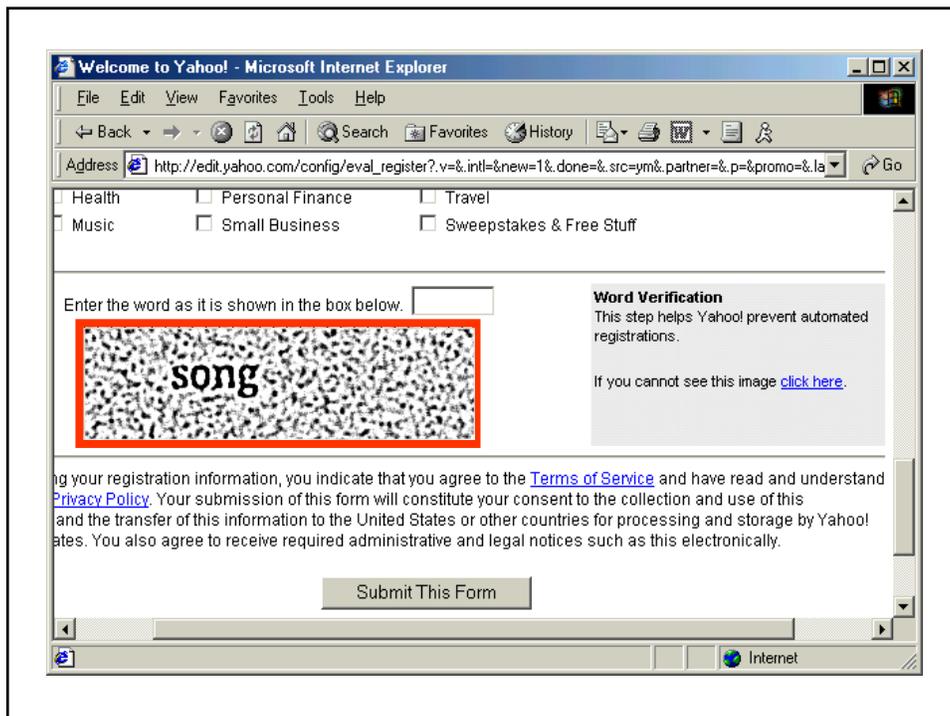


# CAPTCHA: Telling Humans and Computers Apart Automatically

Luis von Ahn  
Manuel Blum  
Nicholas Hopper  
John Langford

School of Computer Science  
Carnegie Mellon University



## CAPTCHA (2000)

A program that can tell whether its user is a [human or a computer](#)

## CAPTCHA

A [program](#) that can generate and grade tests that:

- A. Most humans can pass
- B. Current computer programs cannot pass

Example

Example

Picks a [random string](#)  
of letters

oamg

## Example

Picks a **random string**  
of letters

oamg



Renders the string into a  
**randomly distorted** image



## Example

...and generates a test:



Type the characters that appear in the image

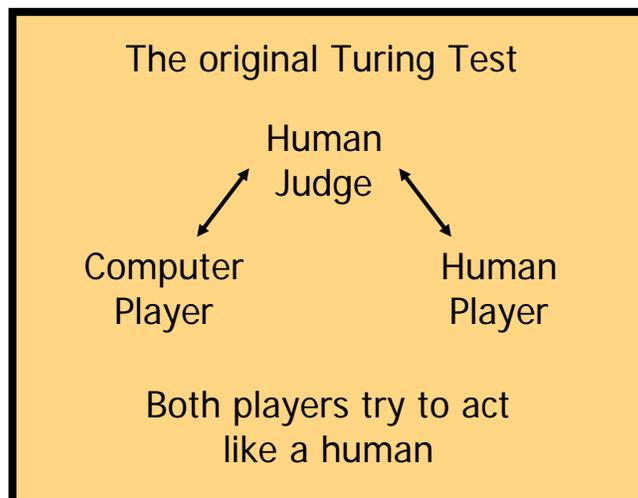
## P stands for Public

All code and data used by a CAPTCHA  
should be publicly available

Completely Automated Public  
Turing Test to Tell Computers and  
Humans Apart

Are CAPTCHAs Reverse Turing Tests?

Are CAPTCHAs Reverse Turing Tests?



Are CAPTCHAs Reverse Turing Tests?

**NO!**

The Rest Of The Talk

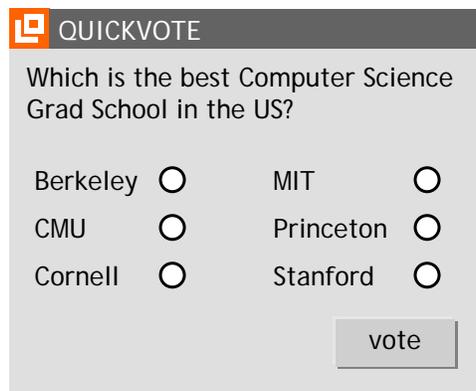
Applications

Examples of CAPTCHAs

Advancing AI

A Spin-off Idea

# Applications



QUICKVOTE

Which is the best Computer Science Grad School in the US?

Berkeley	<input type="radio"/>	MIT	<input type="radio"/>
CMU	<input type="radio"/>	Princeton	<input type="radio"/>
Cornell	<input type="radio"/>	Stanford	<input type="radio"/>

vote

(from [www.slashdot.com](http://www.slashdot.com))

## Applications

Free E-mail Services

Data Collection

Worms and Spam

Preventing Dictionary Attacks  
(Pinkas and Sander '02)



## Applications

Free E-mail Services

Data Collection

Worms and Spam

Preventing Dictionary Attacks  
(Pinkas and Sander '02)



## Applications

Free E-mail Services

Data Collection

Worms and Spam

Preventing Dictionary Attacks  
(Pinkas and Sander '02)



## Applications

Free E-mail Services

Data Collection

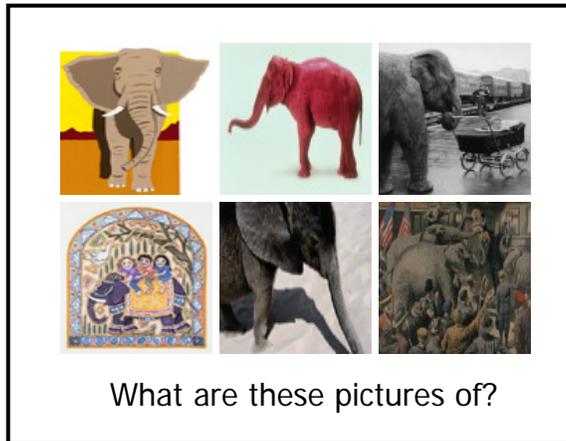
Worms and Spam

Preventing Dictionary Attacks  
(Pinkas and Sander '02)



## Examples of CAPTCHAs

Pix



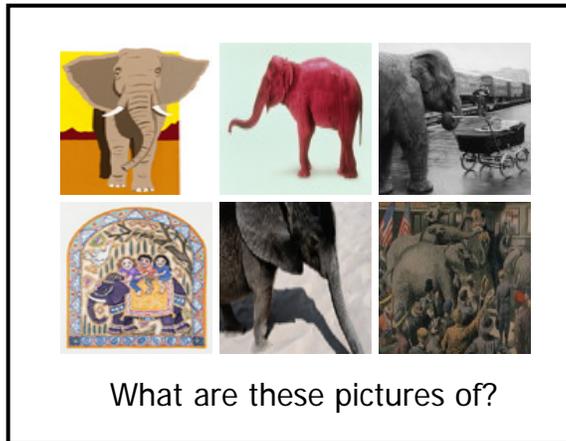
What are these pictures of?

Pix



The images need to be randomly distorted

Pix

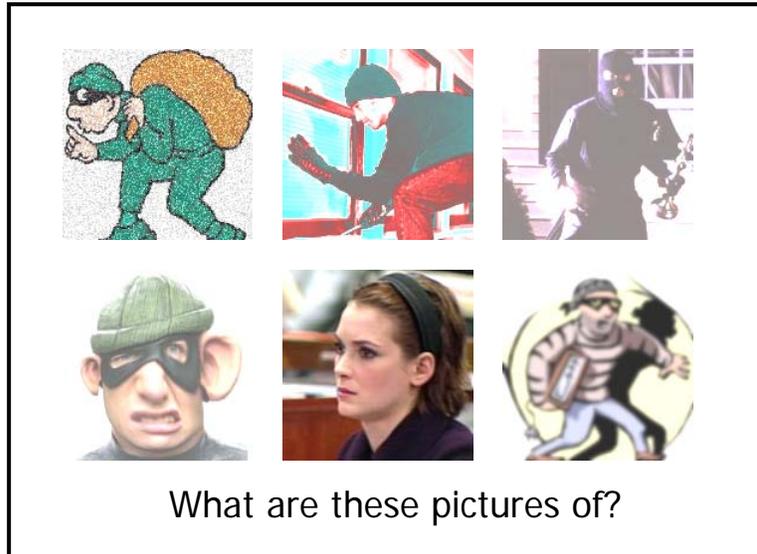


The images need to be randomly distorted

# Pix

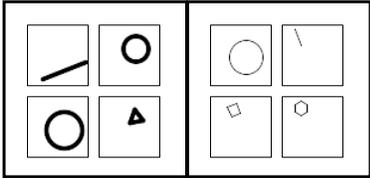


The images need to be randomly distorted



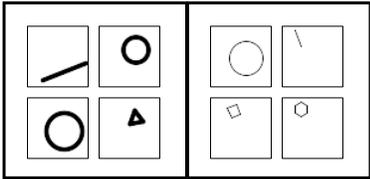
# Bongo

Visual Analogy  
Problems  
(Bongard 1951)



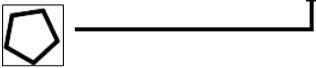
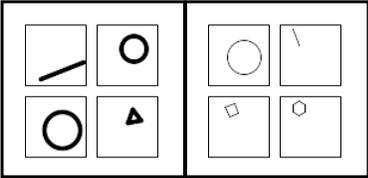
# Bongo

Visual Analogy  
Problems  
(Bongard 1951)



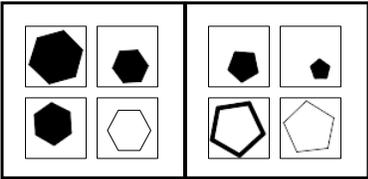
# Bongo

Visual Analogy  
Problems  
(Bongard 1951)



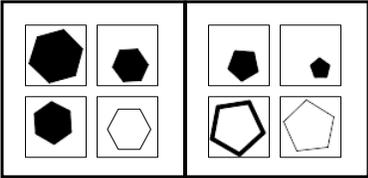
# Bongo

Visual Analogy  
Problems  
(Bongard 1951)



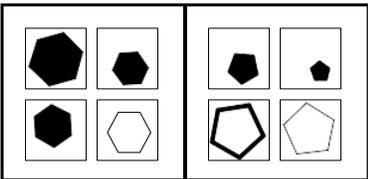
# Bongo

Visual Analogy  
Problems  
(Bongard 1951)



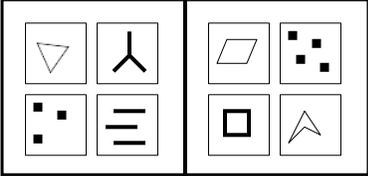
# Bongo

Visual Analogy  
Problems  
(Bongard 1951)



# Bongo

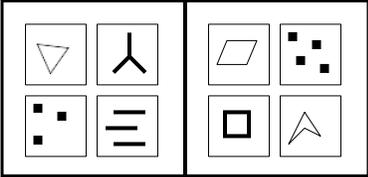
Visual Analogy  
Problems  
(Bongard 1951)



four

# Bongo

Visual Analogy  
Problems  
(Bongard 1951)



four



## Sound Oriented CAPTCHA

Humans are better than computers at understanding [spoken language](#)

Question: which English digits are being said?

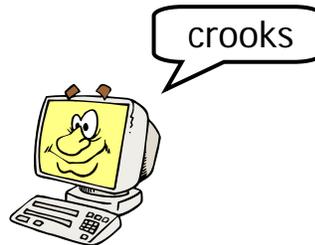
 (Written by Nancy Chan)

## Open Problem

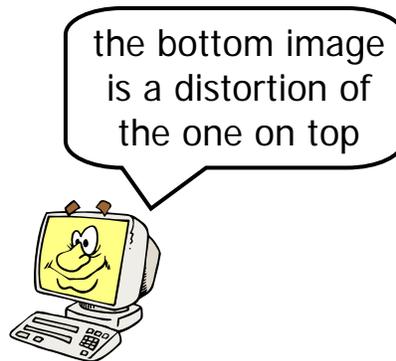
Create a CAPTCHA based on language [understanding](#)

Advancing AI

Any program that passes the tests  
generated by a CAPTCHA can be  
used to do [something good](#)



Any program that passes certain versions of PIX can be used to do **weak watermarking**



Any program that passes certain versions of PIX can be used to do **weak watermarking**

## CAPTCHAs Are a Win-Win Situation

Either a CAPTCHA remains secure **or**  
an open problem becomes solved

CAPTCHAs get **malicious people** to  
work on AI problems!

## Challenges to the AI Community

CAPTCHAs provide [well defined](#) problems for the AI community to work on

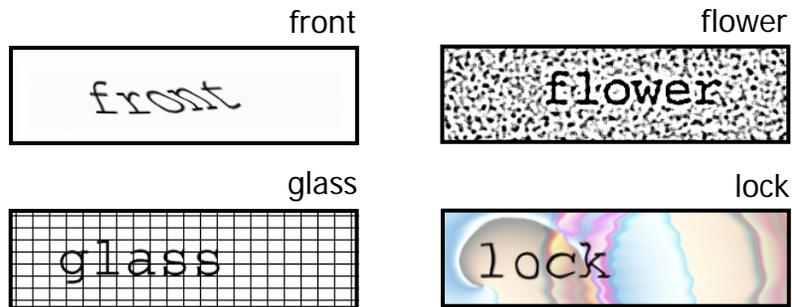
## Challenges to the AI Community

CAPTCHAs provide [well defined](#) problems for the AI community to work on

Algorithms for factoring have vastly improved since factoring started being used for [security](#)

## Advancing AI

Mori and Malik, 2002: 92% accuracy against  
Yahoo! CAPTCHA



## CAPTCHA Sweat Shops

Spam companies hire humans to  
solve CAPTCHAs [all day long](#)

\$5 per hour for each human

720 CAPTCHAs solved per hour per human

---

2/3 cent per account

ter avec Firefox À la une

reCAPTCHA™ **Digitizing Books One Word at a Time**

- HOME
- WHAT IS reCAPTCHA
  - DIGITIZATION ACCURACY
  - WHAT IS A CAPTCHA
  - SECURITY
- GET reCAPTCHA
- MY ACCOUNT
- EMAIL PROTECTION
- RESOURCES

**Submit** The words above come from scanned books.  
By typing them, you help to digitize old texts.

reCAPTCHA is a free CAPTCHA service that helps to digitize books, newspapers and old time radio shows. Check out [our paper](#) in Science about it (or read more below).

A **CAPTCHA** is a program that can tell whether its user is a human or a computer. You've probably seen them — colorful images with distorted text at the bottom of Web registration forms. CAPTCHAs are used by many websites to prevent abuse from "bots," or automated programs usually written to generate spam. No computer program can read distorted text as well as humans can, so bots cannot navigate sites protected by CAPTCHAs.

About 200 million CAPTCHAs are solved by humans around the world every day. In each case, roughly ten seconds of human time are being spent. Individually, that's not a lot of time, but in aggregate these little puzzles consume more than 150,000 hours of work each day. What if we could make positive use of this human effort? reCAPTCHA does exactly that by channeling the effort spent solving CAPTCHAs online into "reading" books.

To archive human knowledge and to make information more accessible to the world, multiple projects are currently digitizing physical books that were written before the computer age. The book pages are being photographically scanned, and then transformed into text using "Optical Character Recognition" (OCR). The transformation into text is useful because scanning a book produces images, which are difficult to store on small devices, expensive to download, and cannot be

Scanned type → **This aged portion of society were distinguished from**

OCR reads as → "niis aged pntkm at society were distinguished frow."

reCAPTCHA improves the process of digitizing books by sending words that cannot be read by computers to the Web in the form of CAPTCHAs for humans to decipher. More specifically, each word that cannot be read correctly by OCR is placed on an image and used as a CAPTCHA. This is possible because most OCR programs alert you when a word cannot be read correctly.

But if a computer can't read such a CAPTCHA, how does the system know the correct answer to the puzzle? Here's how: Each new word that cannot be read correctly by OCR is given to a user in conjunction with another word for which the answer is already known. The user is then asked to read both words. If they solve the one for which the answer is known, the system assumes their answer is correct for the new one. The system then gives the new image to a number of other people to determine, with higher confidence, whether the original answer was correct.

Currently, we are helping to digitize old editions of the [New York Times](#) and books from [Google Books](#).

**How can I help?**

In order to achieve our goal of digitizing books, we need your help.

**If you run a website** that suffers from problems with spam, you can [put reCAPTCHA on your site](#). For some applications (such as [WordPress](#) and [MediaWiki](#)), we have [plugins](#) that allow you to use reCAPTCHA without writing any code. We also have easy-to-use code for common web programming languages such as [PHP](#).

**If you get email spam** we have a method that will help you to reduce it. Many spammers crawl the web looking for email addresses. When they see an email address on a web page, they send spam to the address. [Mailhide](#) allows you to safely post your email address on the web. Mailhide takes an address such as `jsmith@example.com` and turns it into `jsm_@example.com`. In order to reveal the address, a user must click on the "..." and solve a reCAPTCHA. If you use the Mailhide version of your email address, spammers won't be able to find your real email address and you'll get less spam.

[Blog](#) | [About Us](#) | [Contact](#) | [FAQ](#) | [Terms](#) | [Privacy](#)

A Spin-off Idea

Labeling Images With Words

## Labeling Images With Words



## Labeling Images With Words



→ Martha Stewart  
Flowers  
Super Evil

## Labeling Images With Words



→ Martha Stewart  
Flowers  
Super Evil

Completely Open Problem

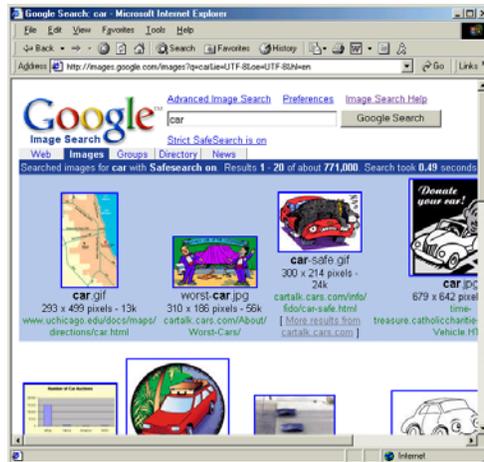
## Image Search on The Web



## Image Search on The Web

Uses **filenames** and  
surrounding **text**

Doesn't look at the  
actual image



## Desiderata

A method for labeling images that:

1. Actually looks at the images
2. For any image gives **several keywords** that make sense
3. Is very fast (Google has 425,000,000 images)

## Stealing Cycles From Humans

## Stealing Cycles From Humans

Over **50 million** people in the United States play computer games on a regular basis!

The ESP Game will allow us to label  
all images on the web in 30 days!

## The ESP Game

Two-player online game

Partners don't know each other  
and can't communicate

## The ESP Game

Two-player online game

Partners don't know each other  
and can't communicate

Object of the game: **type the  
same word**

The only thing in common  
is an **image**

## The ESP Game

Player 1



Player 2



## The ESP Game

Player 1



Guessing: car

Player 2



Guessing: boy

## The ESP Game

Player 1



Guessing: car

Guessing: hat

Player 2



Guessing: boy

## The ESP Game

Player 1



Guessing: car

Guessing: hat

Guessing: kid

Player 2



Guessing: boy

## The ESP Game

Player 1



Guessing: car

Guessing: hat

Guessing: kid

Success!

You both agree on car

Player 2

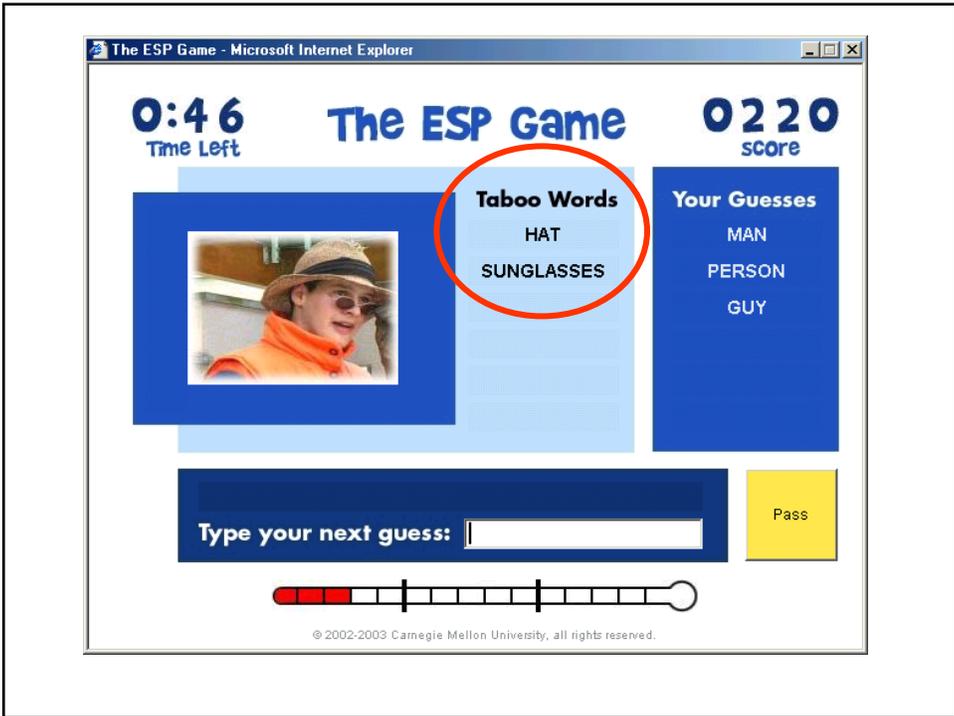
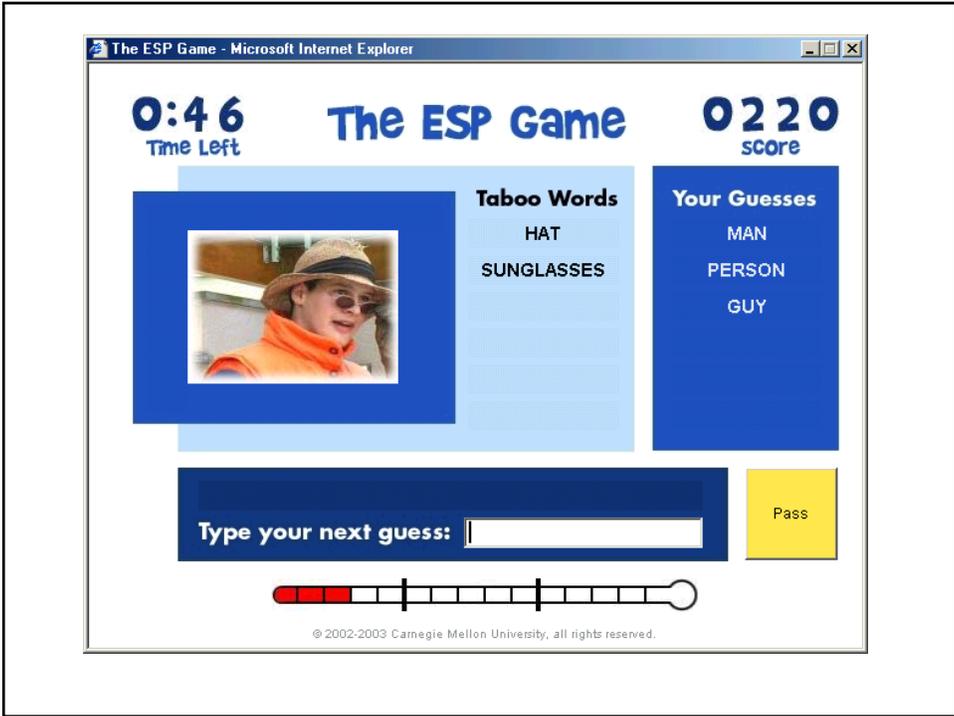


Guessing: boy

Guessing: car

Success!

You both agree on car



## The ESP Game

Taboos guarantee that each image will get [many different keywords](#)

## The ESP Game

Taboos guarantee that each image will get [many different keywords](#)

Preliminary studies suggest that [people find the game fun](#)

## The ESP Game

Average labeling rate: 4 images per minute

5000 people simultaneously playing the game  
would label all the images on Google in 30 days!

$$\frac{5000}{2} \times 4 \times 60 \times 24 \times 30 = 432,000,000$$

Individual games in Yahoo!, Pogo.com or MSN  
average well over 10,000 players at a time

## Take-Home Message #2

There are lots of people doing  
useless stuff on the internet

Stealing Cycles From Humans is a  
More General Idea...

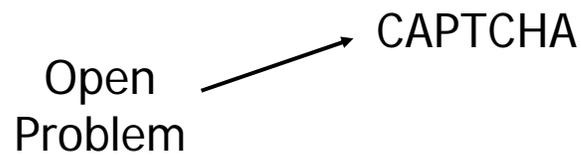
This talk hints at a [paradigm](#) for  
dealing with unsolved AI problems:

This talk hints at a [paradigm](#) for dealing with unsolved AI problems: getting others to do the work for you

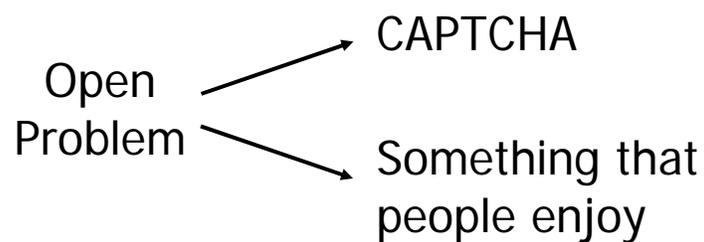
This talk hints at a [paradigm](#) for dealing with unsolved AI problems: getting others to do the work for you

Open  
Problem

This talk hints at a [paradigm](#) for dealing with unsolved AI problems: getting others to do the work for you



This talk hints at a [paradigm](#) for dealing with unsolved AI problems: getting others to do the work for you



[www.captcha.net](http://www.captcha.net)

[www.espgame.org](http://www.espgame.org)

