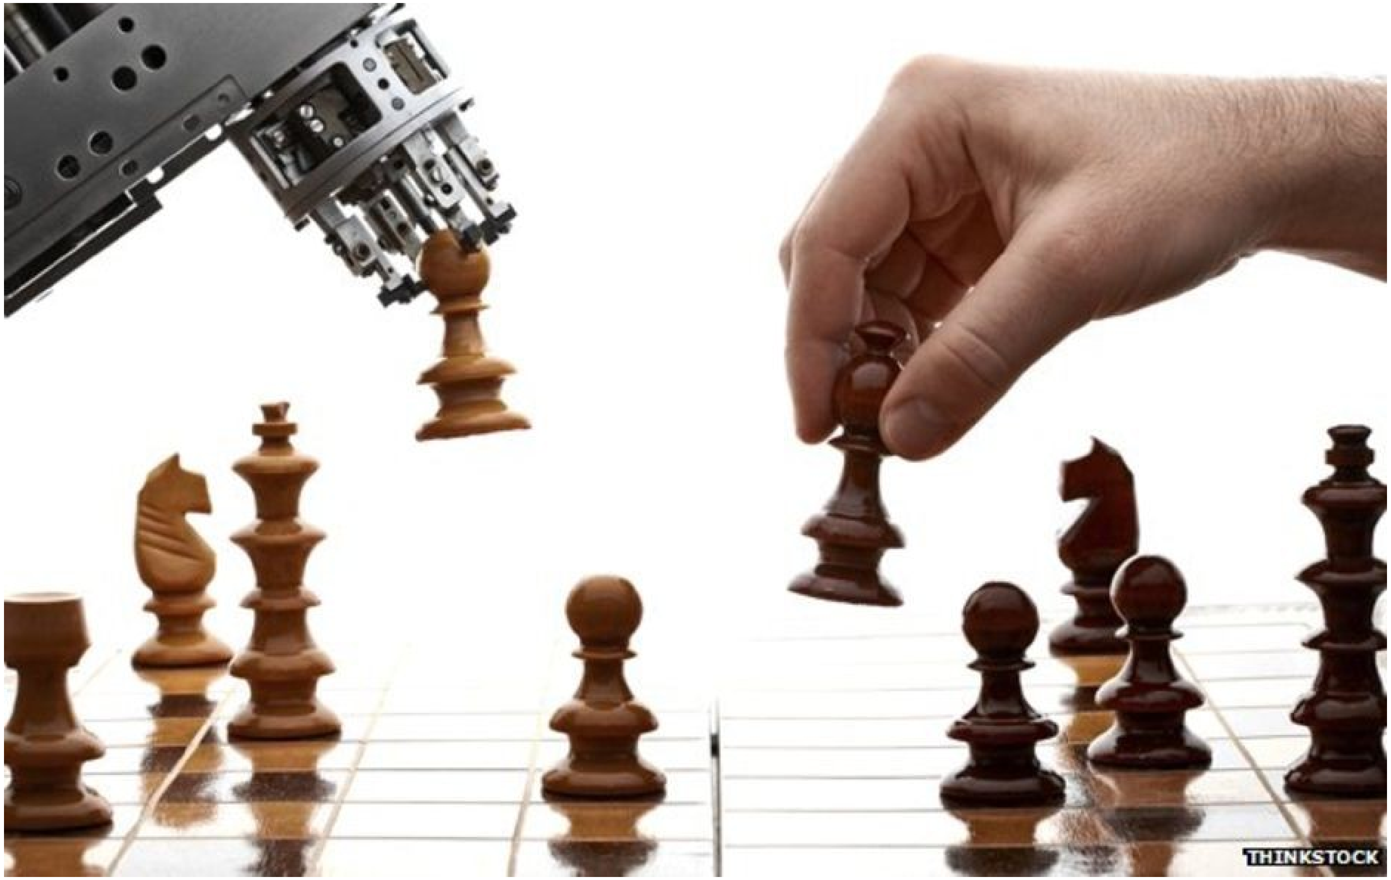

Introduction to Artificial Intelligence

COSC 4550 / COSC 5550

Professor Cheney

11/17/17

unsupervised learning continued…

Neural Photo Editing

Andrew Brock

EDINBURGH CENTRE FOR
ROBOTICS
Innovation Ready

Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks

Jun-Yan Zhu*   Taesung Park*   Phillip Isola   Alexei A. Efros

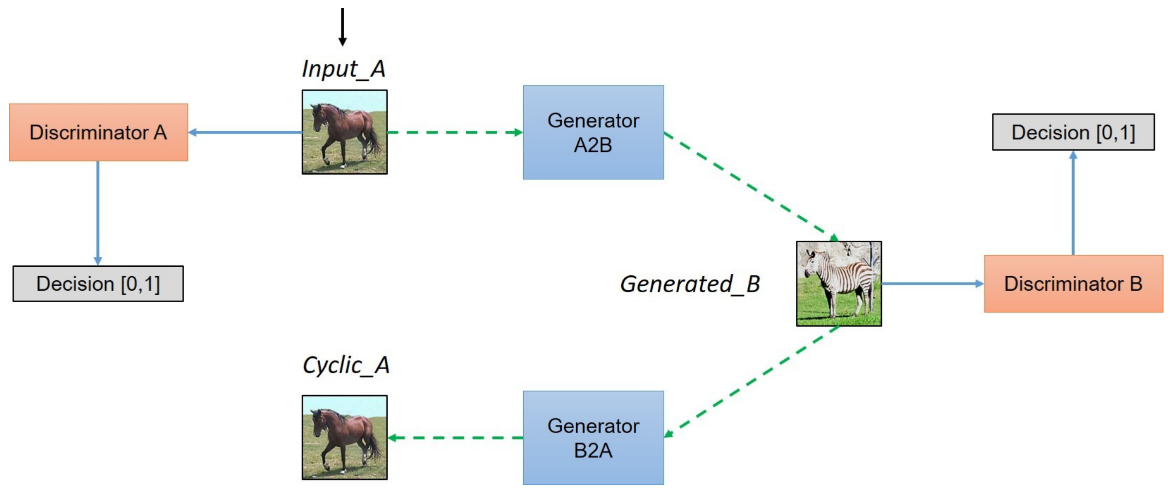

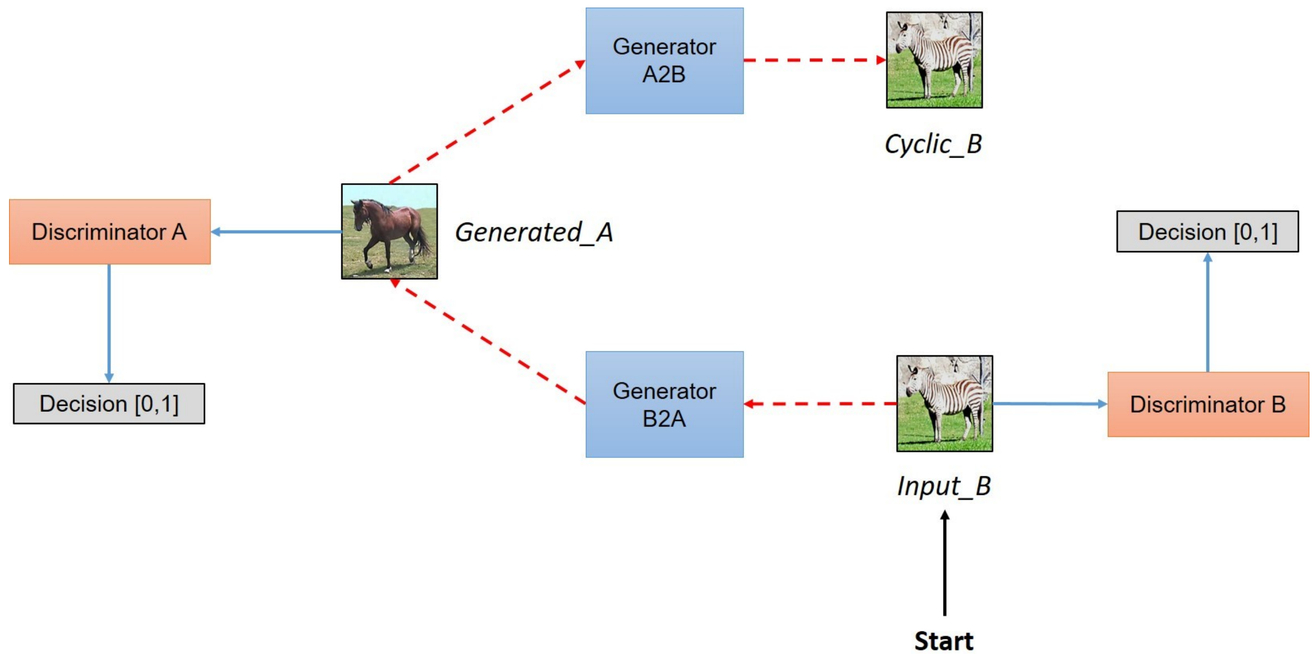

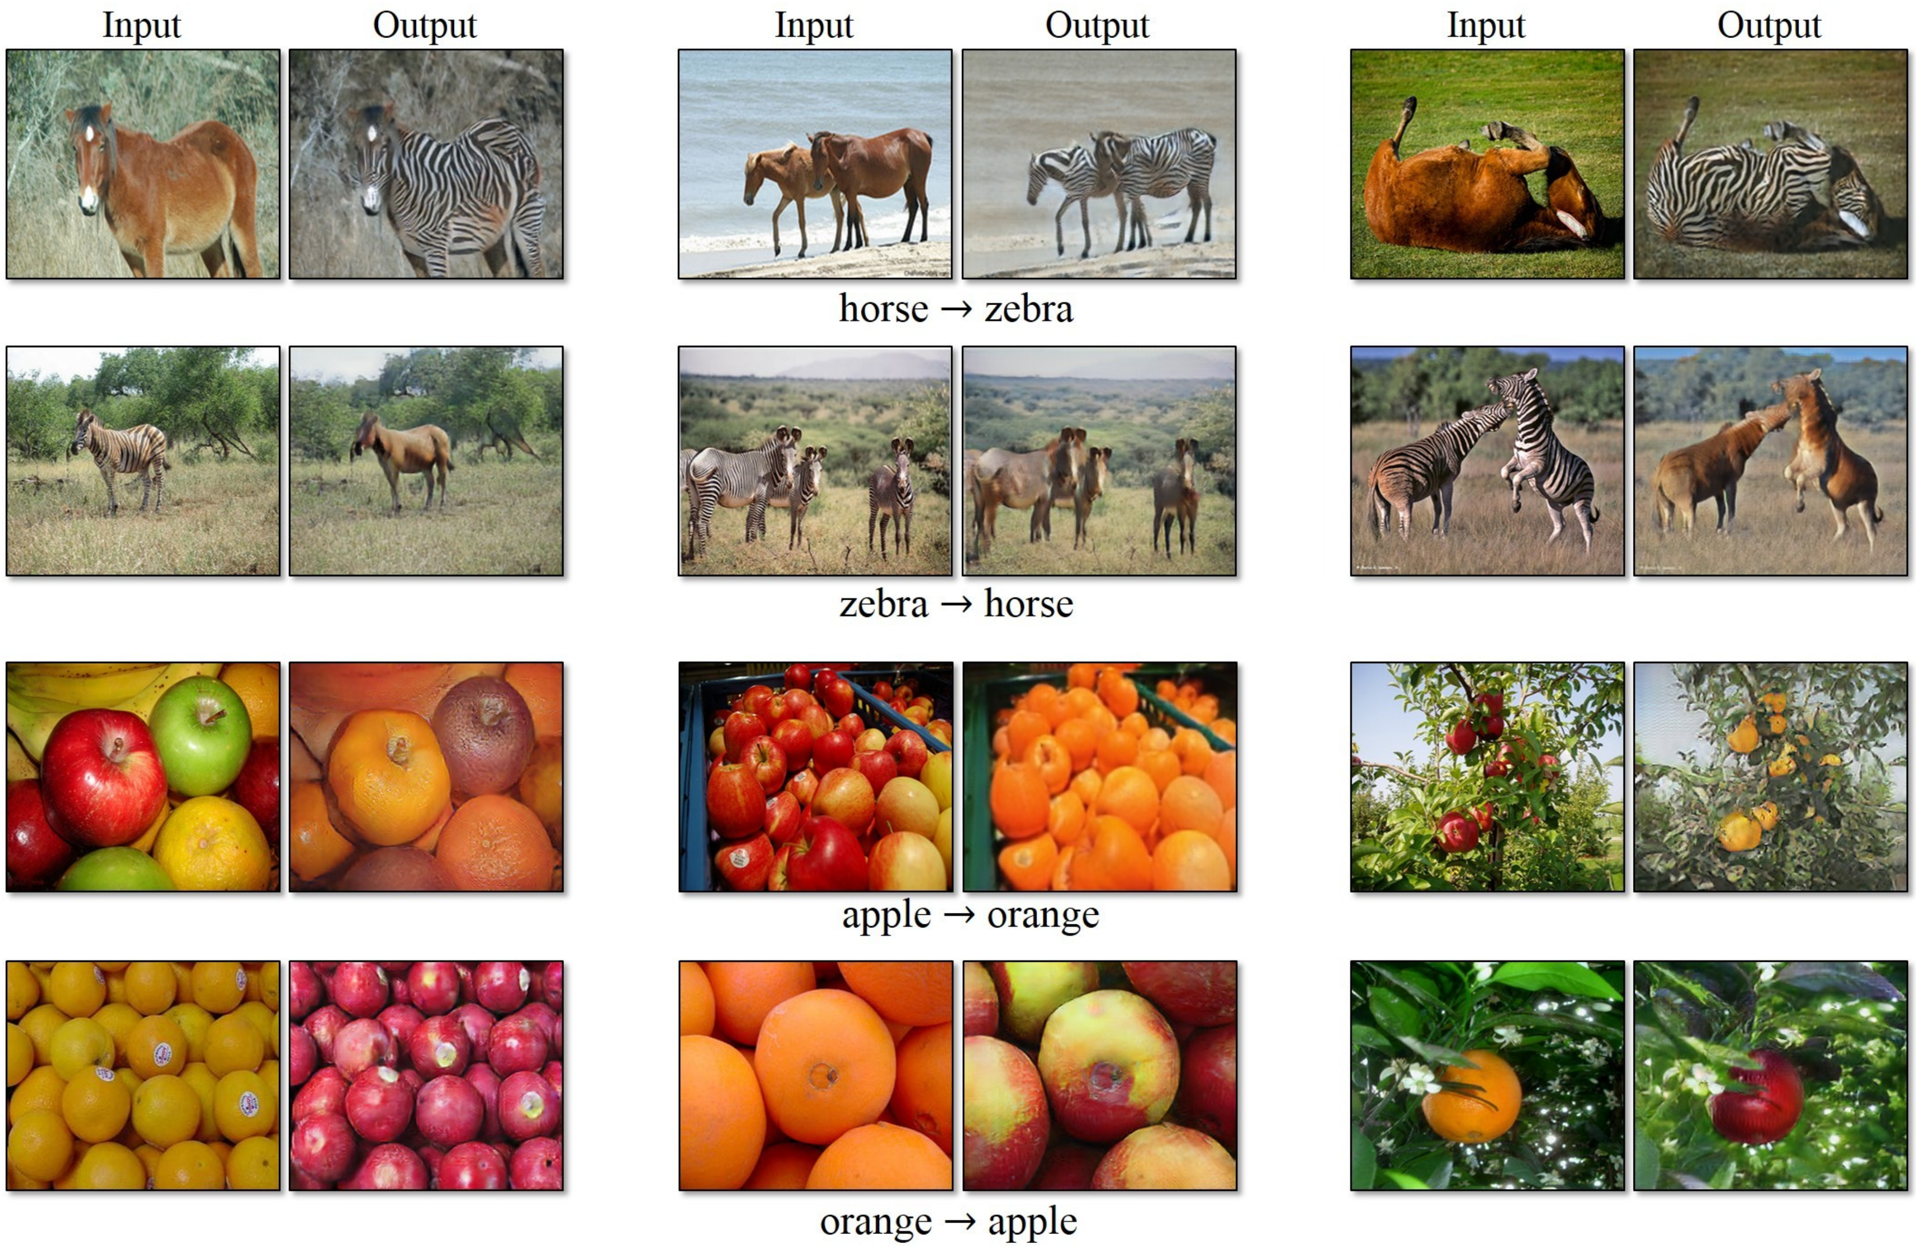
Input
Output
Input
Output
Input
Output
horse → zebra
zebra → horse
apple → orange
orange → apple

winter Yosemite → summer Yosemite

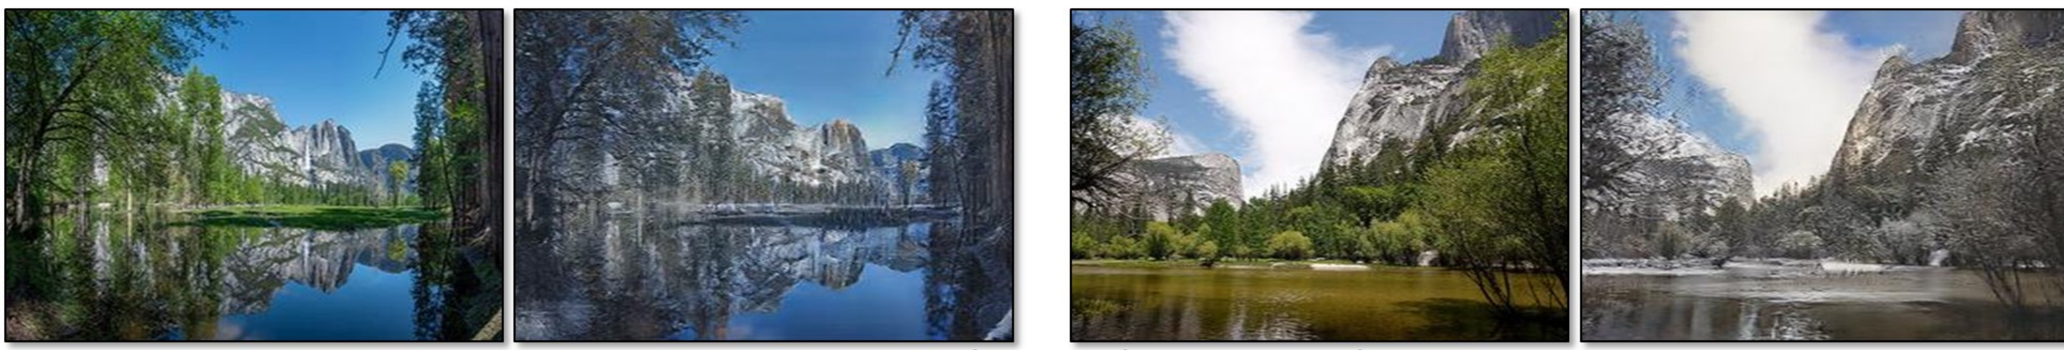

summer Yosemite → winter Yosemite

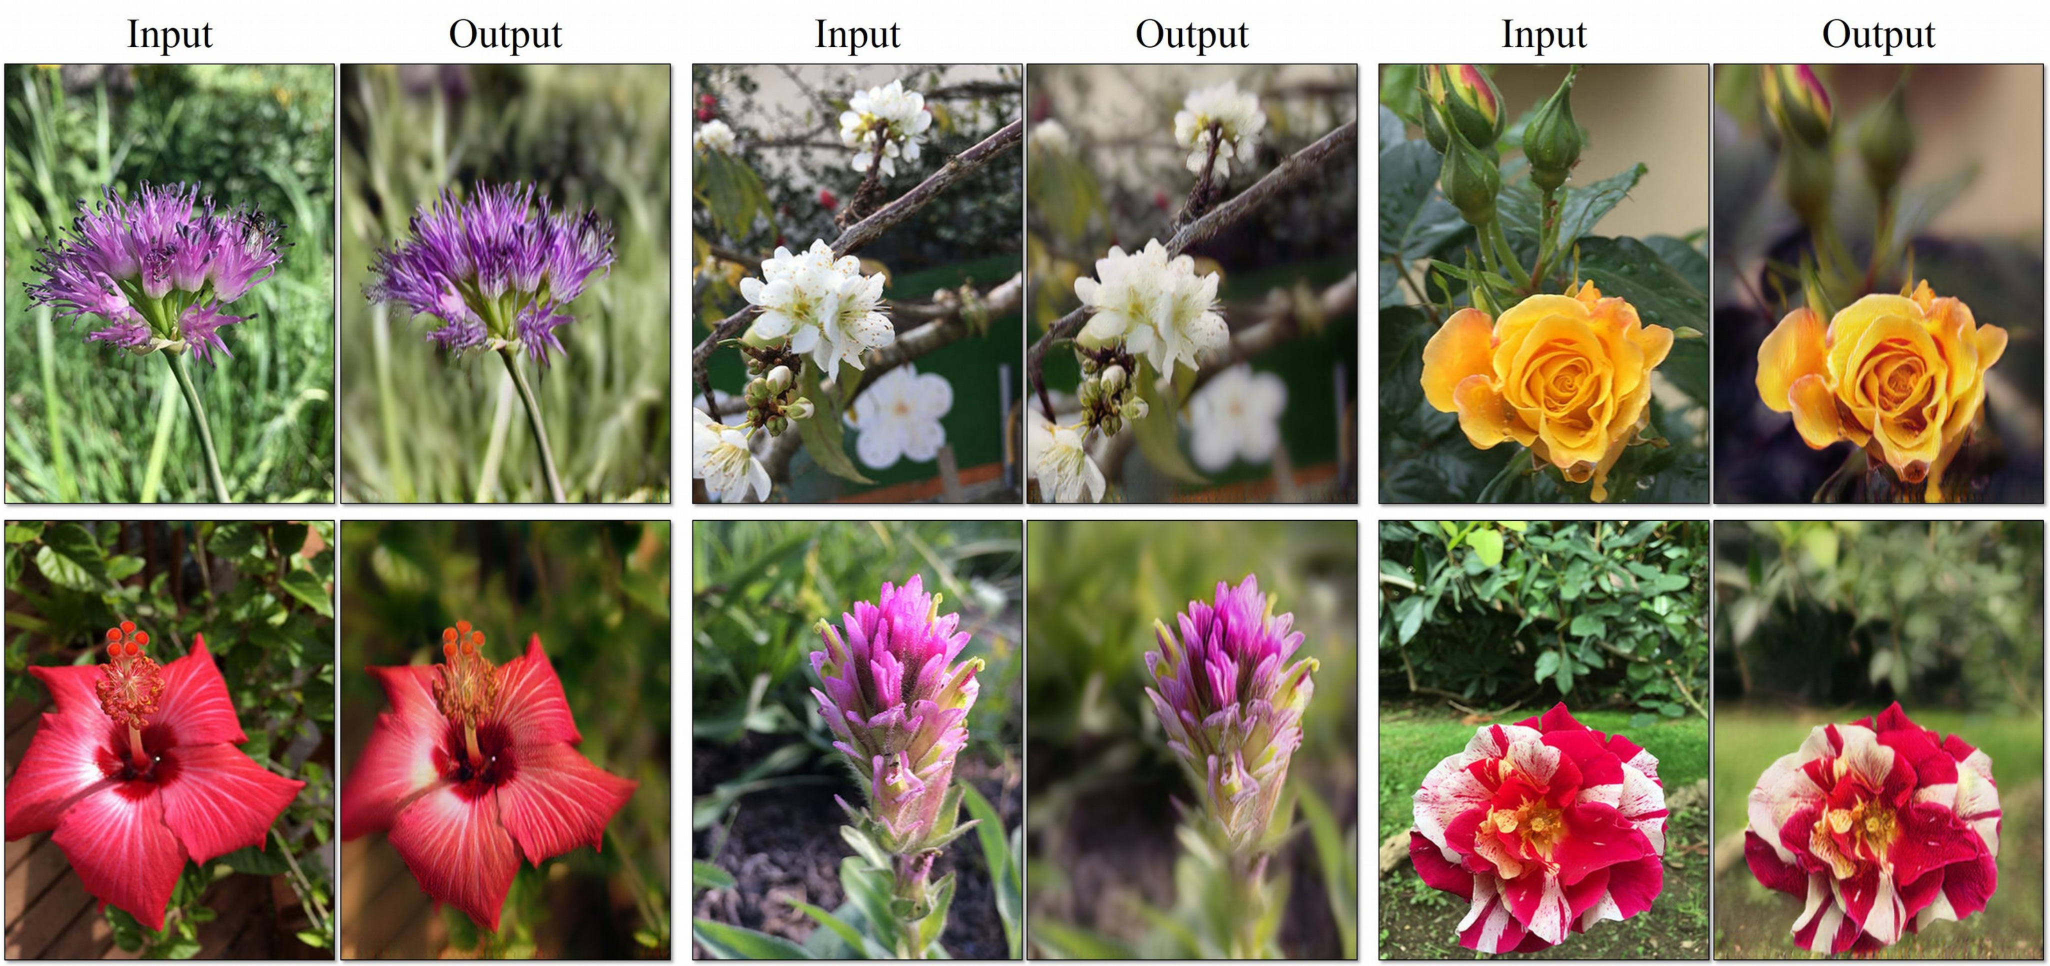
Input
Output
Input
Output
Input
Output

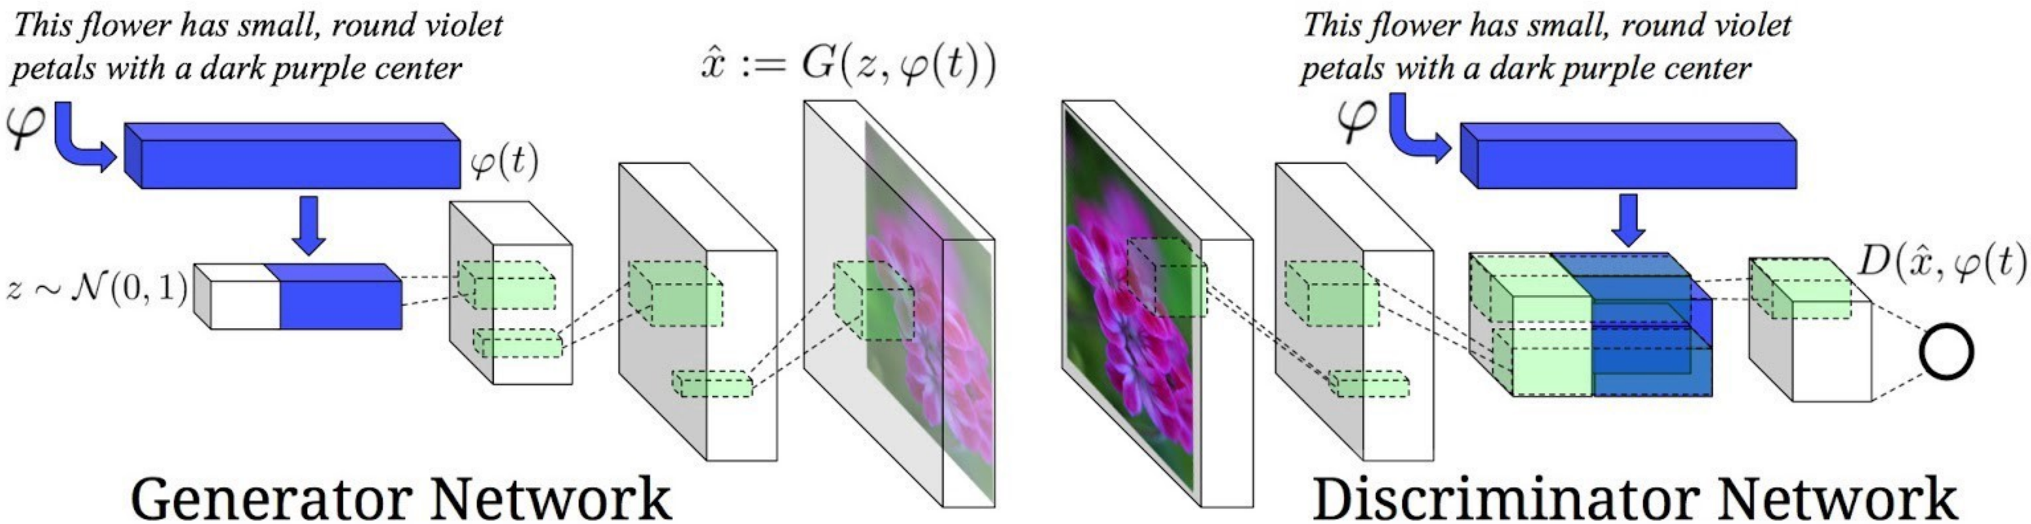

Generative Adversarial Text to Image Synthesis

Scott Reed, Zeynep Akata, Xinchen Yan, Lajanugen Logeswaran REEDSCOT[1], AKATA[2], XCYAN[1], LLAJAN[1]
Bernt Schiele, Honglak Lee SCHIELE[2],HONGLAK[1]

[1] University of Michigan, Ann Arbor, MI, USA (UMICH.EDU)
[2] Max Planck Institute for Informatics, Saarbrücken, Germany (MPI-INF.MPG.DE)

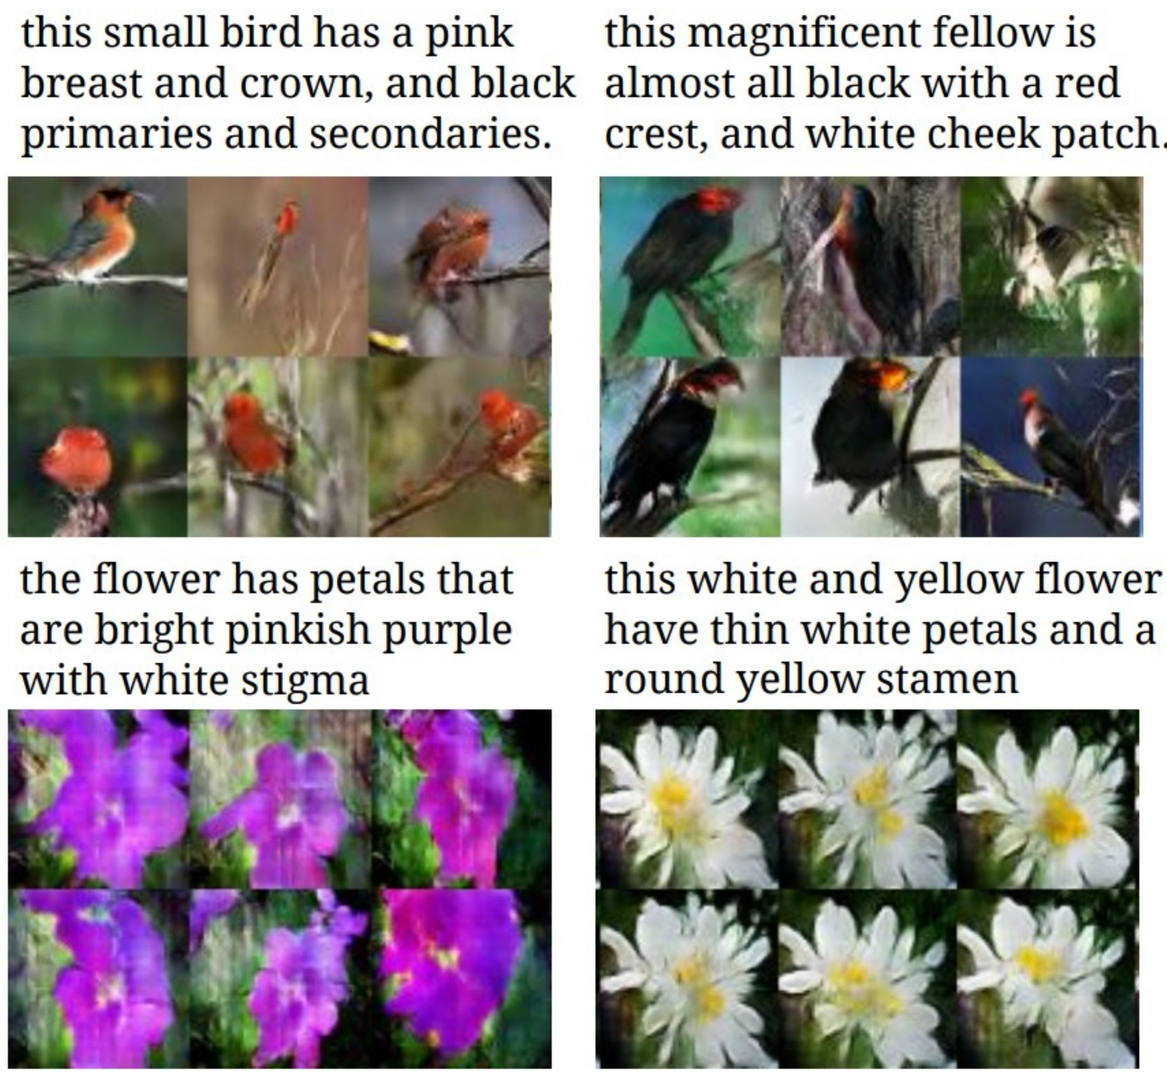

Generative Adversarial Text to Image Synthesis

Scott Reed, Zeynep Akata, Xinchen Yan, Lajanugen Logeswaran REEDSCOT[1], AKATA[2], XCYAN[1], LLAJAN[1]
Bernt Schiele, Honglak Lee SCHIELE[2],HONGLAK[1]
[1] University of Michigan, Ann Arbor, MI, USA (UMICH.EDU)
[2] Max Planck Institute for Informatics, Saarbrücken, Germany (MPI-INF.MPG.DE)

learning word embeddings with word2vec

AUDIO

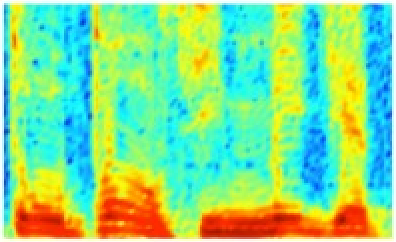

Audio Spectrogram

DENSE

IMAGES

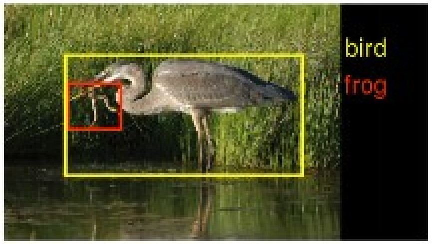

Image pixels

DENSE

TEXT

0	0	0	0.2	0	0.7	0	0	0	...	...

Word, context, or
document vectors

SPARSE

we want to learn to associate nearby words

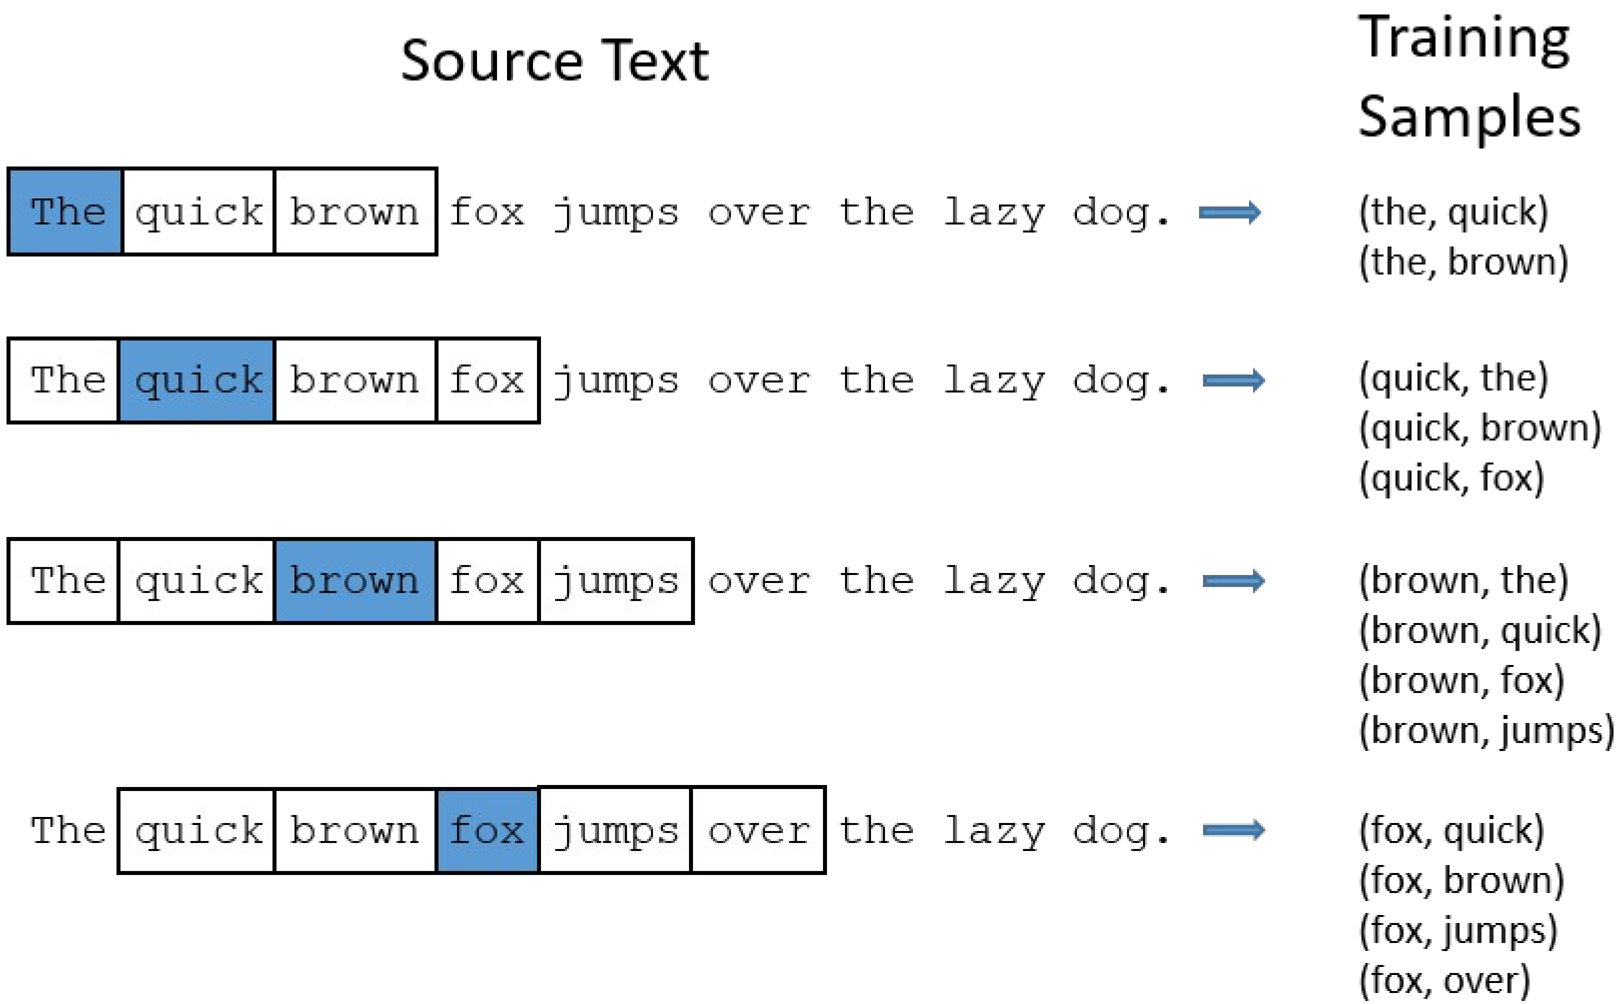

we can predict context around a given word

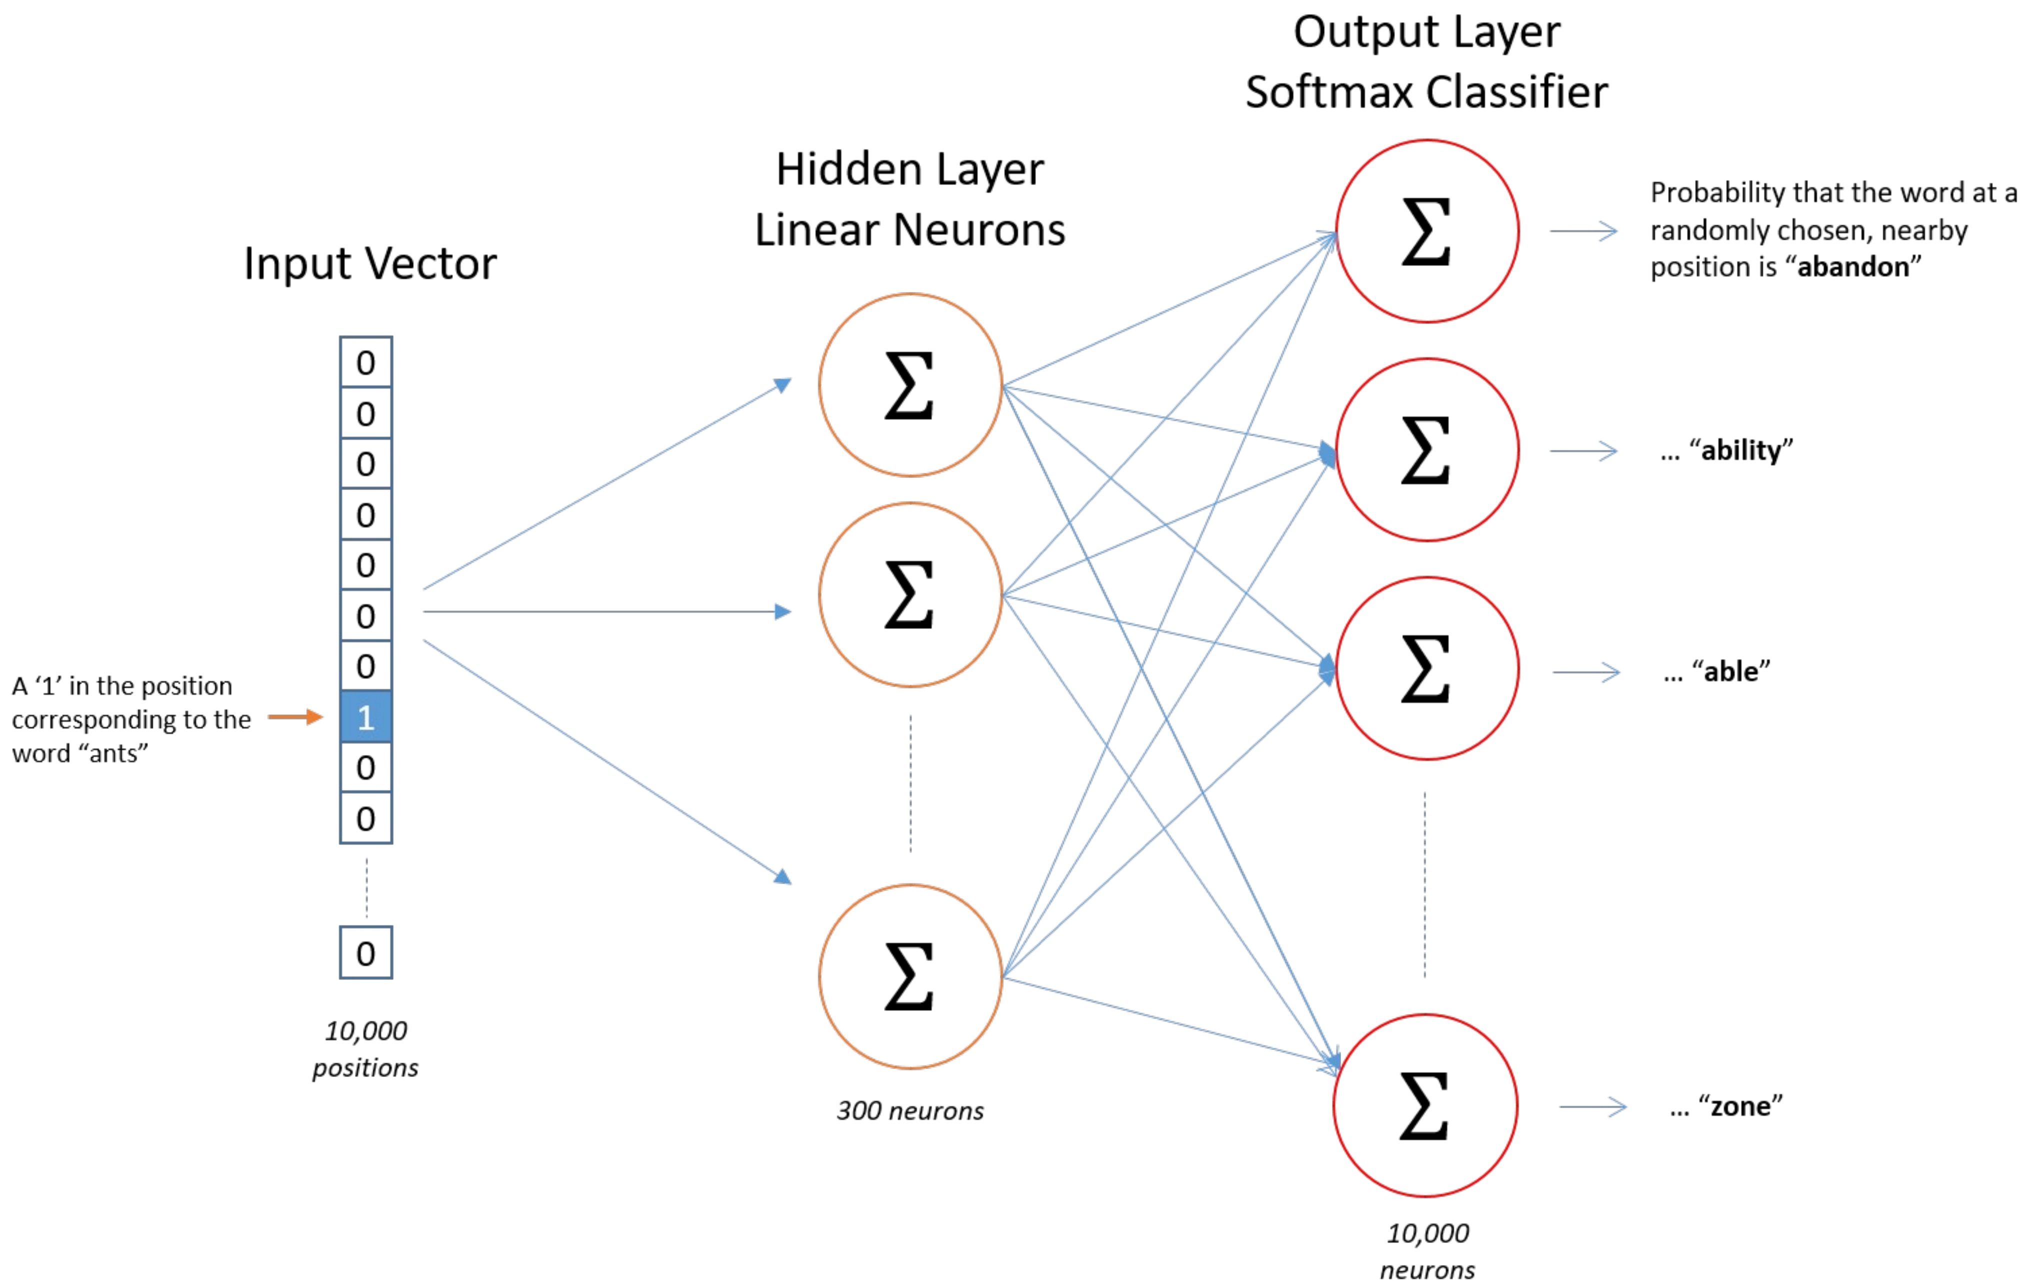

or predict the word, given it's surrounding context

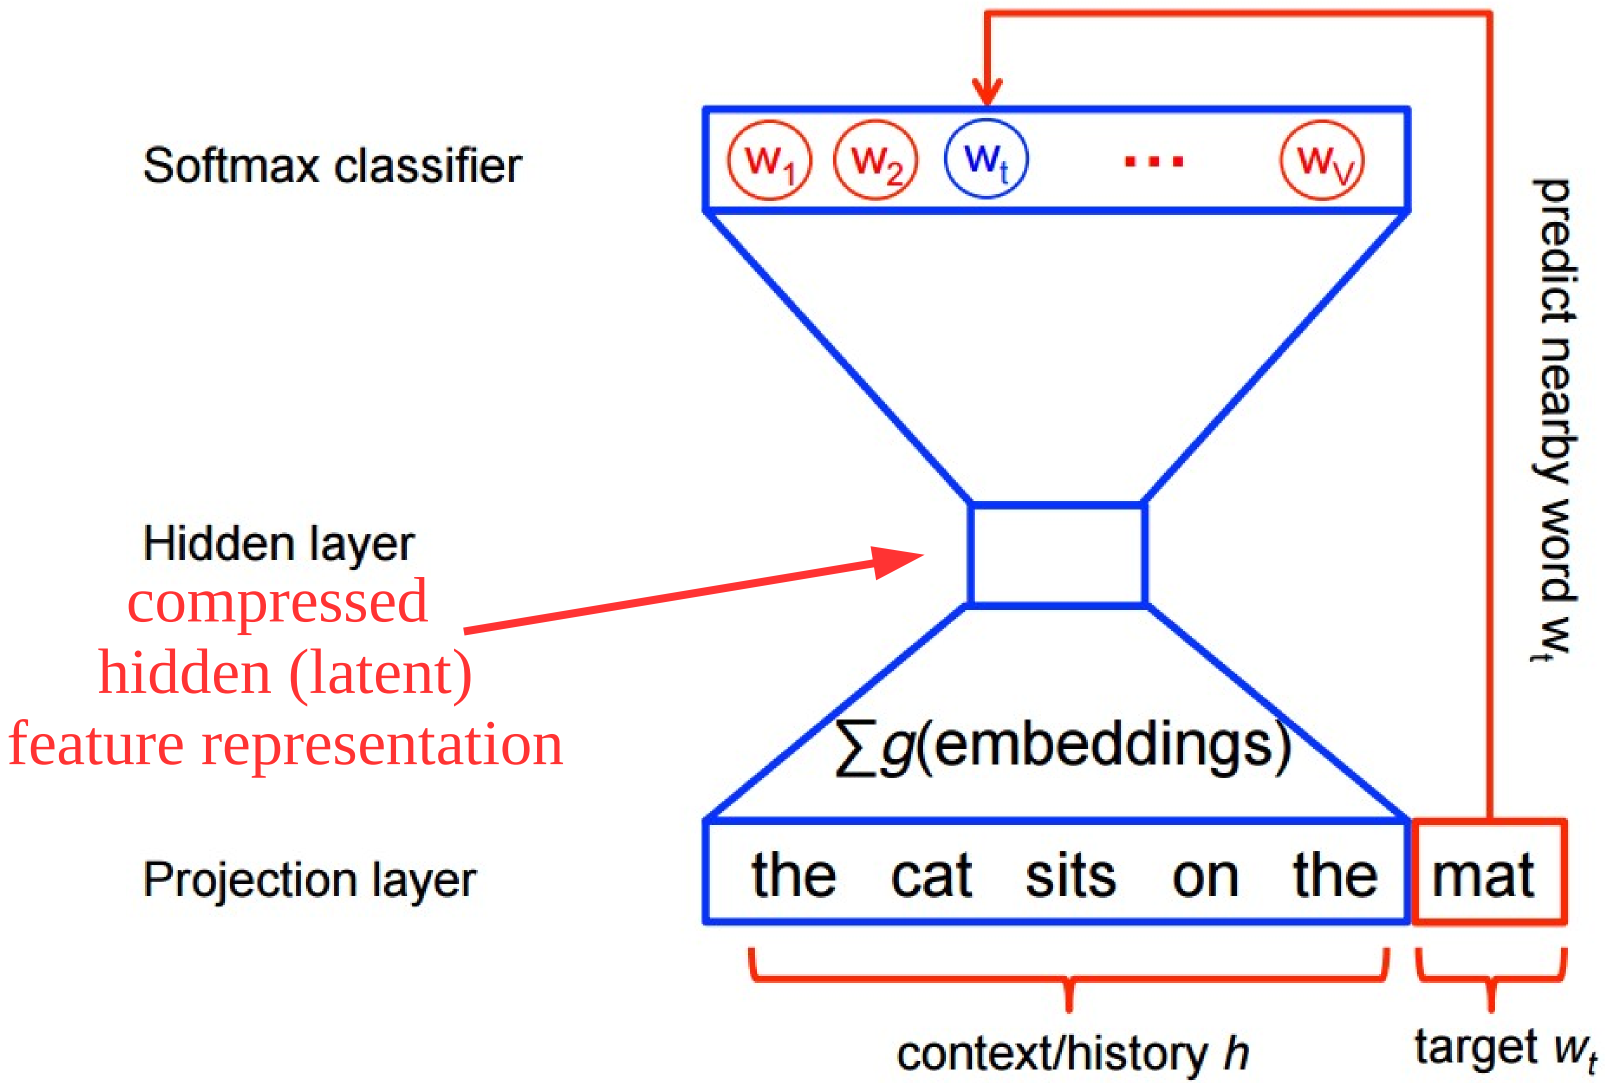

if we plot the latent representation vectors
in a lower dimensional space...

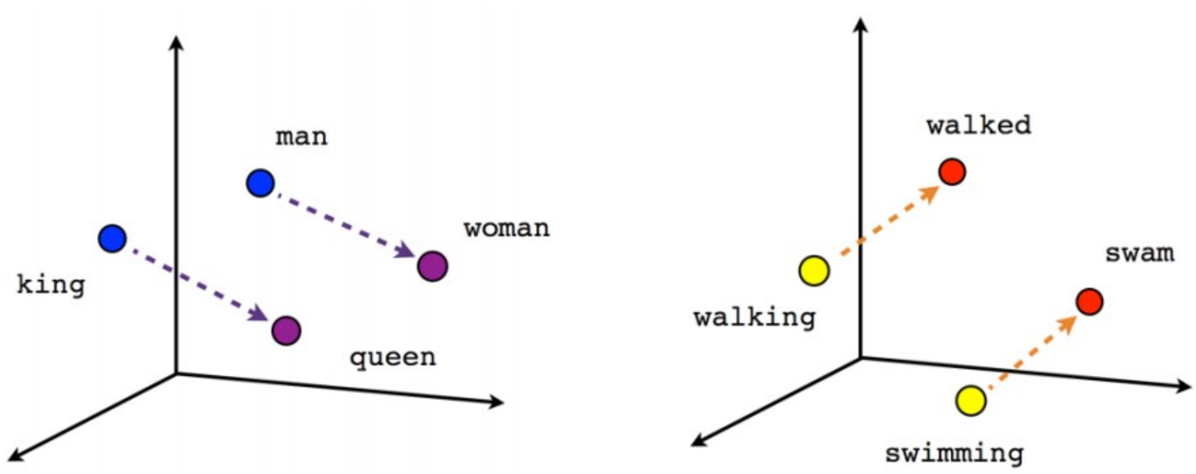

Male-Female

Verb tense

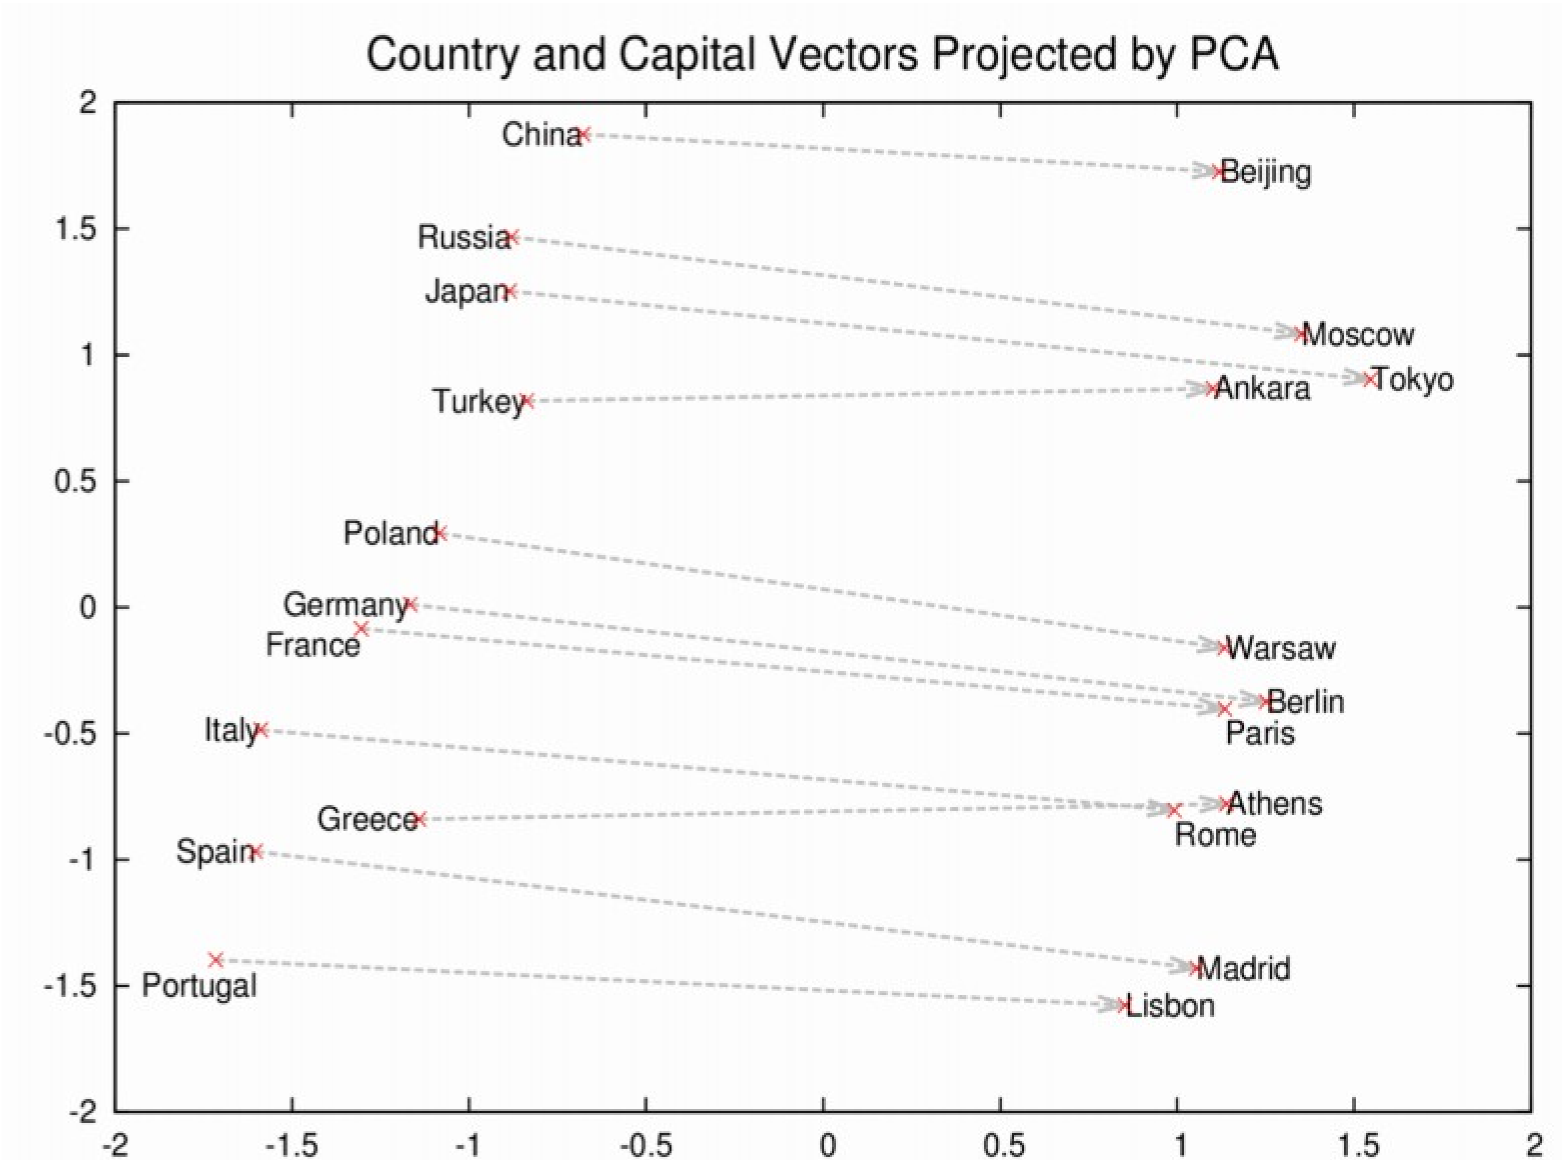
Country and Capital Vectors Projected by PCA
China
Beijing
Russia
Japan
Moscow
Tokyo
Turkey
Ankara
Poland
Germany
France
Warsaw
Berlin
Paris
Italy
Greece
Spain
Athens
Rome
Madrid
Portugal
Lisbon
2
1.5
1
0.5
0
-0.5
-1
-1.5
-2
-2
-1.5
-1
-0.5
0
0.5
1
1.5
2

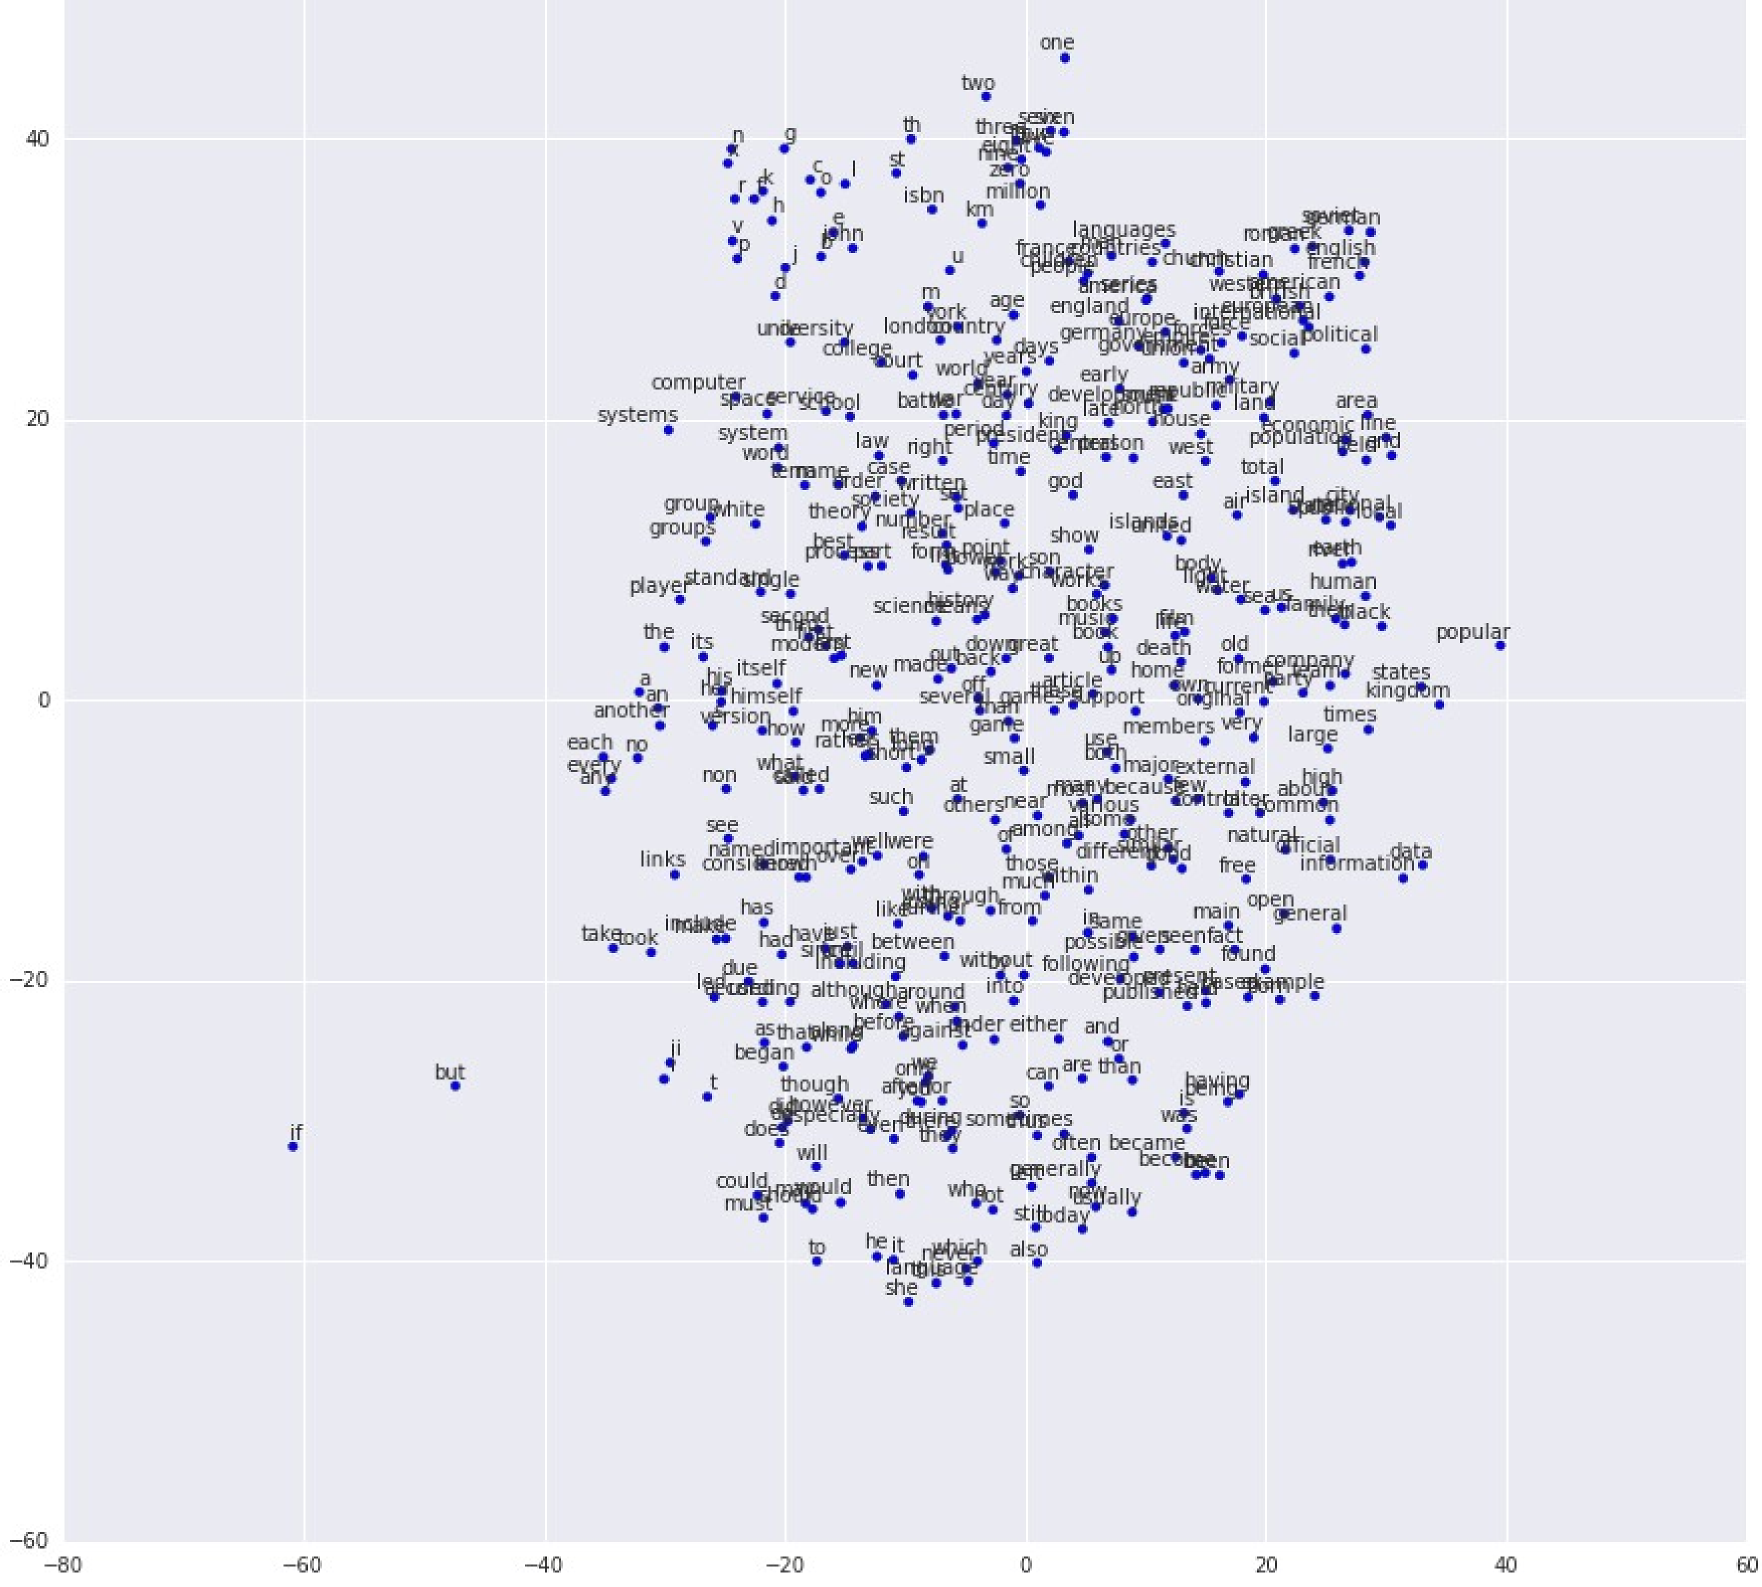

40
20
0
−20
−40
−60
−80
−60
−40
−20
0
20
40
60
one
two
th
st
isbn
km
million
zero
nine
n
g
c
o
l
r
h
e
john
v
p
j
d
u
m
age
languages
england
europe
german
english
french
christian
american
british
social
political
international
college
court
world
computer
systems
system
word
law
right
time
god
king
west
east
area
population
total
island
city
group
groups
white
theory
society
number
place
show
islands
player
standard
single
second
science
history
books
music
book
death
life
old
former
company
human
popular
states
kingdom
the
its
a
an
another
his
itself
himself
version
new
made
back
article
support
members
very
large
times
each
every
any
no
non
what
how
such
at
others
near
among
use
both
major
external
high
about
common
natural
official
data
information
free
see
links
named
considered
well
were
on
those
much
within
from
with
through
main
open
general
found
has
take
took
include
make
had
have
just
between
without
about
following
possible
into
due
led
although
around
where
when
either
and
or
as
before
against
ii
but
if
began
t
can
are
than
though
however
so
often
became
having
is
was
will
does
could
must
should
would
then
who
not
generally
usually
still
today
to
he
it
which
never
also
she

where `:` means "is to" and `::` means "as"; e.g. "Rome is to Italy as Beijing is to China" = `Rome:Italy::Beijing:China`

amusing word2vec analogies

```
king:queen::man:[woman, Attempted abduction, teenager, girl]
//Weird, but you can kind of see it

China:Taiwan::Russia:[Ukraine, Moscow, Moldova, Armenia]
//Two large countries and their small, estranged neighbors

house:roof::castle:[dome, bell_tower, spire, crenellations, turrets]

knee:leg::elbow:[forearm, arm, ulna_bone]

New York Times:Sulzberger::Fox:[Murdoch, Chernin, Bancroft, Ailes]
//The Sulzberger-Ochs family owns and runs the NYT.
//The Murdoch family owns News Corp., which owns Fox News.
//Peter Chernin was News Corp.'s COO for 13 yrs.
//Roger Ailes is president of Fox News.
//The Bancroft family sold the Wall St. Journal to News Corp.

love:indifference::fear:[apathy, callousness, timidity, helplessness, inaction]
//the poetry of this single array is simply amazing...

Donald Trump:Republican::Barack Obama:[Democratic, GOP, Democrats, McCain]
//It's interesting to note that, just as Obama and McCain were rivals,
//so too, Word2vec thinks Trump has a rivalry with the idea Republican.

monkey:human::dinosaur:[fossil, fossilized, Ice_Age_mammals, fossilization]
//Humans are fossilized monkeys? Humans are what's left
//over from monkeys? Humans are the species that beat monkeys
//just as Ice Age mammals beat dinosaurs? Plausible.

building:architect::software:[programmer, SecurityCenter, WinPcap]
```

Amusing Word2Vec Results

- Geopolitics: *Iraq - Violence = Jordan*
- Distinction: *Human - Animal = Ethics*
- *President - Power = Prime Minister*
- *Library - Books = Hall*
- Analogy: *Stock Market ≈ Thermometer*

associative unsupervised learning

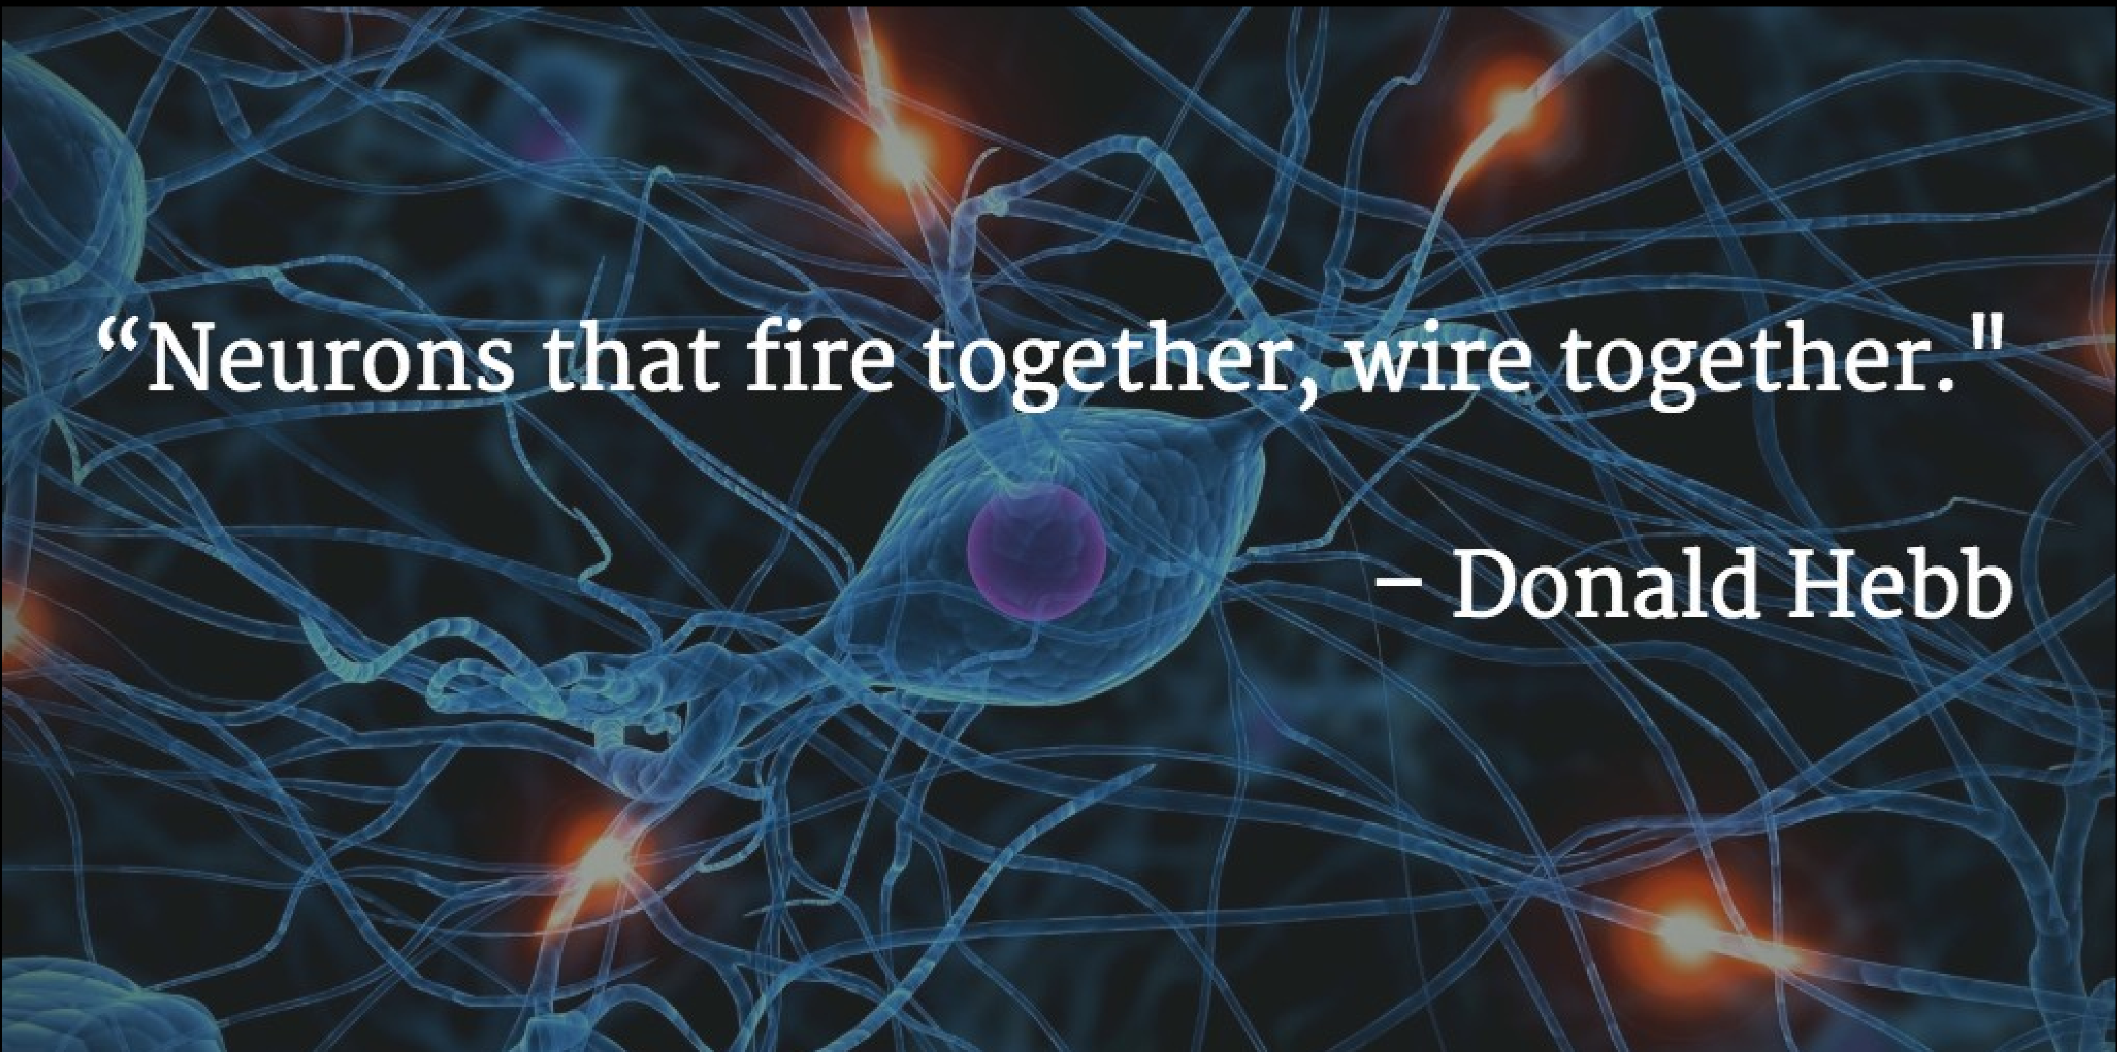
“Neurons that fire together, wire together."
– Donald Hebb

Hebbian learning

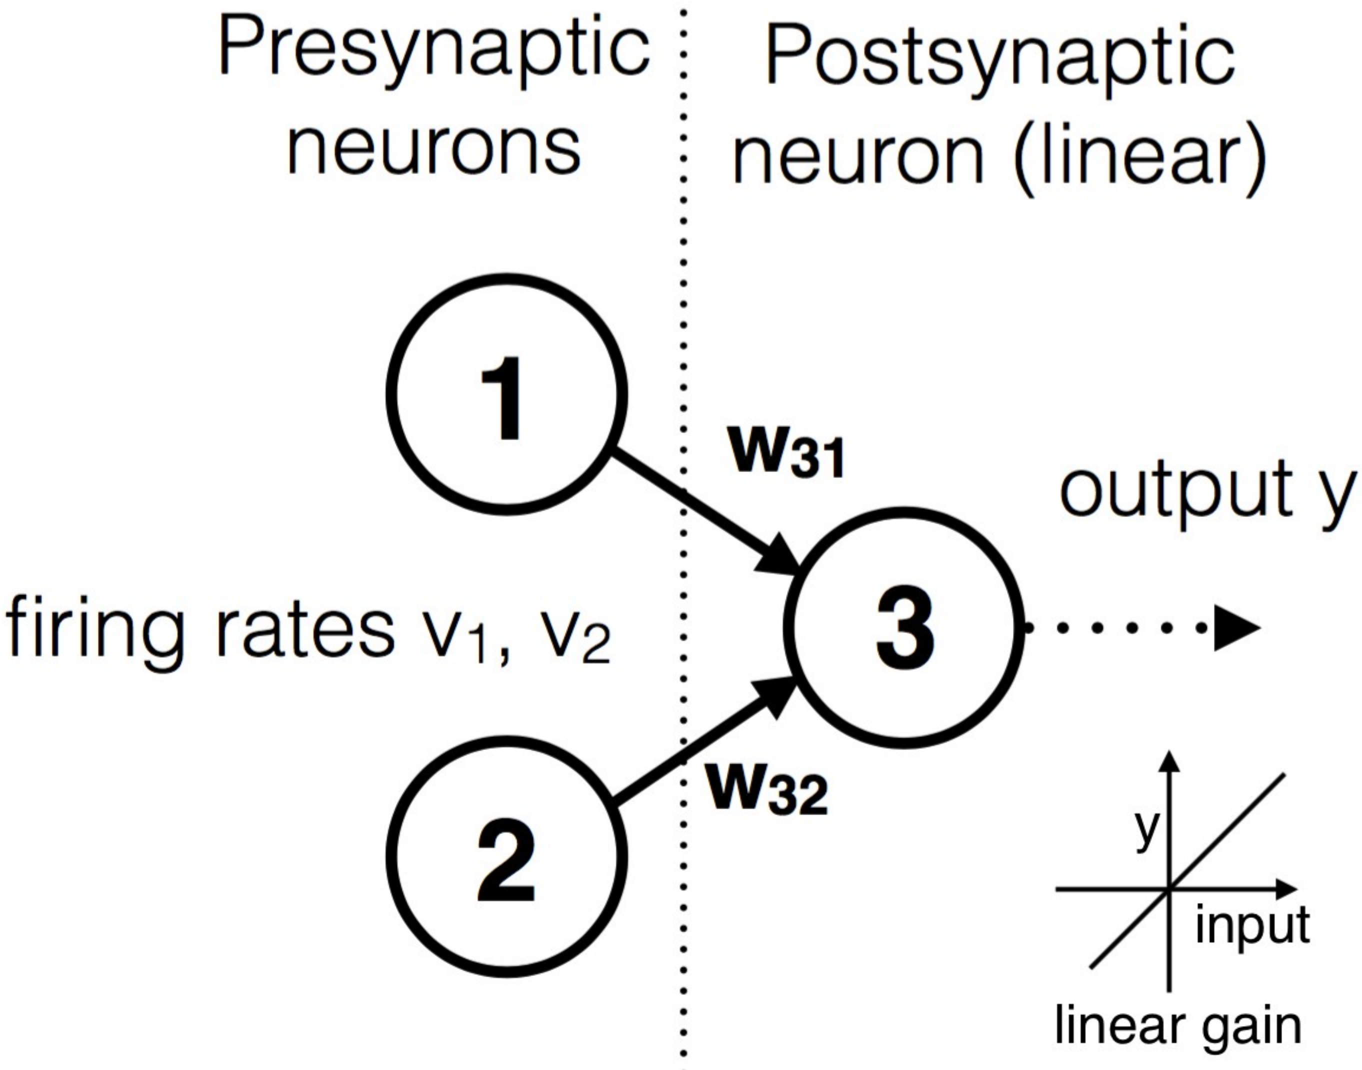

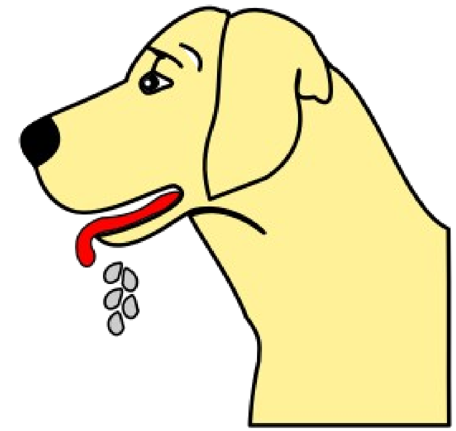

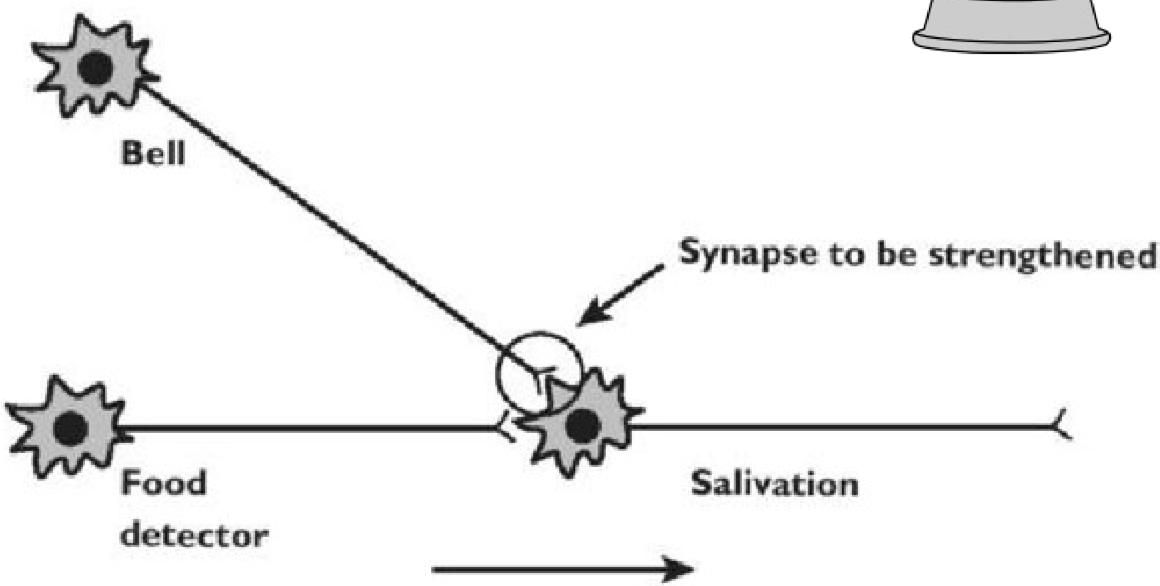
Bell
Synapse to be strengthened
Food
detector
Salivation

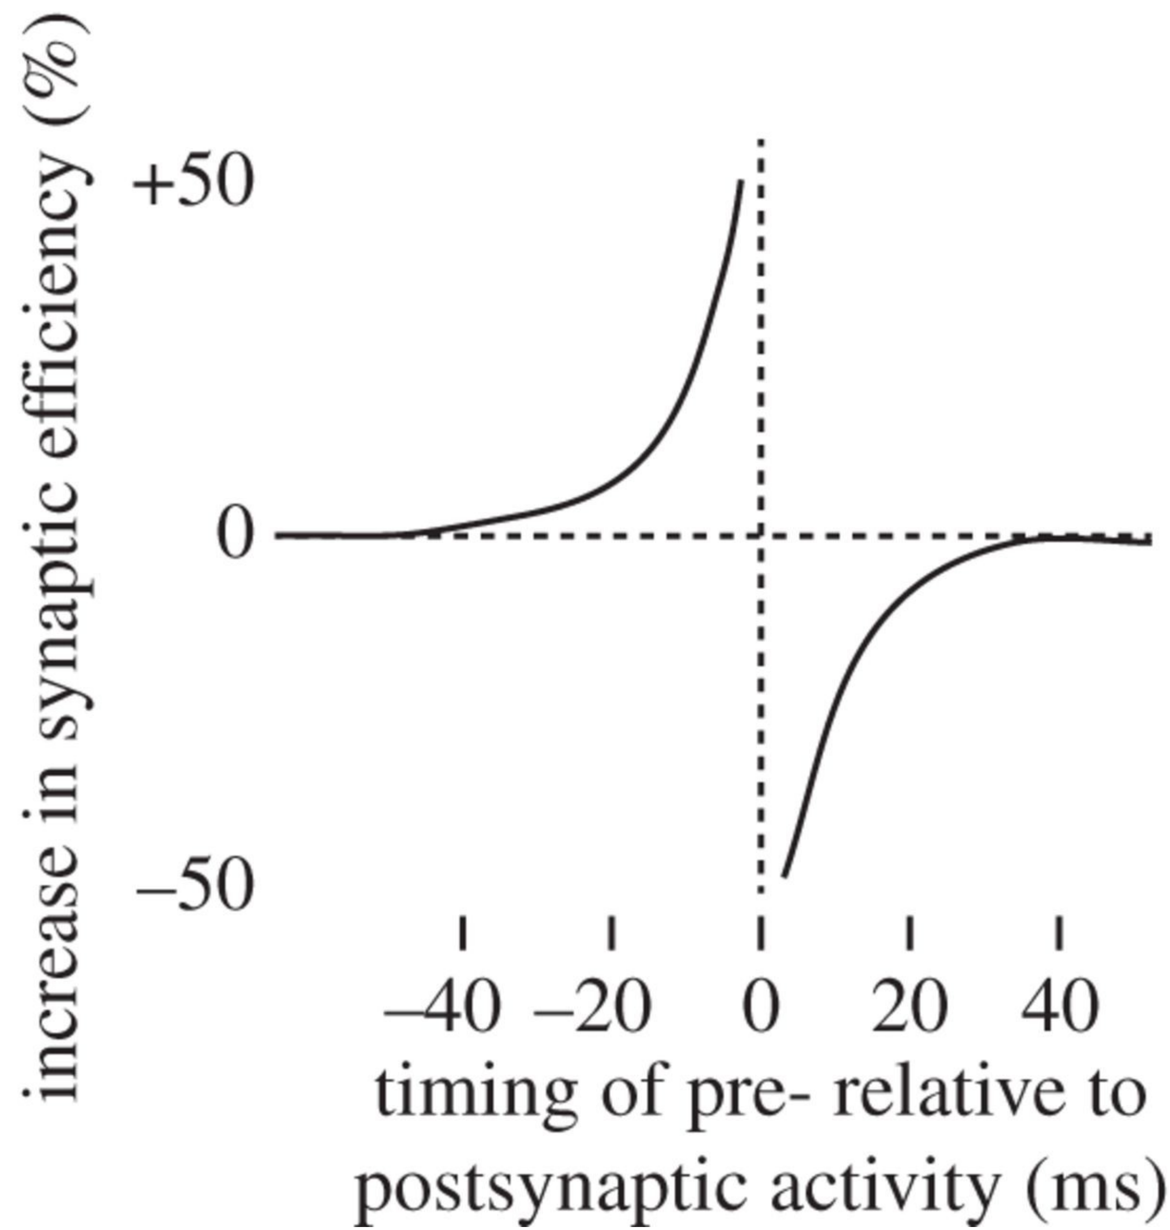
increase in synaptic efficiency (%)
+50
0
–50
–40
–20
0
20
40
timing of pre- relative to postsynaptic activity (ms)

auto-associative learning
(Hopfield network)

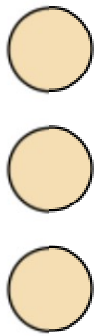

auto-associative learning
(Hopfield network)

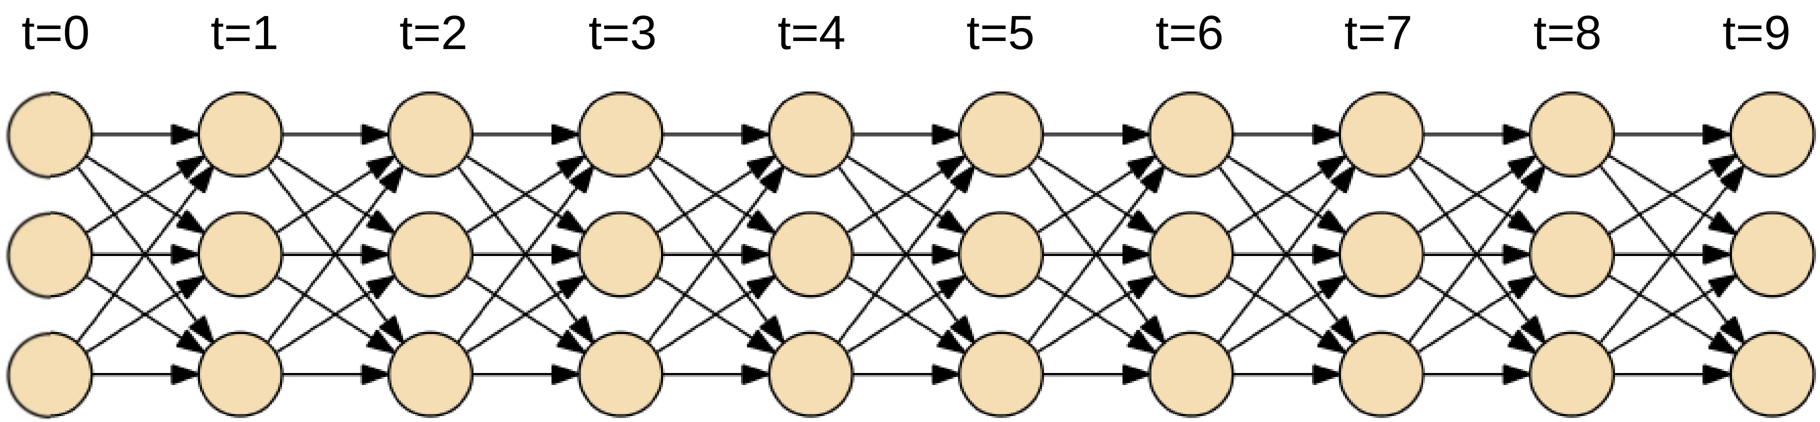

Hopfield Network

- Recurrent network
  - Feedback from output to input
- Fully connected
  - Every neuron connected to every other neuron

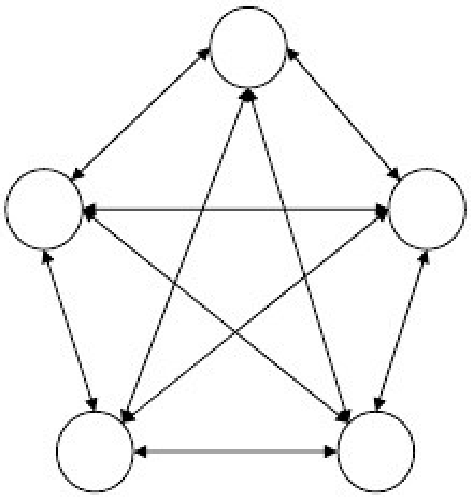

Hopfield Network

- Recurrent network
  - Feedback from output to input
- Fully connected
  - Every neuron connected to every other neuron

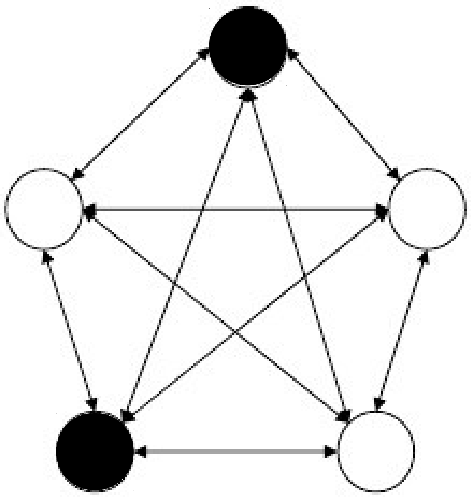

Hopfield Network

- Recurrent network
  - Feedback from output to input
- Fully connected
  - Every neuron connected to every other neuron

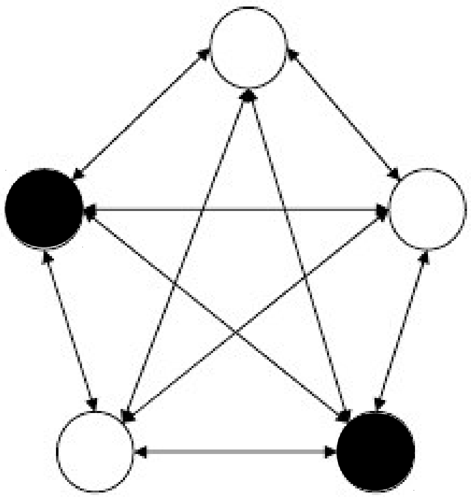

Hopfield Network

- Recurrent network
  - Feedback from output to input
- Fully connected
  - Every neuron connected to every other neuron

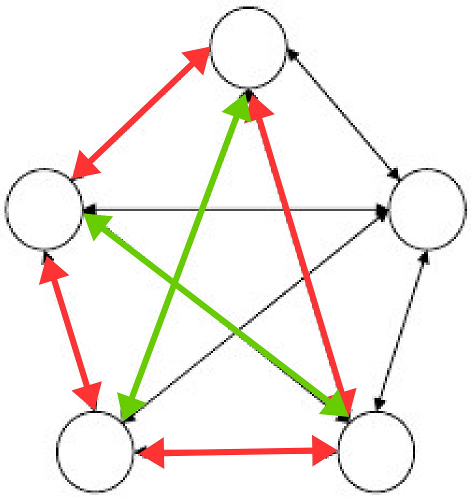

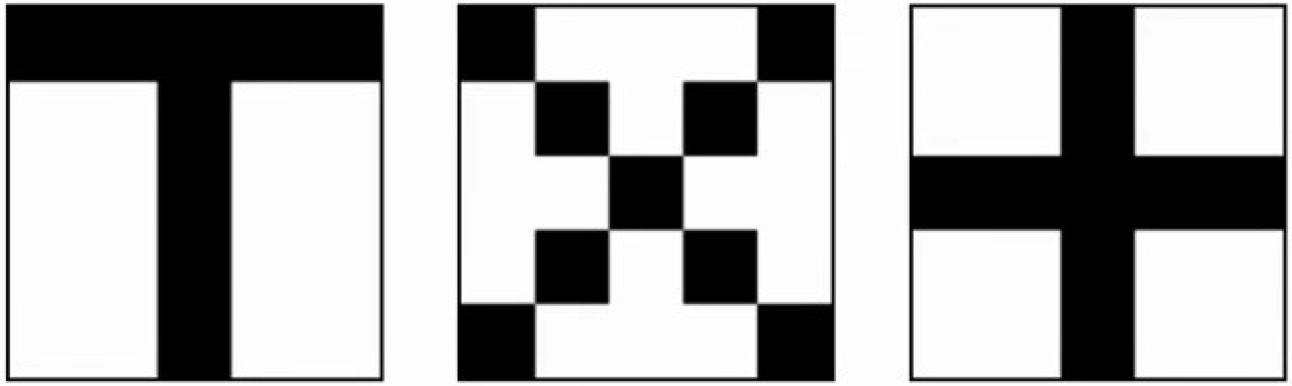

Start Running
Train Network

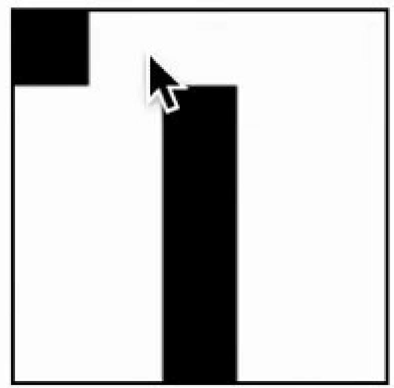

unsupervised clustering

simple, but widely helpful

there are lots of types of clustering algorithms
(e.g. hierarchical clustering, self organizing maps)
but the one we will introduce here is
K-means clustering, because it is the simplest
(and it's surprisingly effective in practice)

Step 1:
choose a number of classes (k),
and randomly assign all points to one of the classes

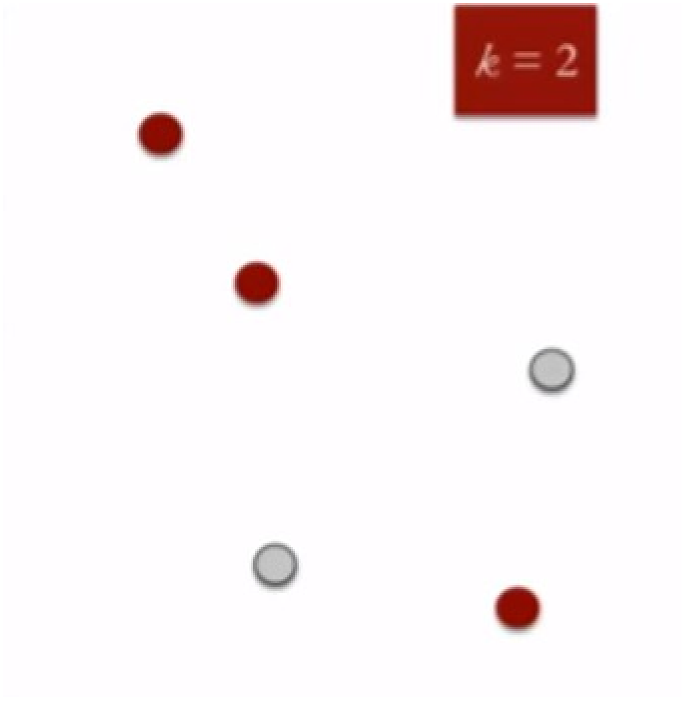

Step 2:
average the position of all of the points in a given class to find that class' centroid (mean features representation)

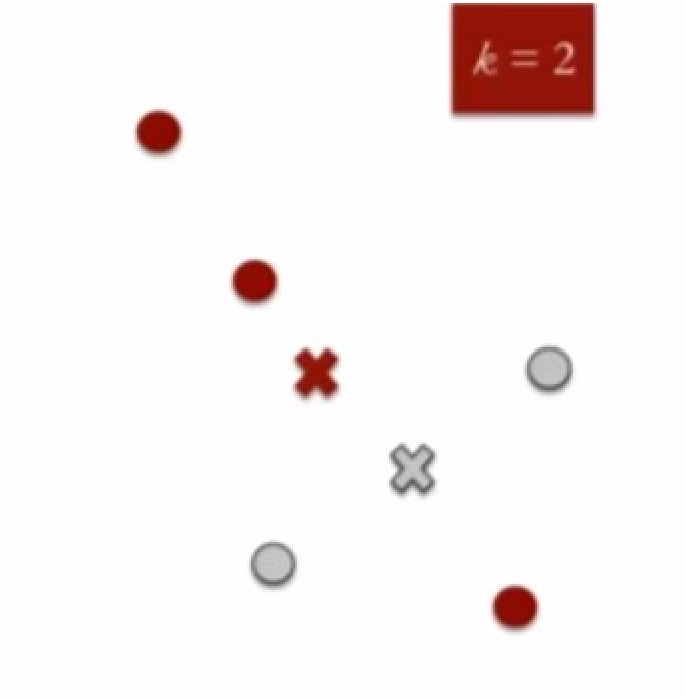

Step 3:
re-assign each point to the closest centroid

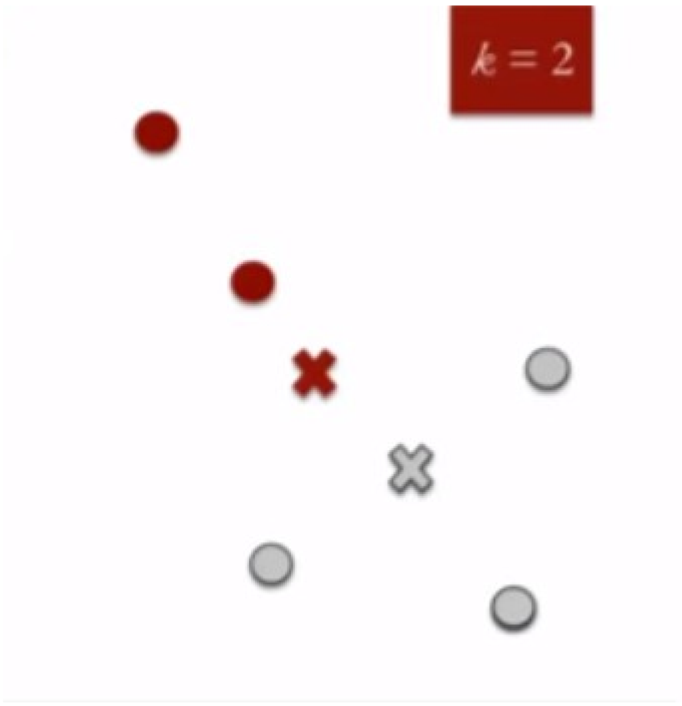

Step …
repeat centroid and point assignment

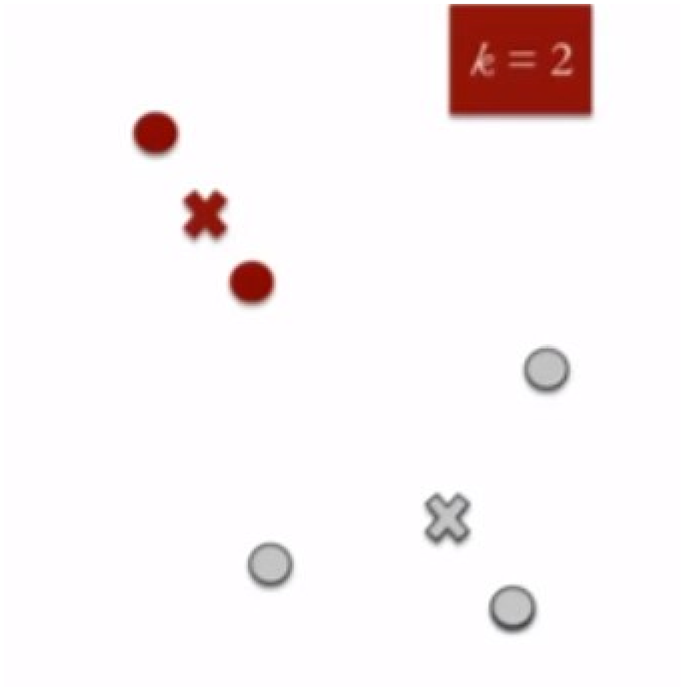

Step …
repeat centroid and point assignment
until a steady state is reached (no assignments change)

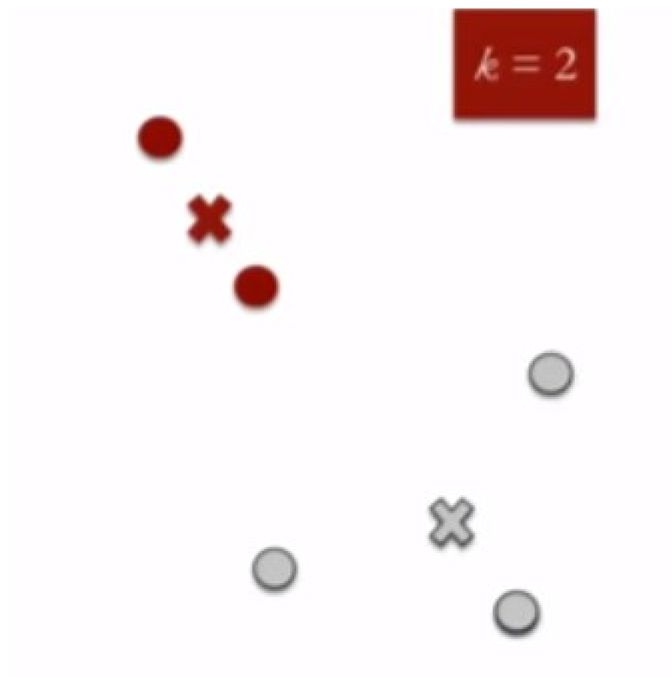

you now have the optimal greedy point assignments for that number of clusters

but you did have to choose the number of clusters, which may not be obvious

how many clusters are here?

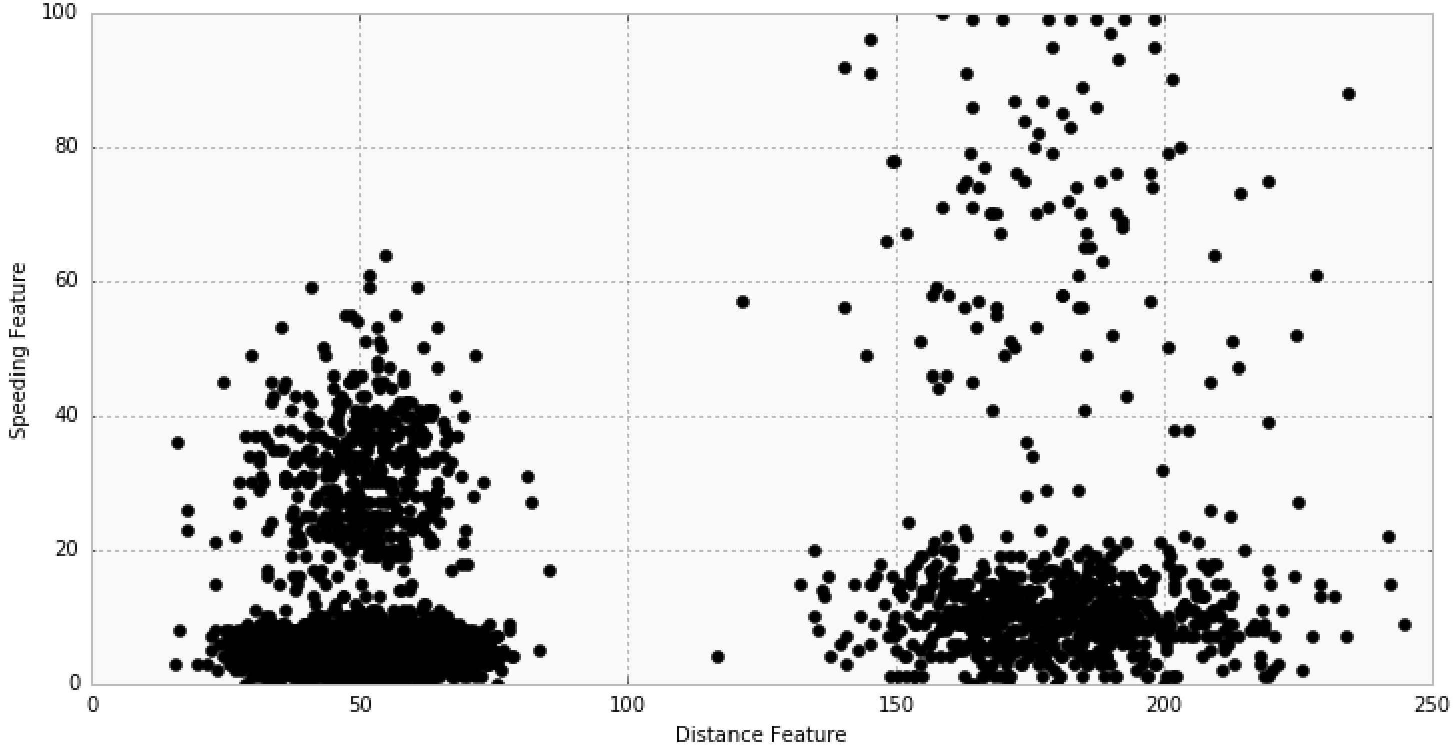

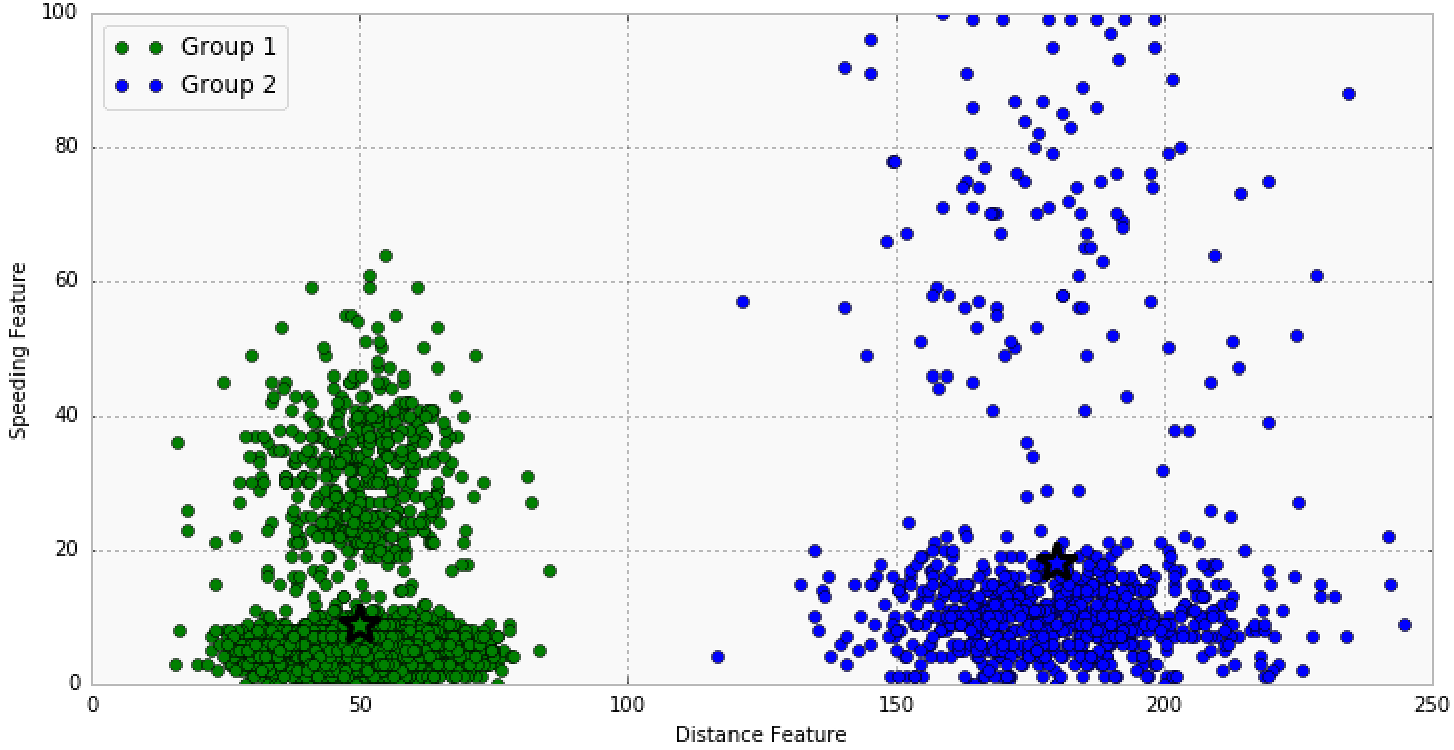

Group 1
Group 2
Speeding Feature
Distance Feature
0
20
40
60
80
100
0
50
100
150
200
250

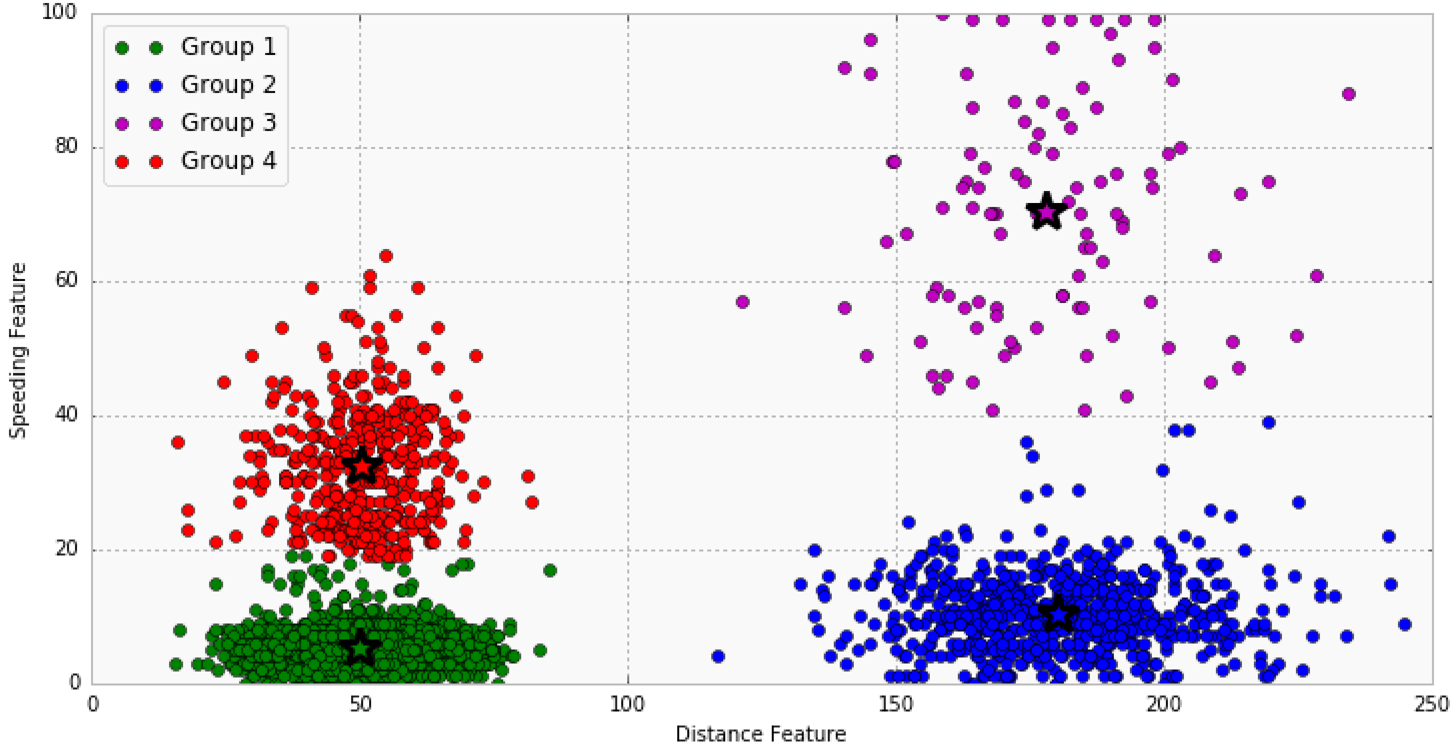

Group 1
Group 2
Group 3
Group 4
Speeding Feature
Distance Feature
0
20
40
60
80
100
0
50
100
150
200
250

k-means is cheap and fast, so let's try a range of values, and look for a good trade-offs between classification and number of classes

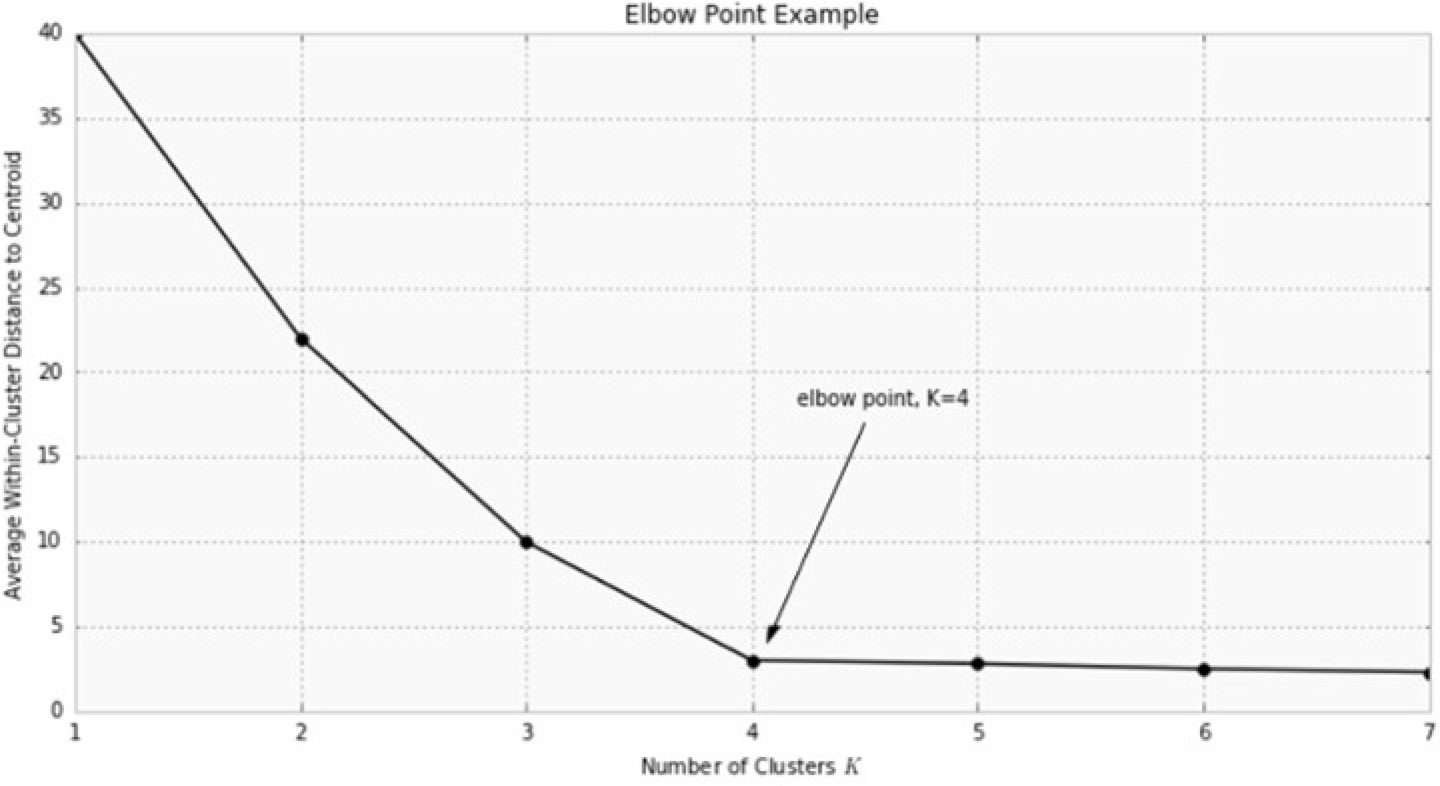

since we are doing arithmetic in the feature space, having good features and preprocessing your data is extremely important to your outcomes (just like it was in KNN)

but we'll cover this after break…