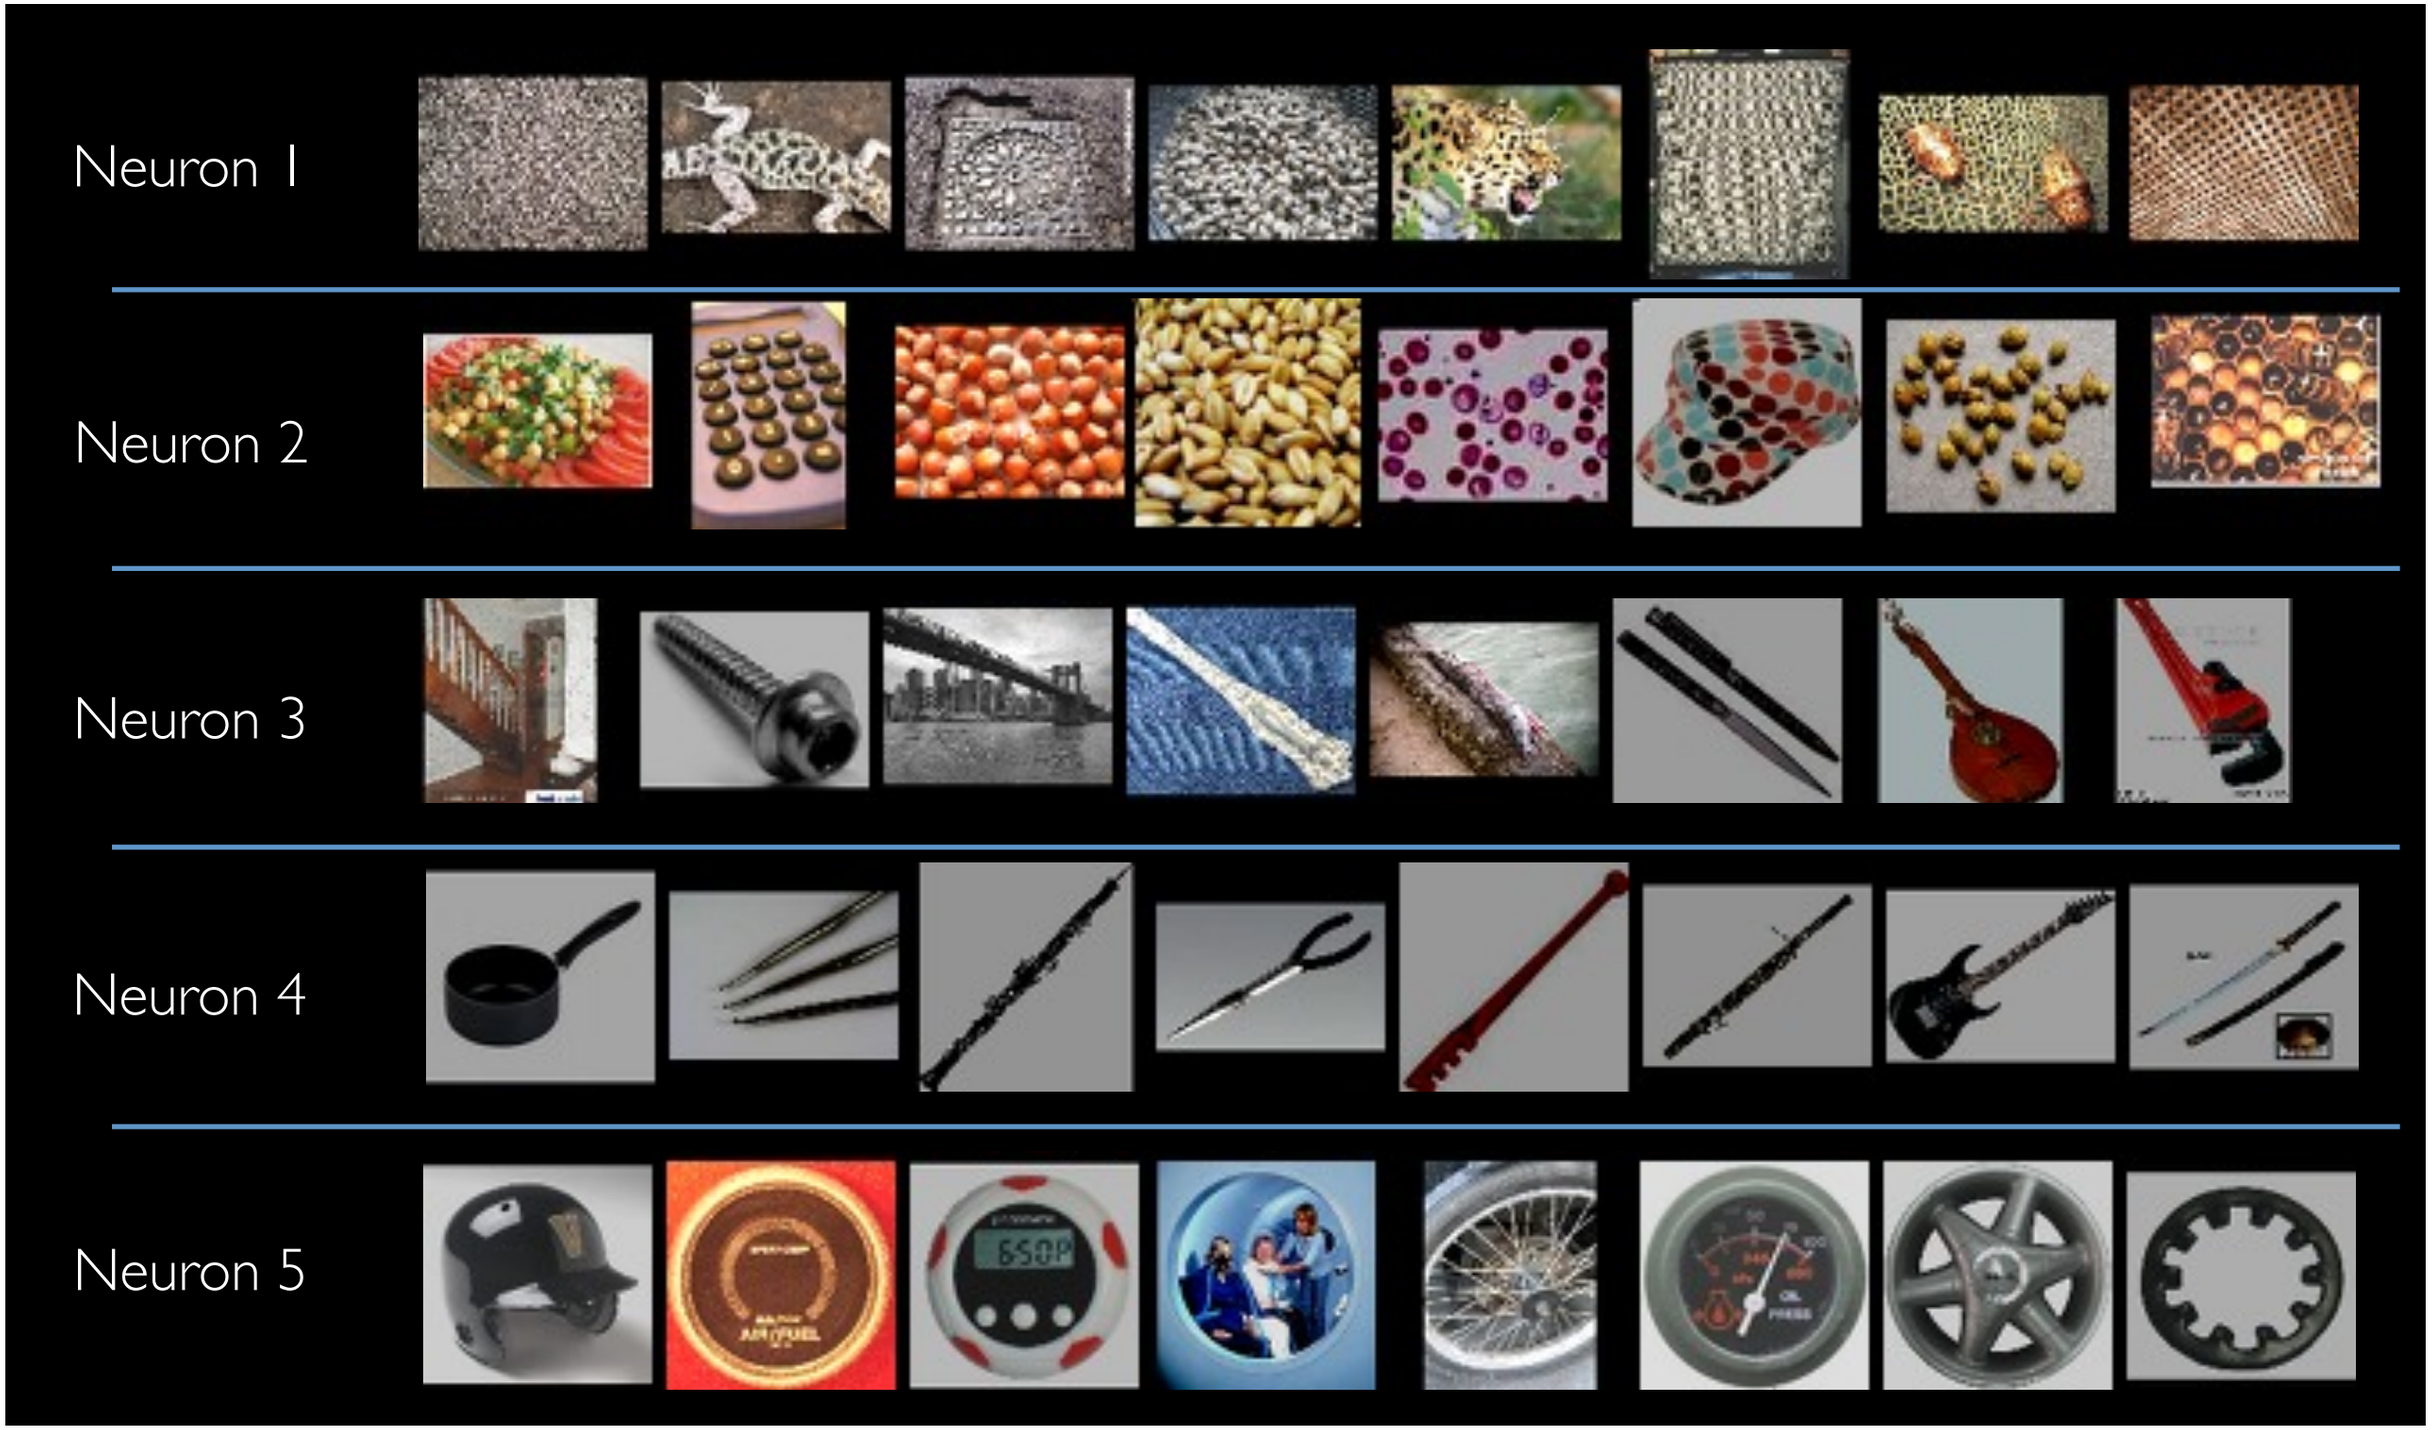

Semi-supervised Feature Learning in Images

Example top stimuli after fine tuning on ImageNet:

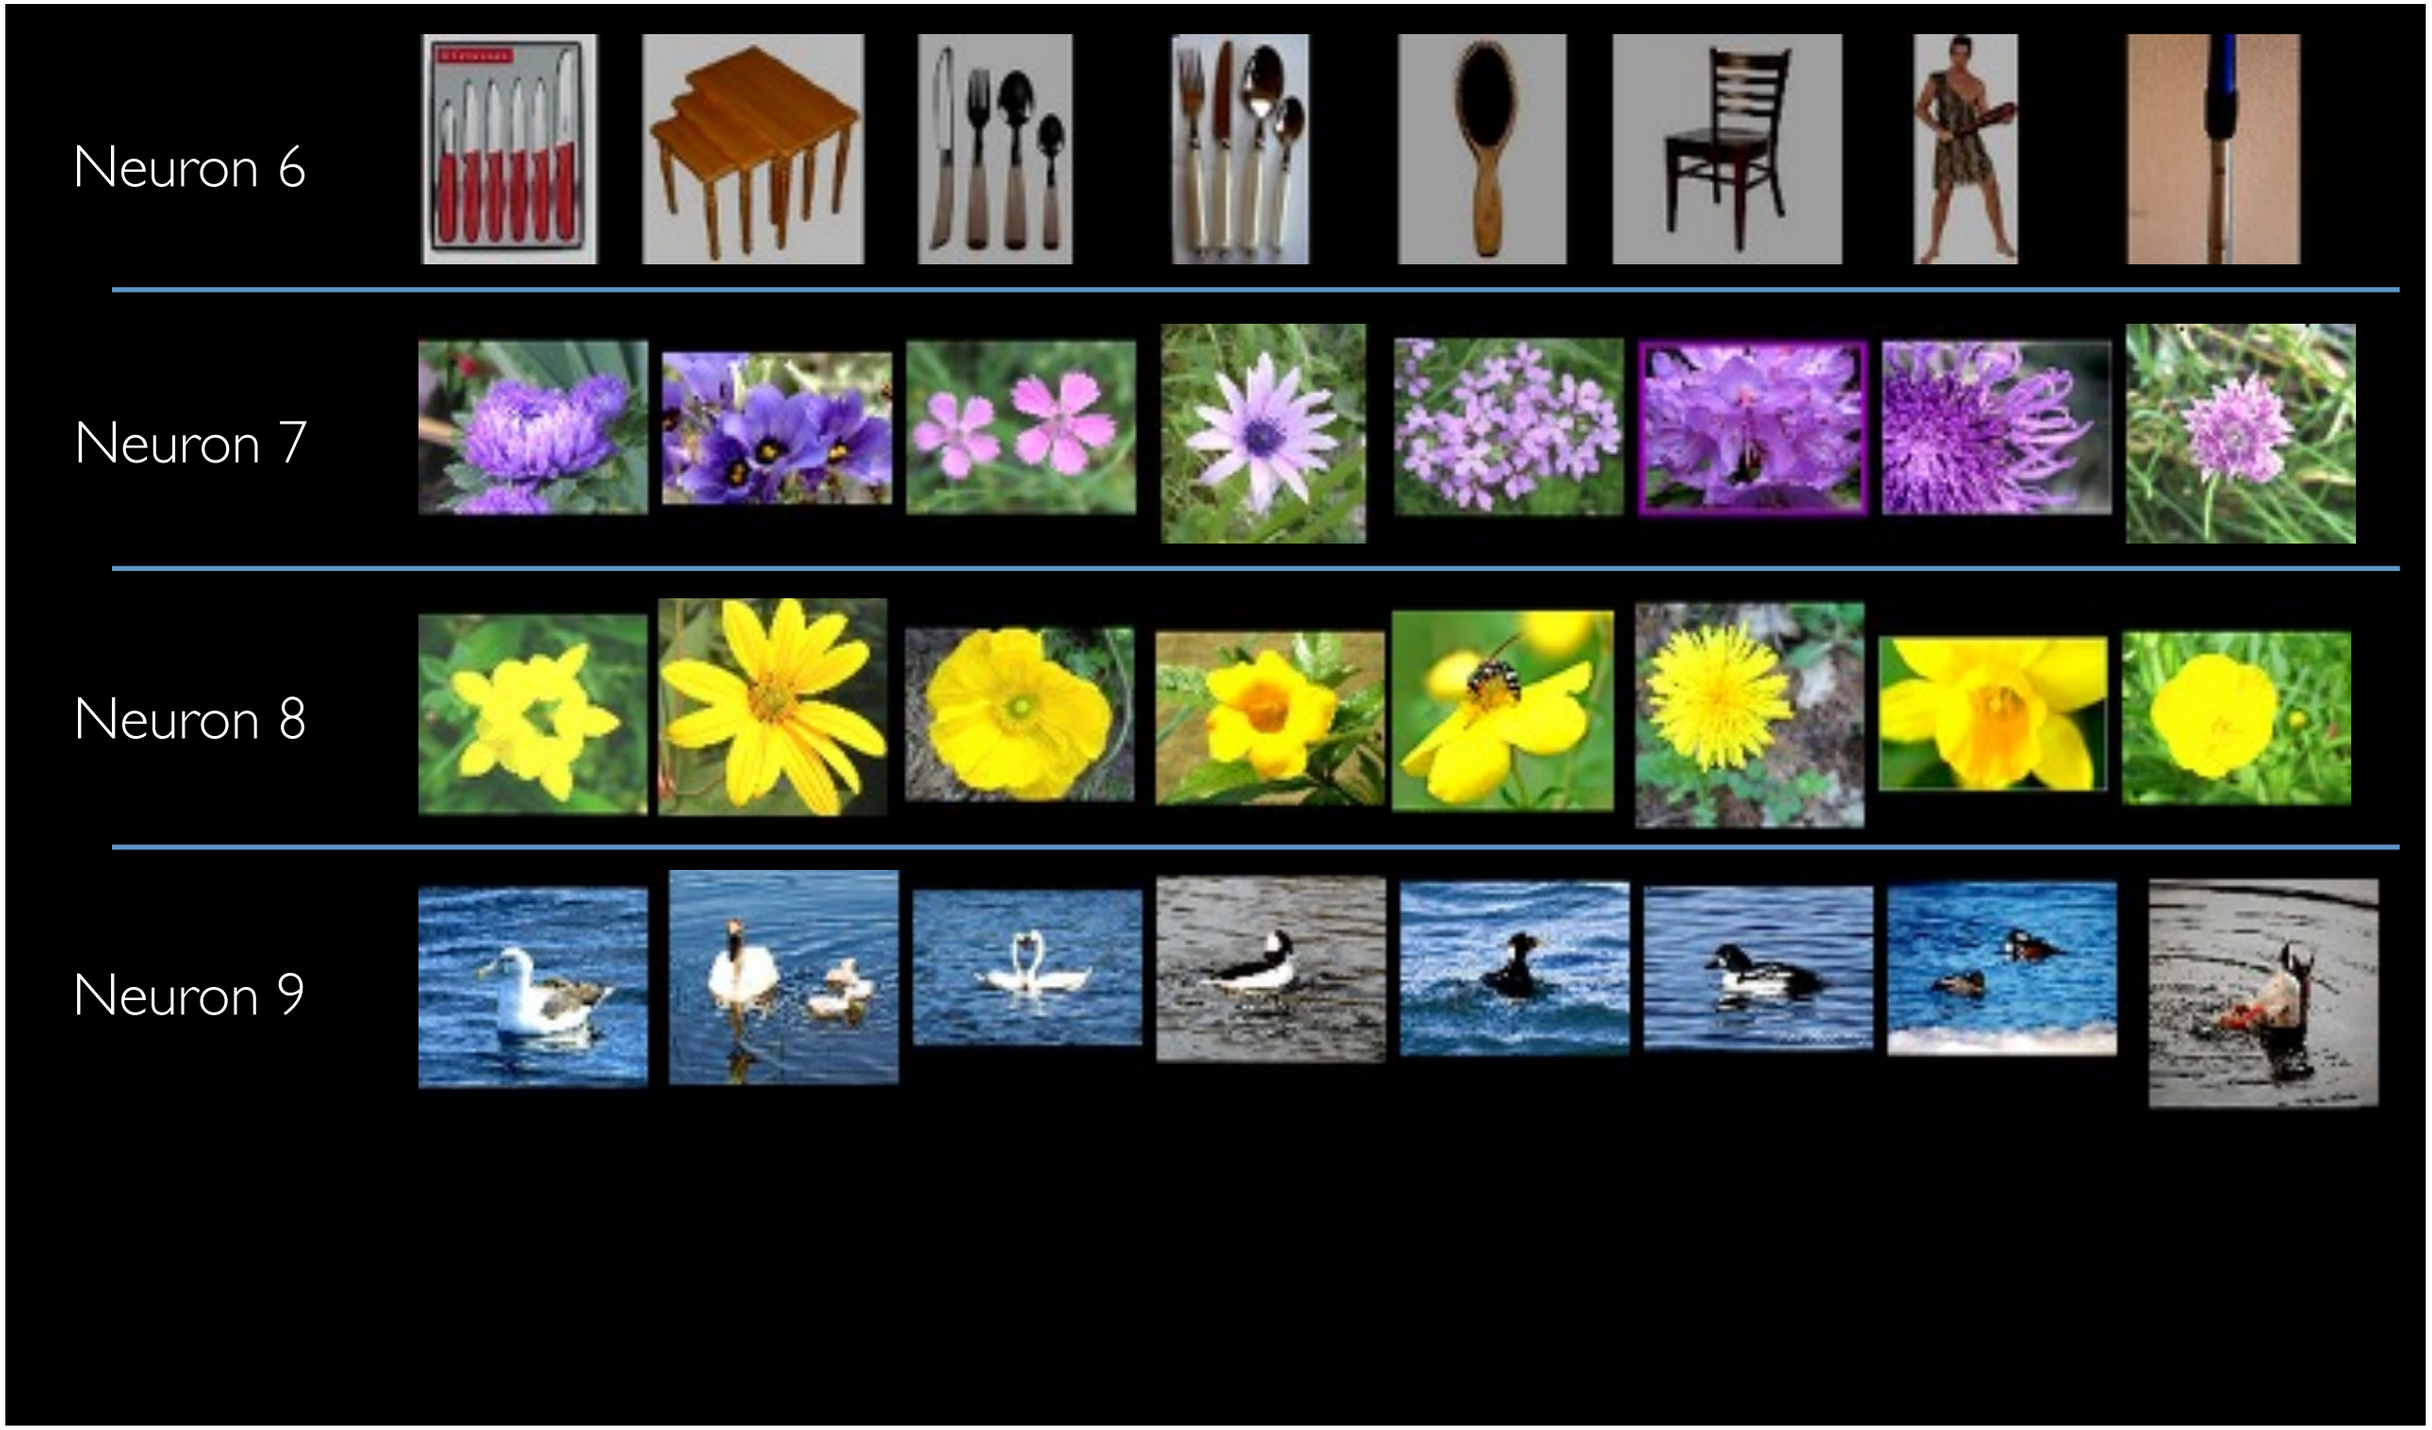

Semi-supervised Feature Learning in Images

Example top stimuli after fine tuning on ImageNet:

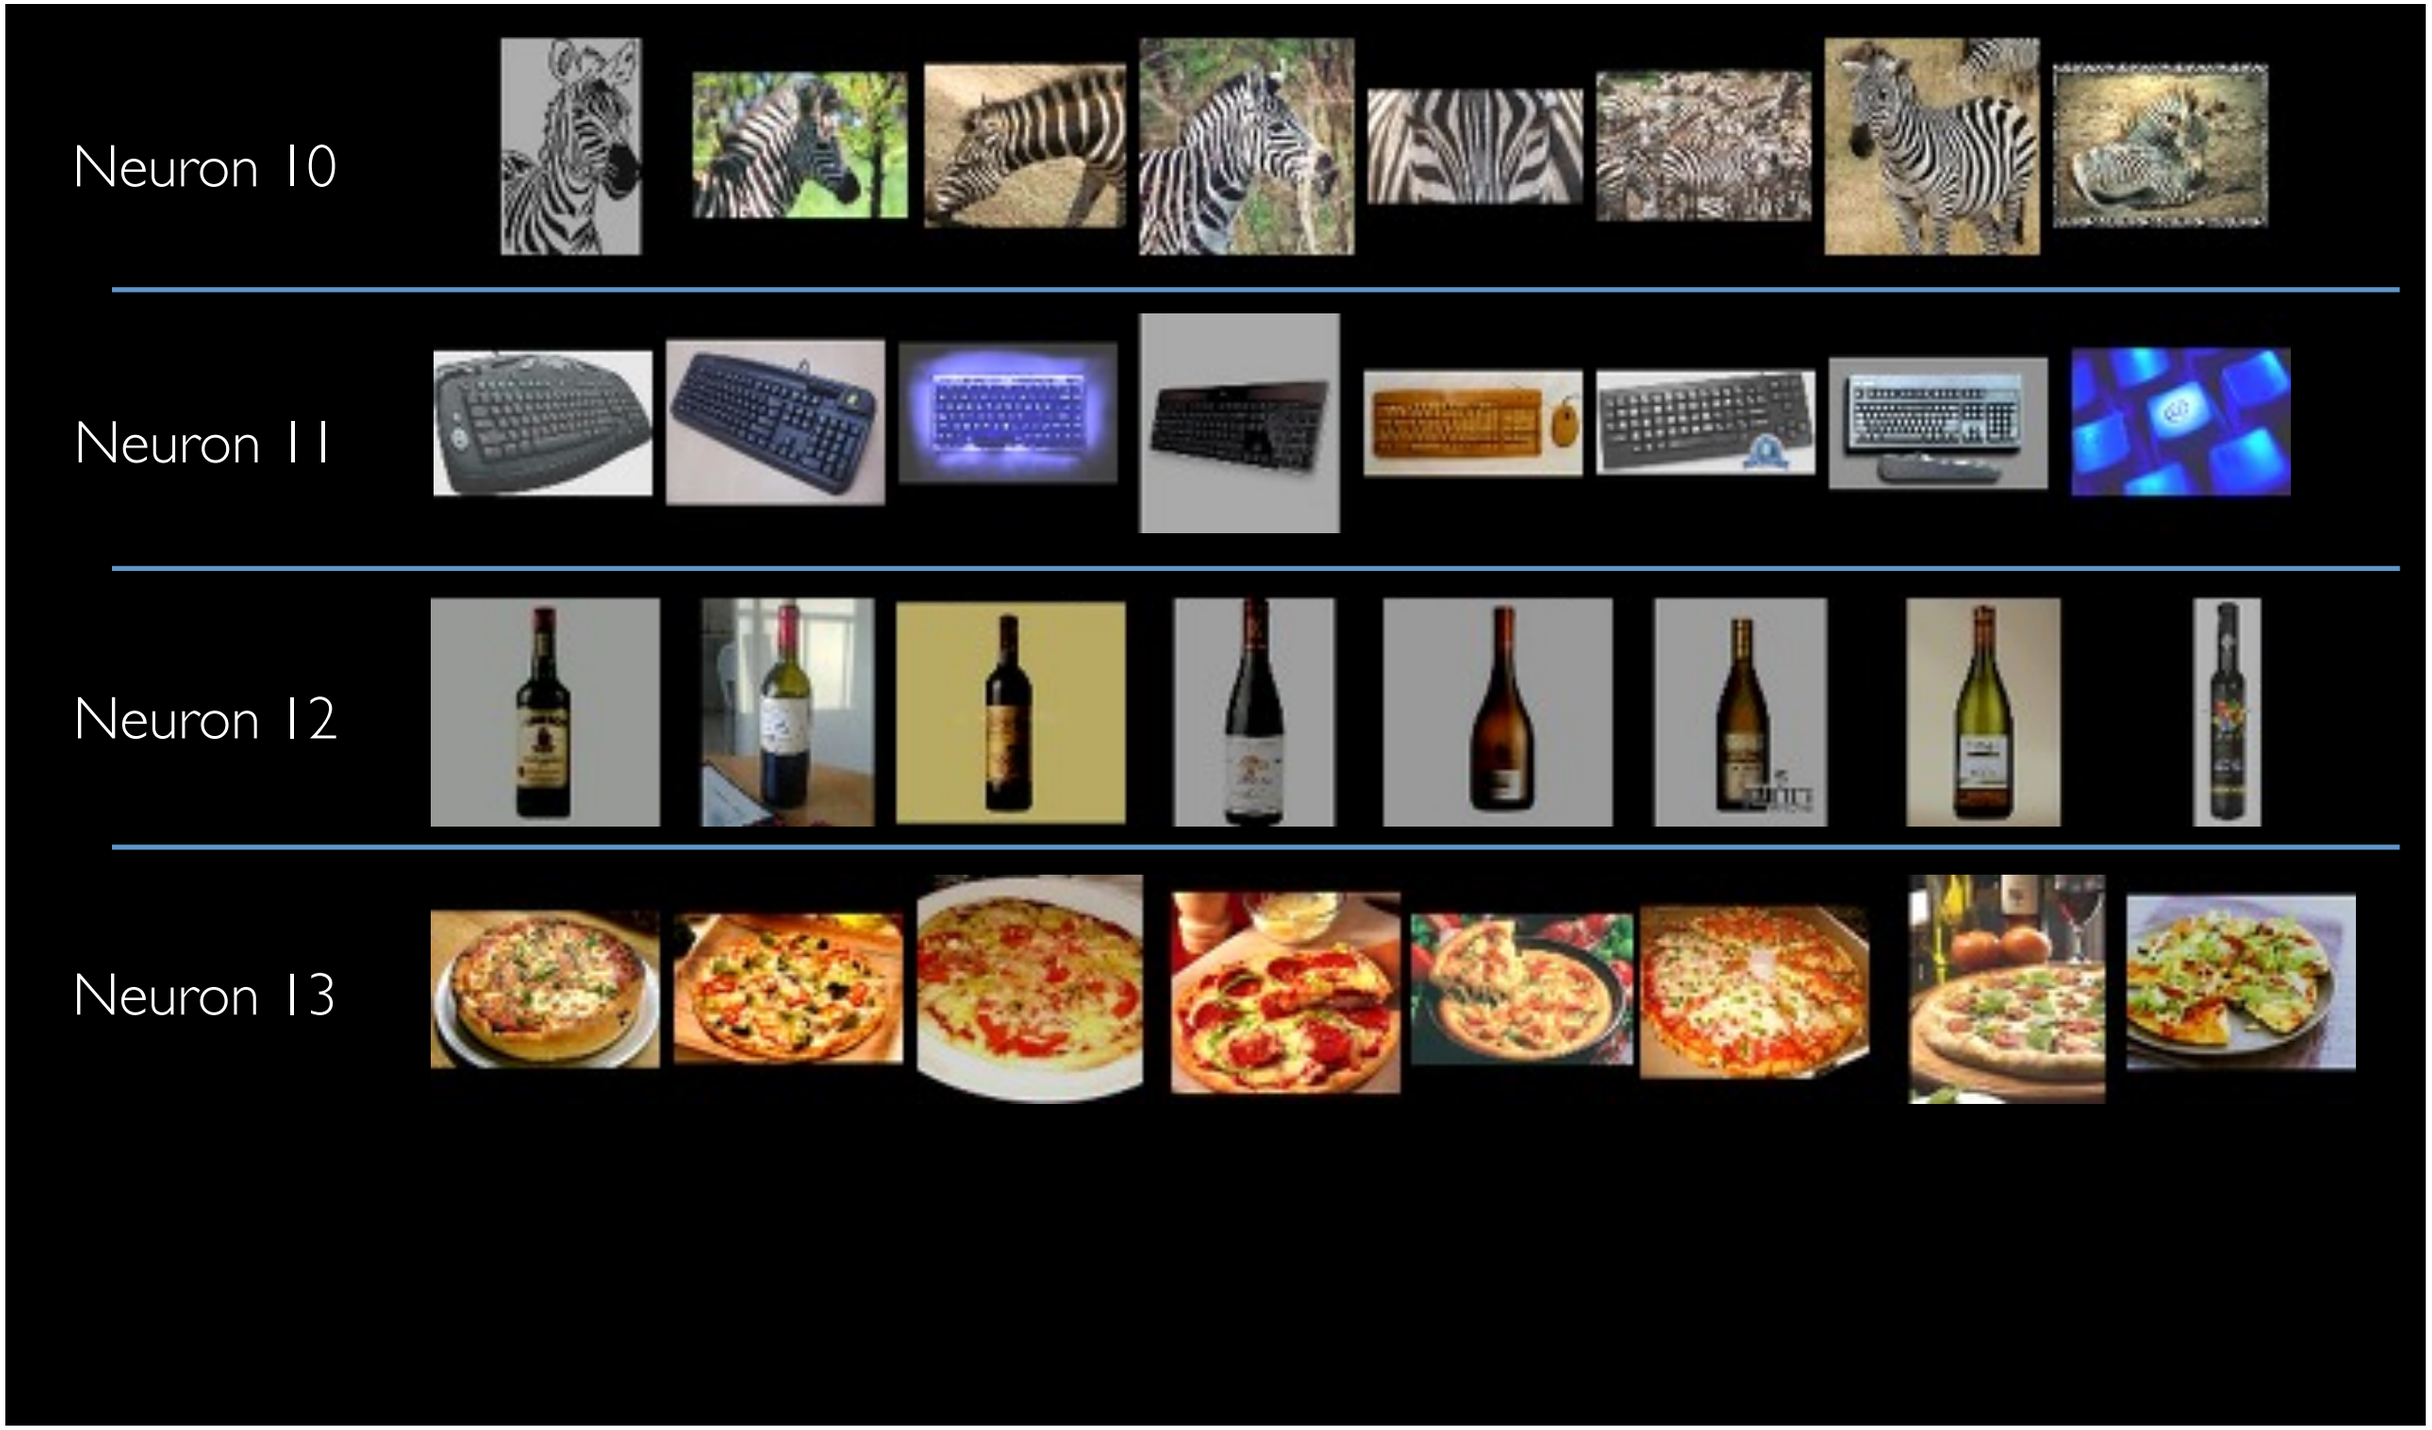

Semi-supervised Feature Learning in Images

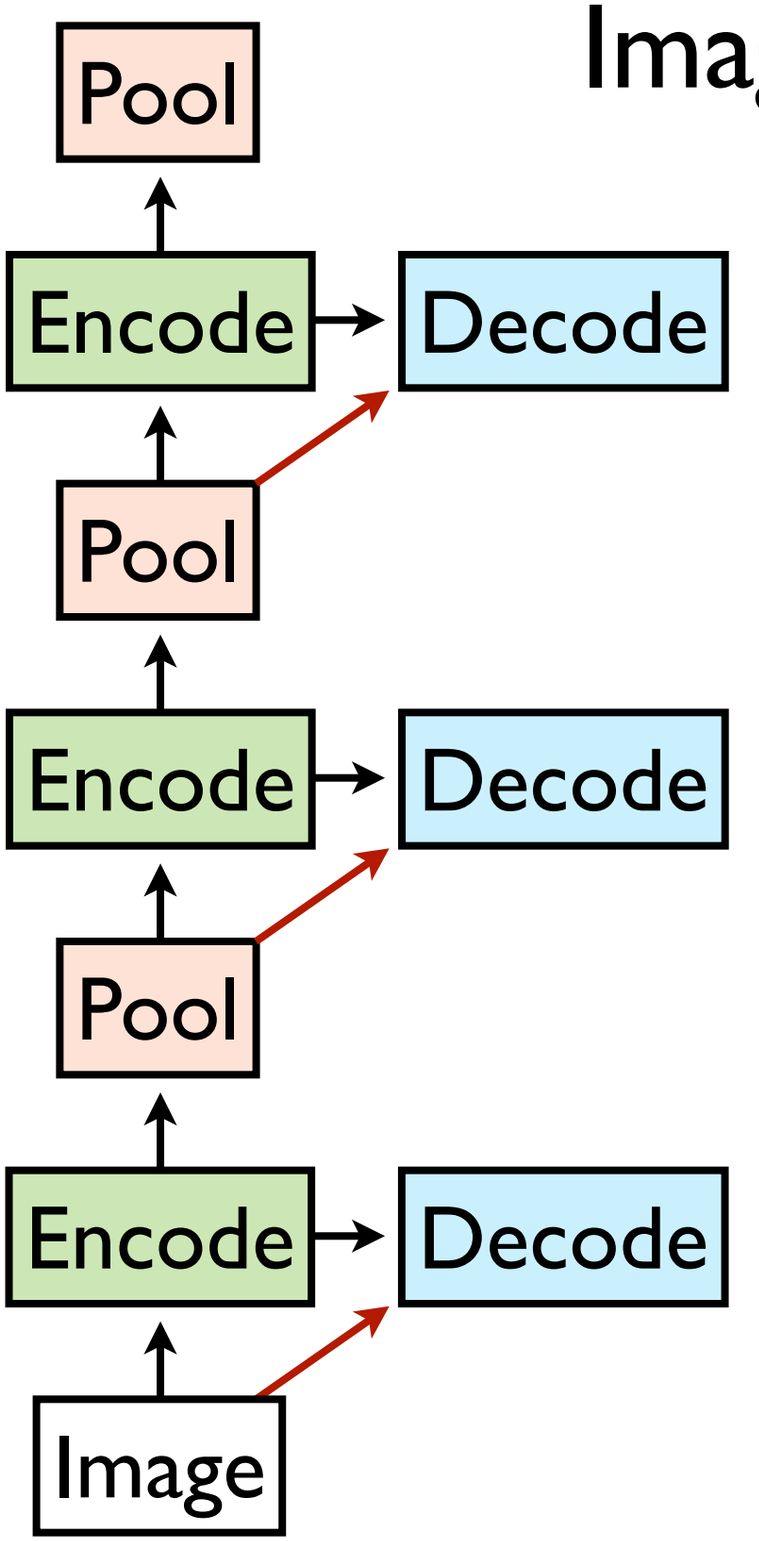

ImageNet Classification Results:

ImageNet 2011 (20k categories)

- Chance: 0.005%
- Best reported: 9.5%
- Our network: 16% (+70% relative)

Applications

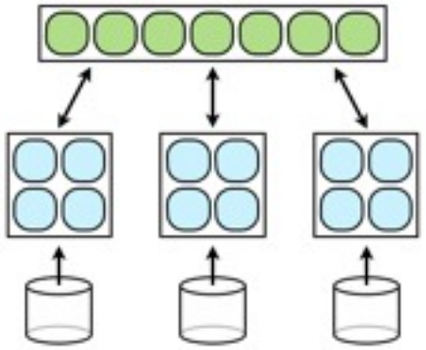

- Acoustic Models for Speech
- Unsupervised Feature Learning for Still Images
- Neural Language Models

Embeddings

~100-D joint embedding space

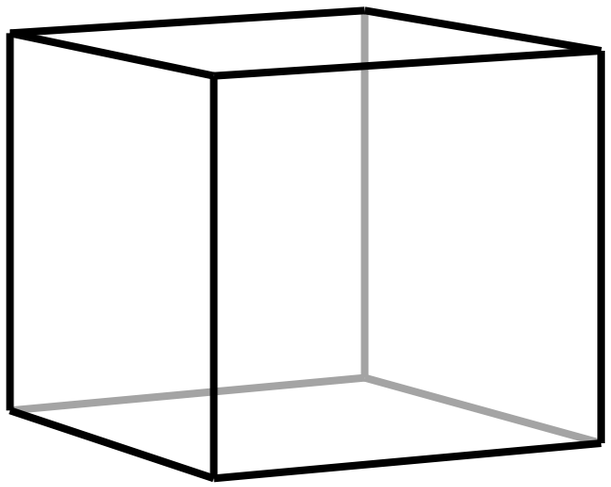

porpoise

dolphin

Embeddings

~100-D joint embedding space

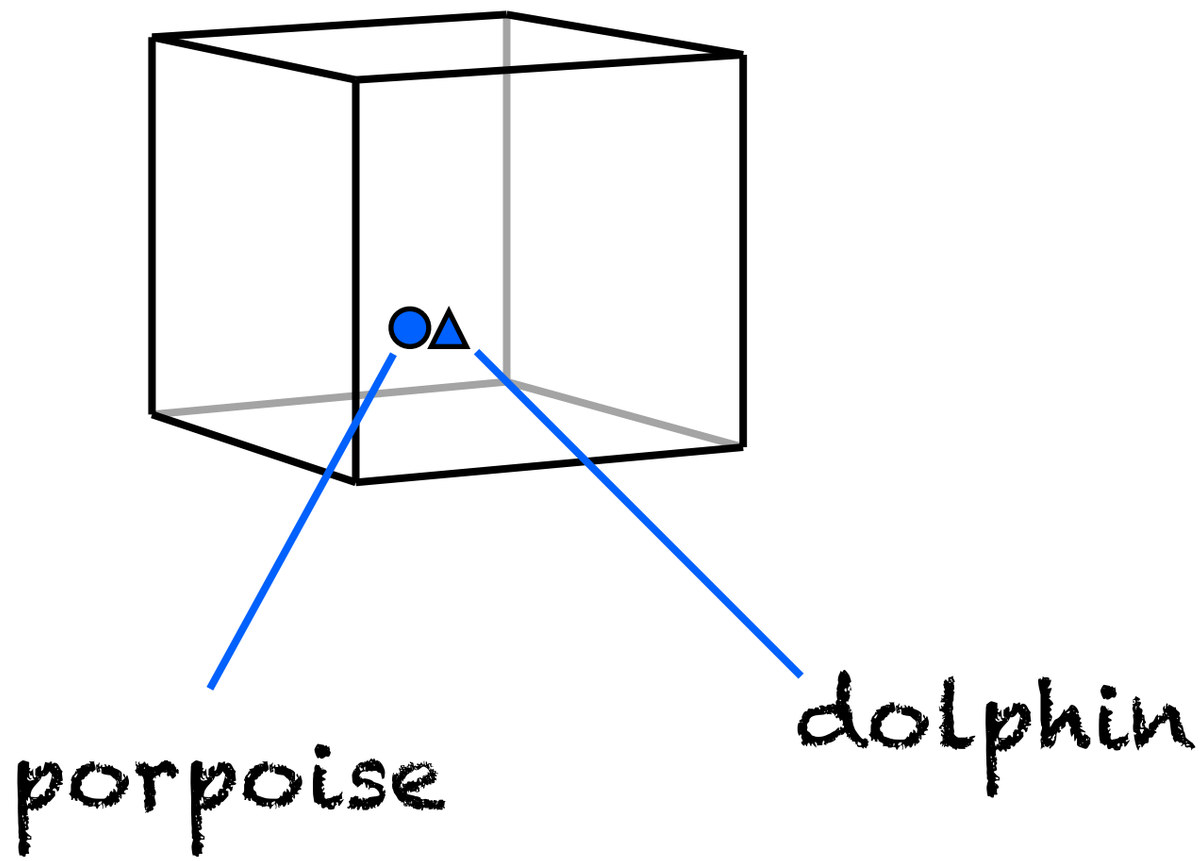

Embeddings

~100-D joint embedding space

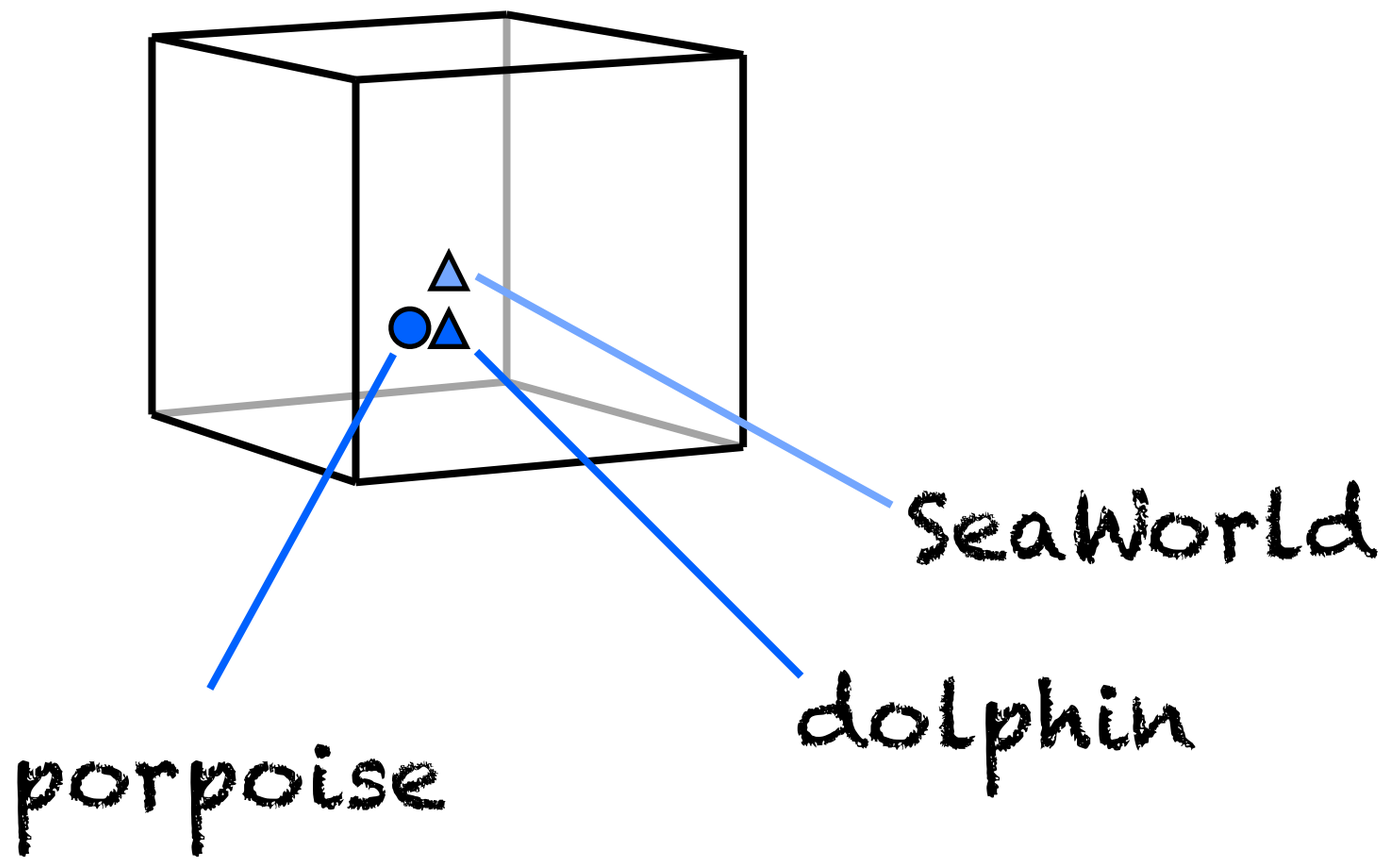

Embeddings

~100-D joint embedding space

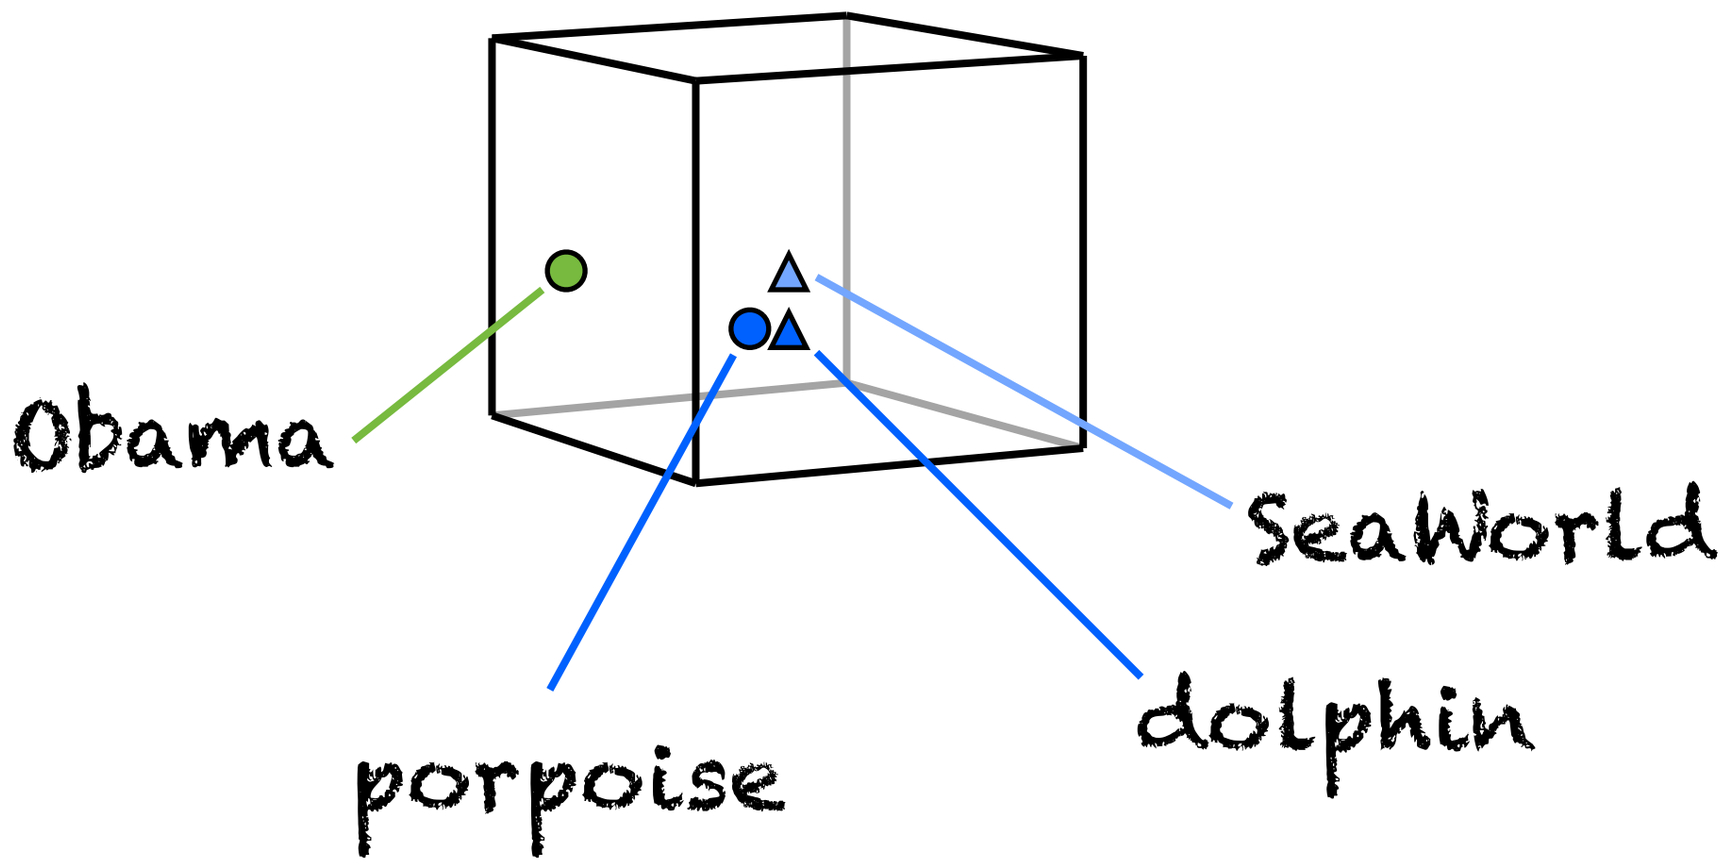

Embeddings

~100-D joint embedding space

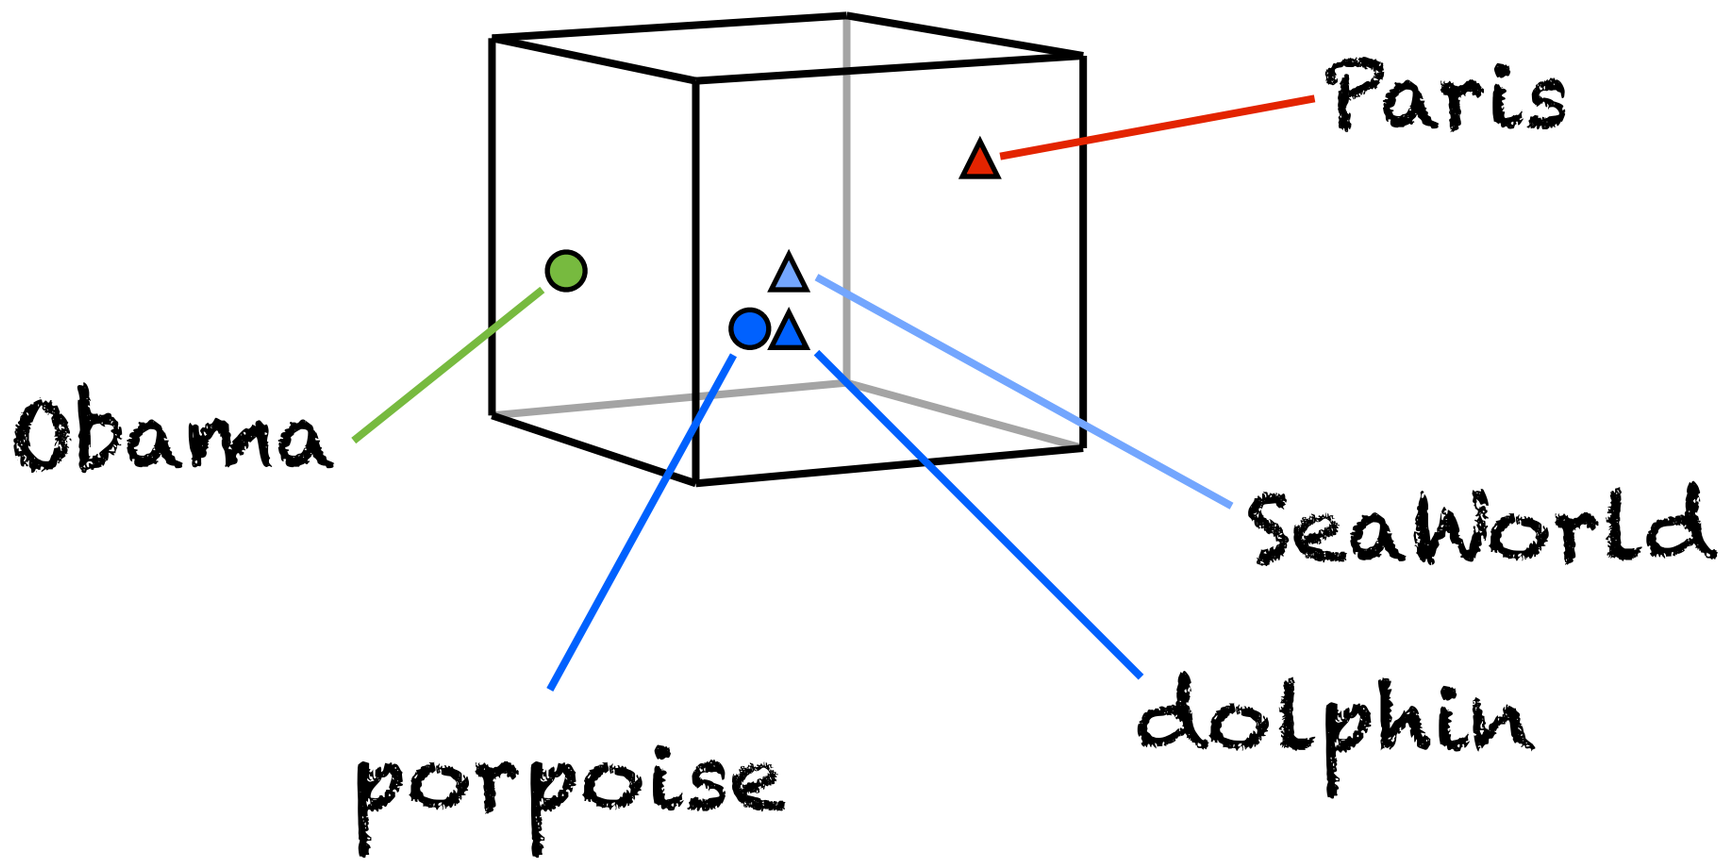

Neural Language Models

Hinge Loss // Softmax

Hidden Layers?

Word Embedding Matrix

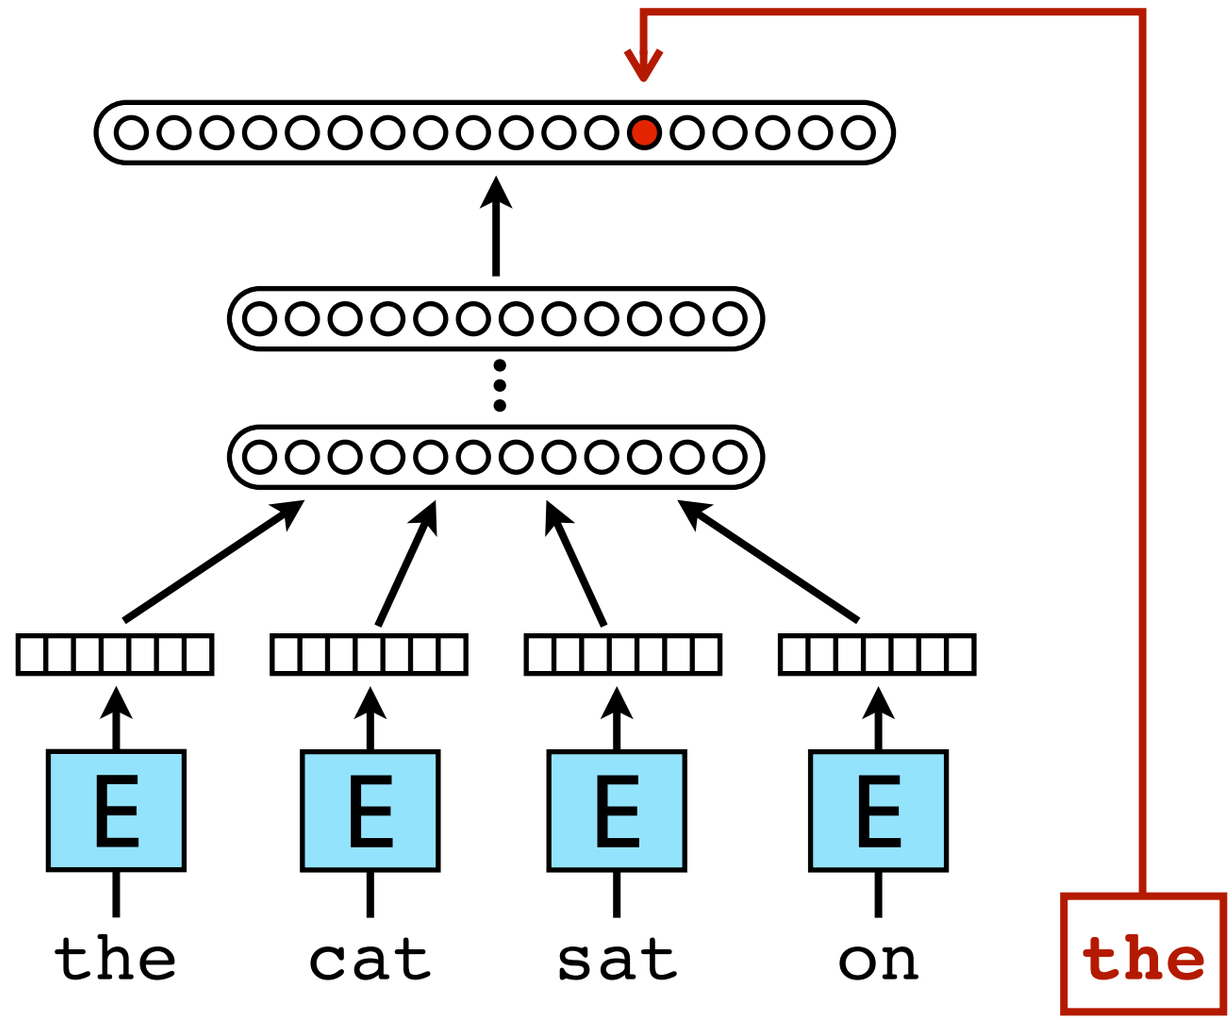

E is a matrix of dimension ||*Vocab*|| x *d*

Top prediction layer has ||*Vocab*|| x *h* parameters.

Most ideas from Bengio et al 2003, Collobert & Weston 2008

Neural Language Models

Hinge Loss // Softmax

Hidden Layers?

Word Embedding Matrix

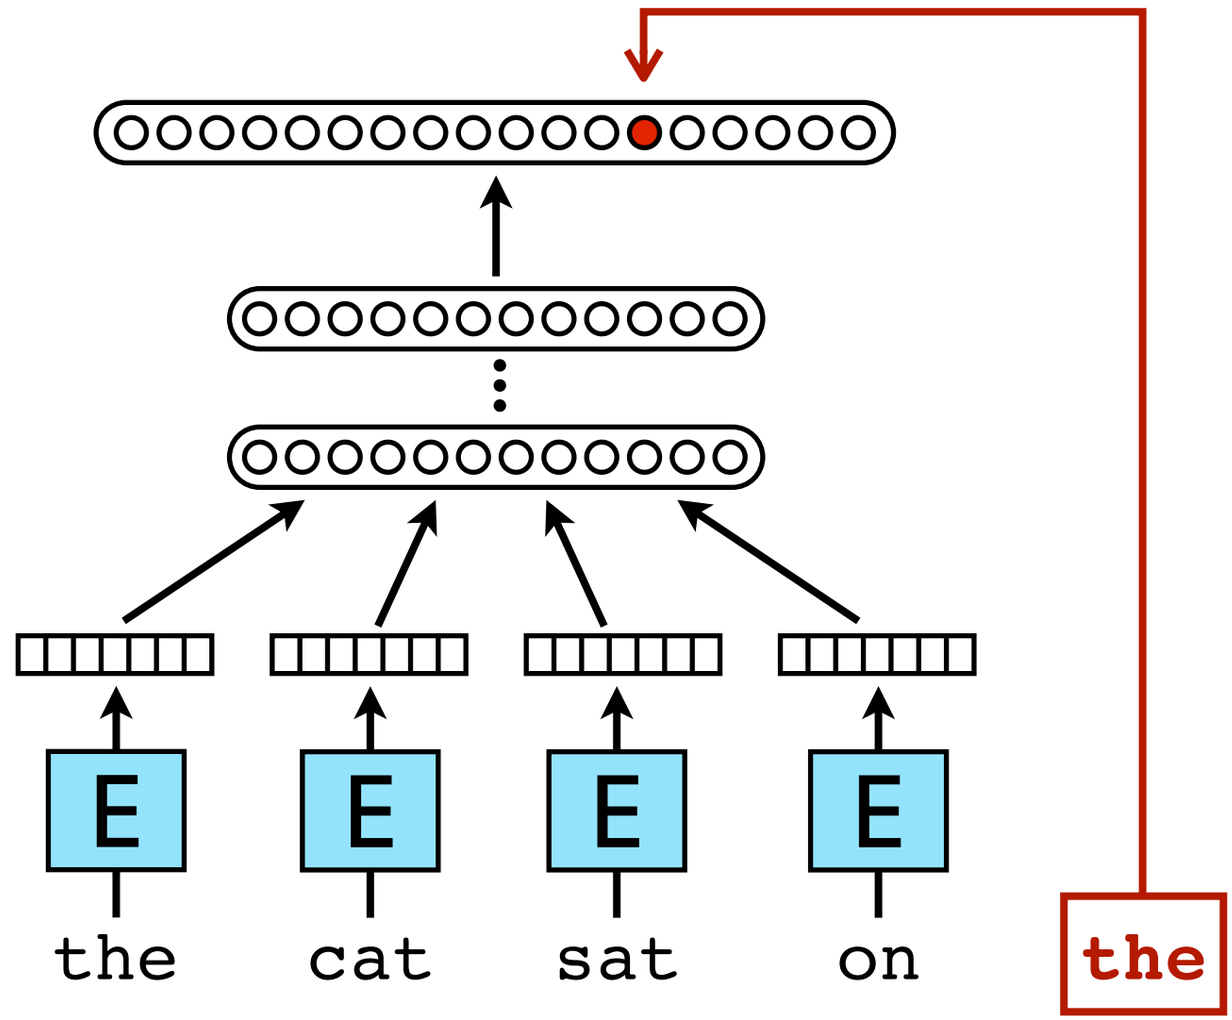

E is a matrix of dimension ||*Vocab*|| x *d*

Top prediction layer has ||*Vocab*|| x *h* parameters.

} 100s of millions of parameters, but gradients very sparse

Most ideas from Bengio et al 2003, Collobert & Weston 2008

Embedding sparse tokens in an N-dimensional space

Example: 50-D embedding trained for semantic similarity

Cluster 1: apple

Cluster 1

Columns | Row filter (regexp)

Id	Distance↑	Adjust	Word
11114	0.000000	Remove	apple
5026	0.652580	Add	fruit
14080	0.699192	Add	apples
48657	0.717818	Add	melon
28498	0.722390	Add	peach
39795	0.729893	Add	blueberry
35570	0.730500	Add	berry
25974	0.739561	Add	strawberry
46156	0.745343	Add	pecan
11907	0.756422	Add	potato
33847	0.759111	Add	pear
30895	0.763317	Add	mango
17848	0.768230	Add	pumpkin
39133	0.770143	Add	almond
14395	0.773105	Add	tomato
18163	0.782610	Add	onion
10470	0.782994	Add	pie
3023	0.787229	Add	tree
20340	0.793602	Add	bean
34968	0.794979	Add	watermelon

Embedding sparse tokens in an N-dimensional space

Example: 50-D embedding trained for semantic similarity

Cluster 1: apple

Cluster 1

Columns | Row filter (regexp)

Id	Distance↑	Adjust	Word
11114	0.000000	Remove	apple
5026	0.652580	Add	fruit
14080	0.699192	Add	apples
48657	0.717818	Add	melon
28498	0.722390	Add	peach
39795	0.729893	Add	blueberry
35570	0.730500	Add	berry
25974	0.739561	Add	strawberry
46156	0.745343	Add	pecan
11907	0.756422	Add	potato
33847	0.759111	Add	pear
30895	0.763317	Add	mango
17848	0.768230	Add	pumpkin
39133	0.770143	Add	almond
14395	0.773105	Add	tomato
18163	0.782610	Add	onion
10470	0.782994	Add	pie
3023	0.787229	Add	tree
20340	0.793602	Add	bean
34968	0.794979	Add	watermelon

Cluster 1: stab

Cluster 1

Columns | Row filter (regexp)

Id	Distance↑	Adjust	Word
14979	0.000000	Remove	stab
7728	0.868853	Add	punch
469	0.909304	Add	shot
12820	0.909750	Add	thrust
8934	0.939908	Add	shell
10880	0.951466	Add	hammer
6975	0.951679	Add	bullet
1848	0.962053	Add	push
10888	0.962319	Add	eyed
718	0.965448	Add	hand
5865	0.966663	Add	grab
4611	0.967574	Add	swing
302	0.975696	Add	hit
869	0.976967	Add	force
1597	0.977625	Add	attempt
5977	0.978384	Add	finger
6162	0.978776	Add	knife
3434	0.980028	Add	sharp
1504	0.980160	Add	struck
39157	0.980219	Add	slug

Embedding sparse tokens in an N-dimensional space

Example: 50-D embedding trained for semantic similarity

Cluster 1: apple

Cluster 1

Columns | Row filter (regexp)

Id	Distance↑	Adjust	Word
11114	0.000000	Remove	apple
5026	0.652580	Add	fruit
14080	0.699192	Add	apples
48657	0.717818	Add	melon
28498	0.722390	Add	peach
39795	0.729893	Add	blueberry
35570	0.730500	Add	berry
25974	0.739561	Add	strawberry
46156	0.745343	Add	pecan
11907	0.756422	Add	potato
33847	0.759111	Add	pear
30895	0.763317	Add	mango
17848	0.768230	Add	pumpkin
39133	0.770143	Add	almond
14395	0.773105	Add	tomato
18163	0.782610	Add	onion
10470	0.782994	Add	pie
3023	0.787229	Add	tree
20340	0.793602	Add	bean
34968	0.794979	Add	watermelon

Cluster 1: stab

Cluster 1

Columns | Row filter (regexp)

Id	Distance↑	Adjust	Word
14979	0.000000	Remove	stab
7728	0.868853	Add	punch
469	0.909304	Add	shot
12820	0.909750	Add	thrust
8934	0.939908	Add	shell
10880	0.951466	Add	hammer
6975	0.951679	Add	bullet
1848	0.962053	Add	push
10888	0.962319	Add	eyed
718	0.965448	Add	hand
5865	0.966663	Add	grab
4611	0.967574	Add	swing
302	0.975696	Add	hit
869	0.976967	Add	force
1597	0.977625	Add	attempt
5977	0.978384	Add	finger
6162	0.978776	Add	knife
3434	0.980028	Add	sharp
1504	0.980160	Add	struck
39157	0.980219	Add	slug

Cluster 1: iPhone

Cluster 1

Columns | Row filter (regexp)

Id	Distance↑	Adjust	Word
2964	0.000000	Remove	iPhone
6377	0.359153	Add	iPad
22542	0.554838	Add	iOS
10081	0.585379	Add	smartphone
5824	0.587948	Add	iPod
43921	0.608292	Add	PlayBook
18025	0.653021	Add	iPhones
6439	0.656983	Add	Android
38104	0.681779	Add	3GS
8088	0.690880	Add	BlackBerry
24581	0.696648	Add	Zune
33435	0.713150	Add	Smartphone
19186	0.714883	Add	Blackberry
9326	0.715027	Add	handset
26020	0.739856	Add	Droid
30557	0.756973	Add	Treo
12057	0.762164	Add	smartphones
6878	0.769016	Add	app
8211	0.779153	Add	iTunes
28120	0.787939	Add	iPads

Neural Language Models

- 7 Billion word Google News training set
- 1 Million word vocabulary
- 8 word history, 50 dimensional embedding
- Three hidden layers each w/200 nodes
- 50-100 asynchronous model workers

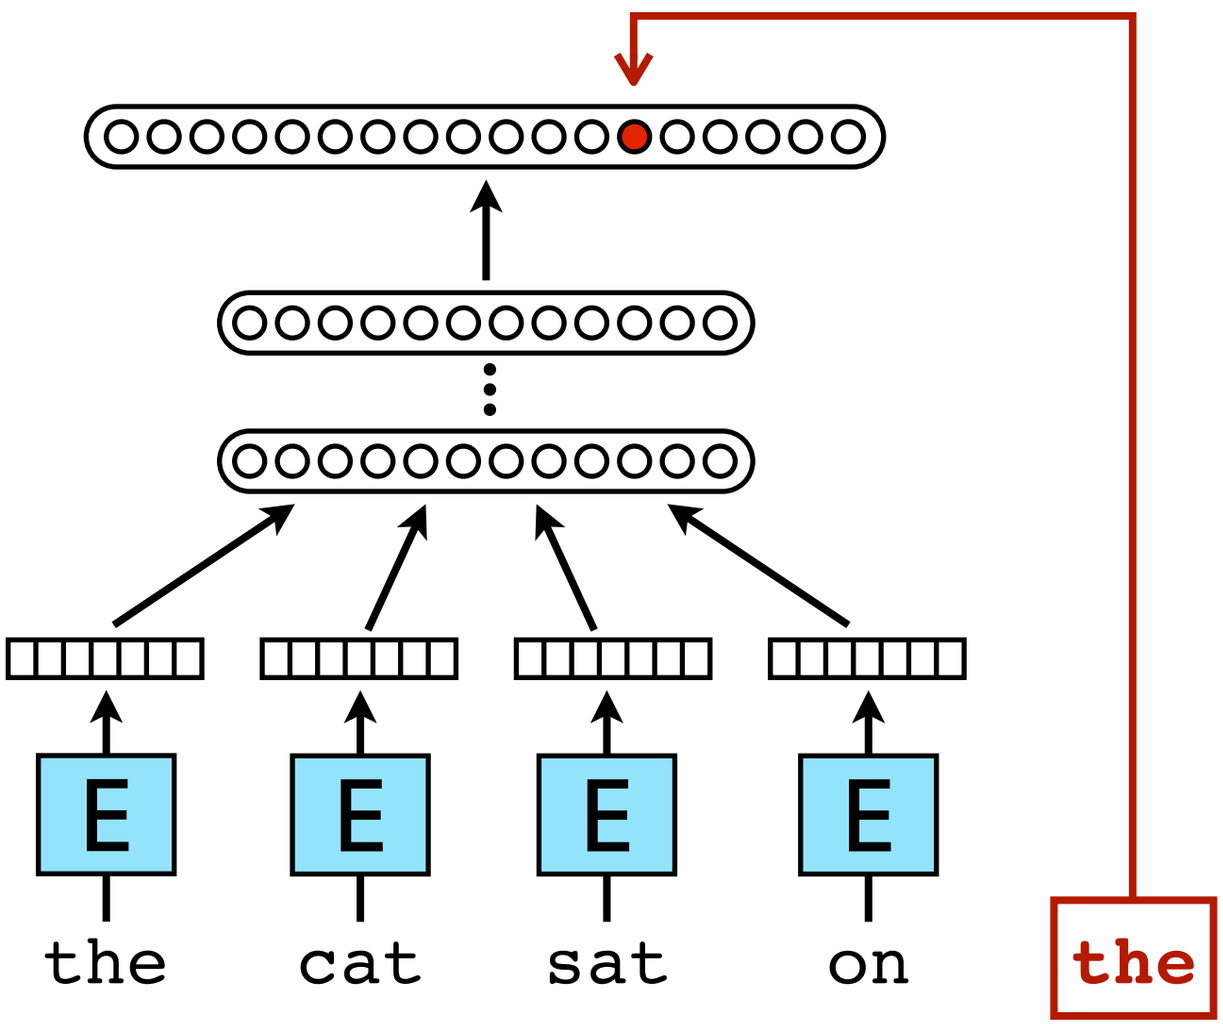

Neural Language Models

- 7 Billion word Google News training set
- 1 Million word vocabulary
- 8 word history, 50 dimensional embedding
- Three hidden layers each w/200 nodes
- 50-100 asynchronous model workers

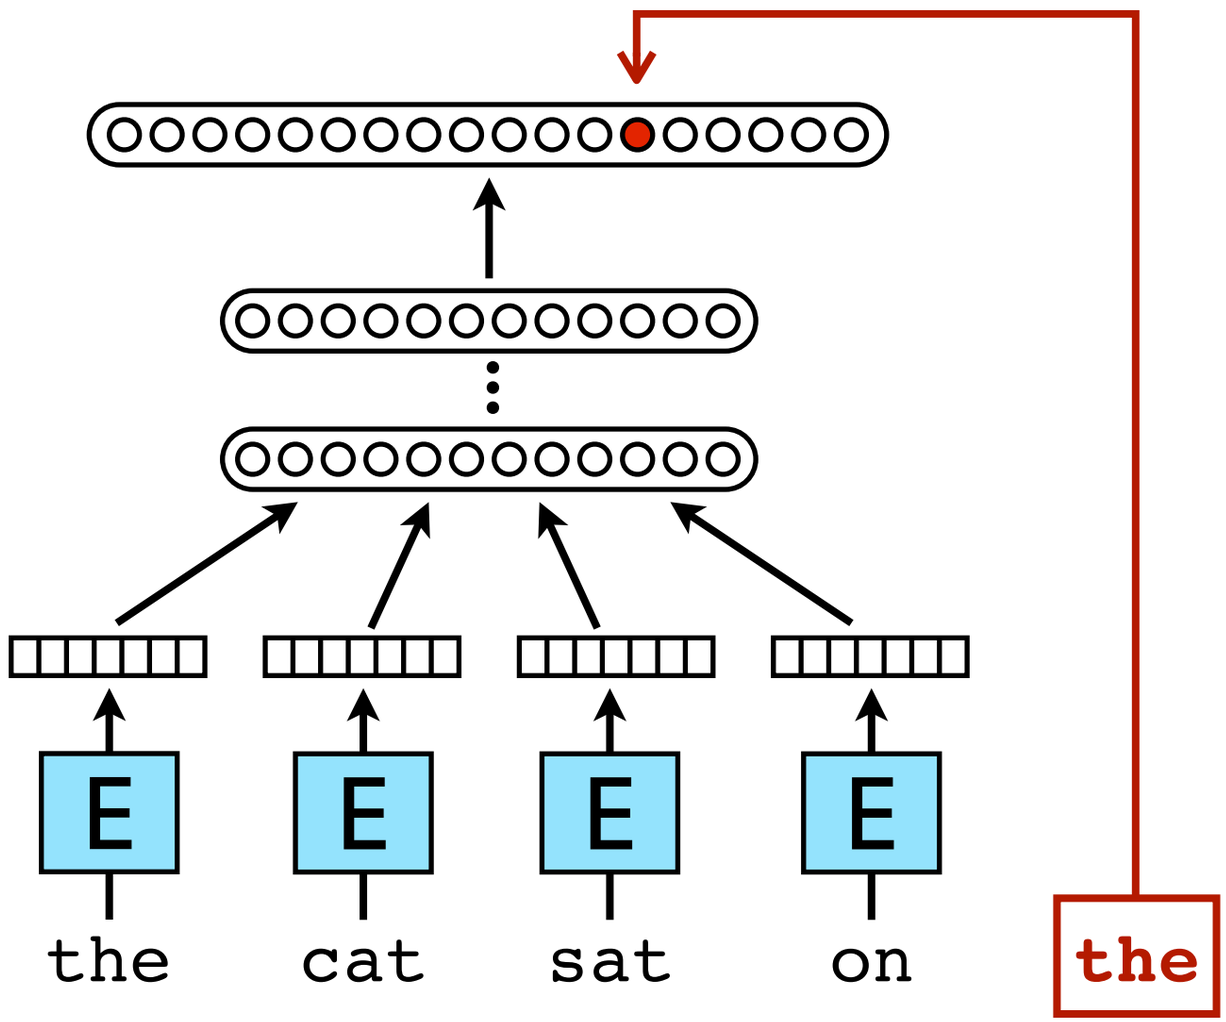

Perplexity Scores

Traditional 5-gram	XXX
NLM	+15%
5-gram + NLM	-33%

Summary

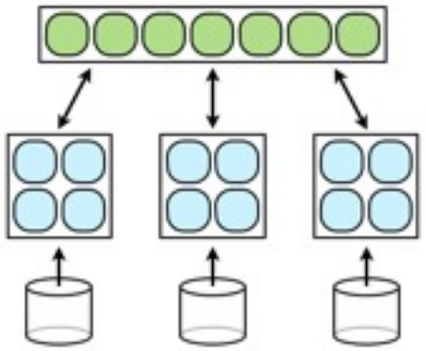

Summary

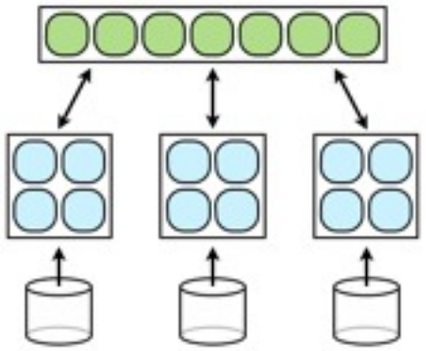

- Two interesting "traditional systems" areas:
  - Building low-latency services out of variable-latency parts
  - Single systems that coordinate activity across the whole world, rather than separate instances per cluster/datacenter

Summary

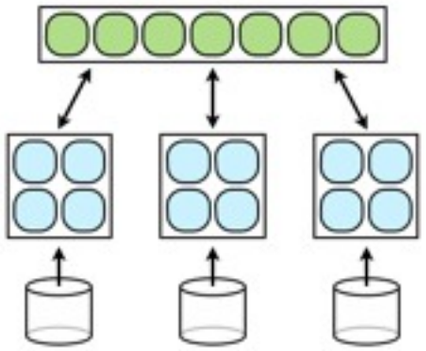

- Two interesting "traditional systems" areas:
  - Building low-latency services out of variable-latency parts
  - Single systems that coordinate activity across the whole world, rather than separate instances per cluster/datacenter

- Systems work to scale deep learning gives promising results:
  - Work at intersection of systems and machine learning
  - Speech: Supervised model gives major improvement in WER.
  - Images: State of the art performance on ImageNet 21K classes
  - Neural language models: complementary to N-gram models

Many Thanks!

Variability reduction collaborators: Luiz Barroso + many others

World-wide storage system collaborators: Jay Corbett, Mike Epstein, Andrew Fikes, Chris Frost, JJ Furman, Sanjay Ghemawat, Andrey Gubarev, Chris Heiser, Peter Hochschild, Wilson Hsieh, Sebastian Kanthak, Eugene Kogan, Hongyi Li, Alex Lloyd, Sergey Melnik, David Mwauara, Sean Quinlan, Rajesh Rao, Lindsay Rolig, Yasushi Saito, Michal Szymaniak, Chris Taylor, Ruth Wang, Dale Woolford + others

Deep learning collaborators: Samy Bengio, Kai Chen, Greg Corrado, Tom Dean, Matthieu Devin, Quoc Le, Mark Mao, Rajat Monga, Andrew Ng, Patrick Nguyen, Marc'Aurelio Ranzato, Paul Tucker, Vincent Vanhoucke, Xiaoyun Wu, Peng Xe, Ke Yang + others