

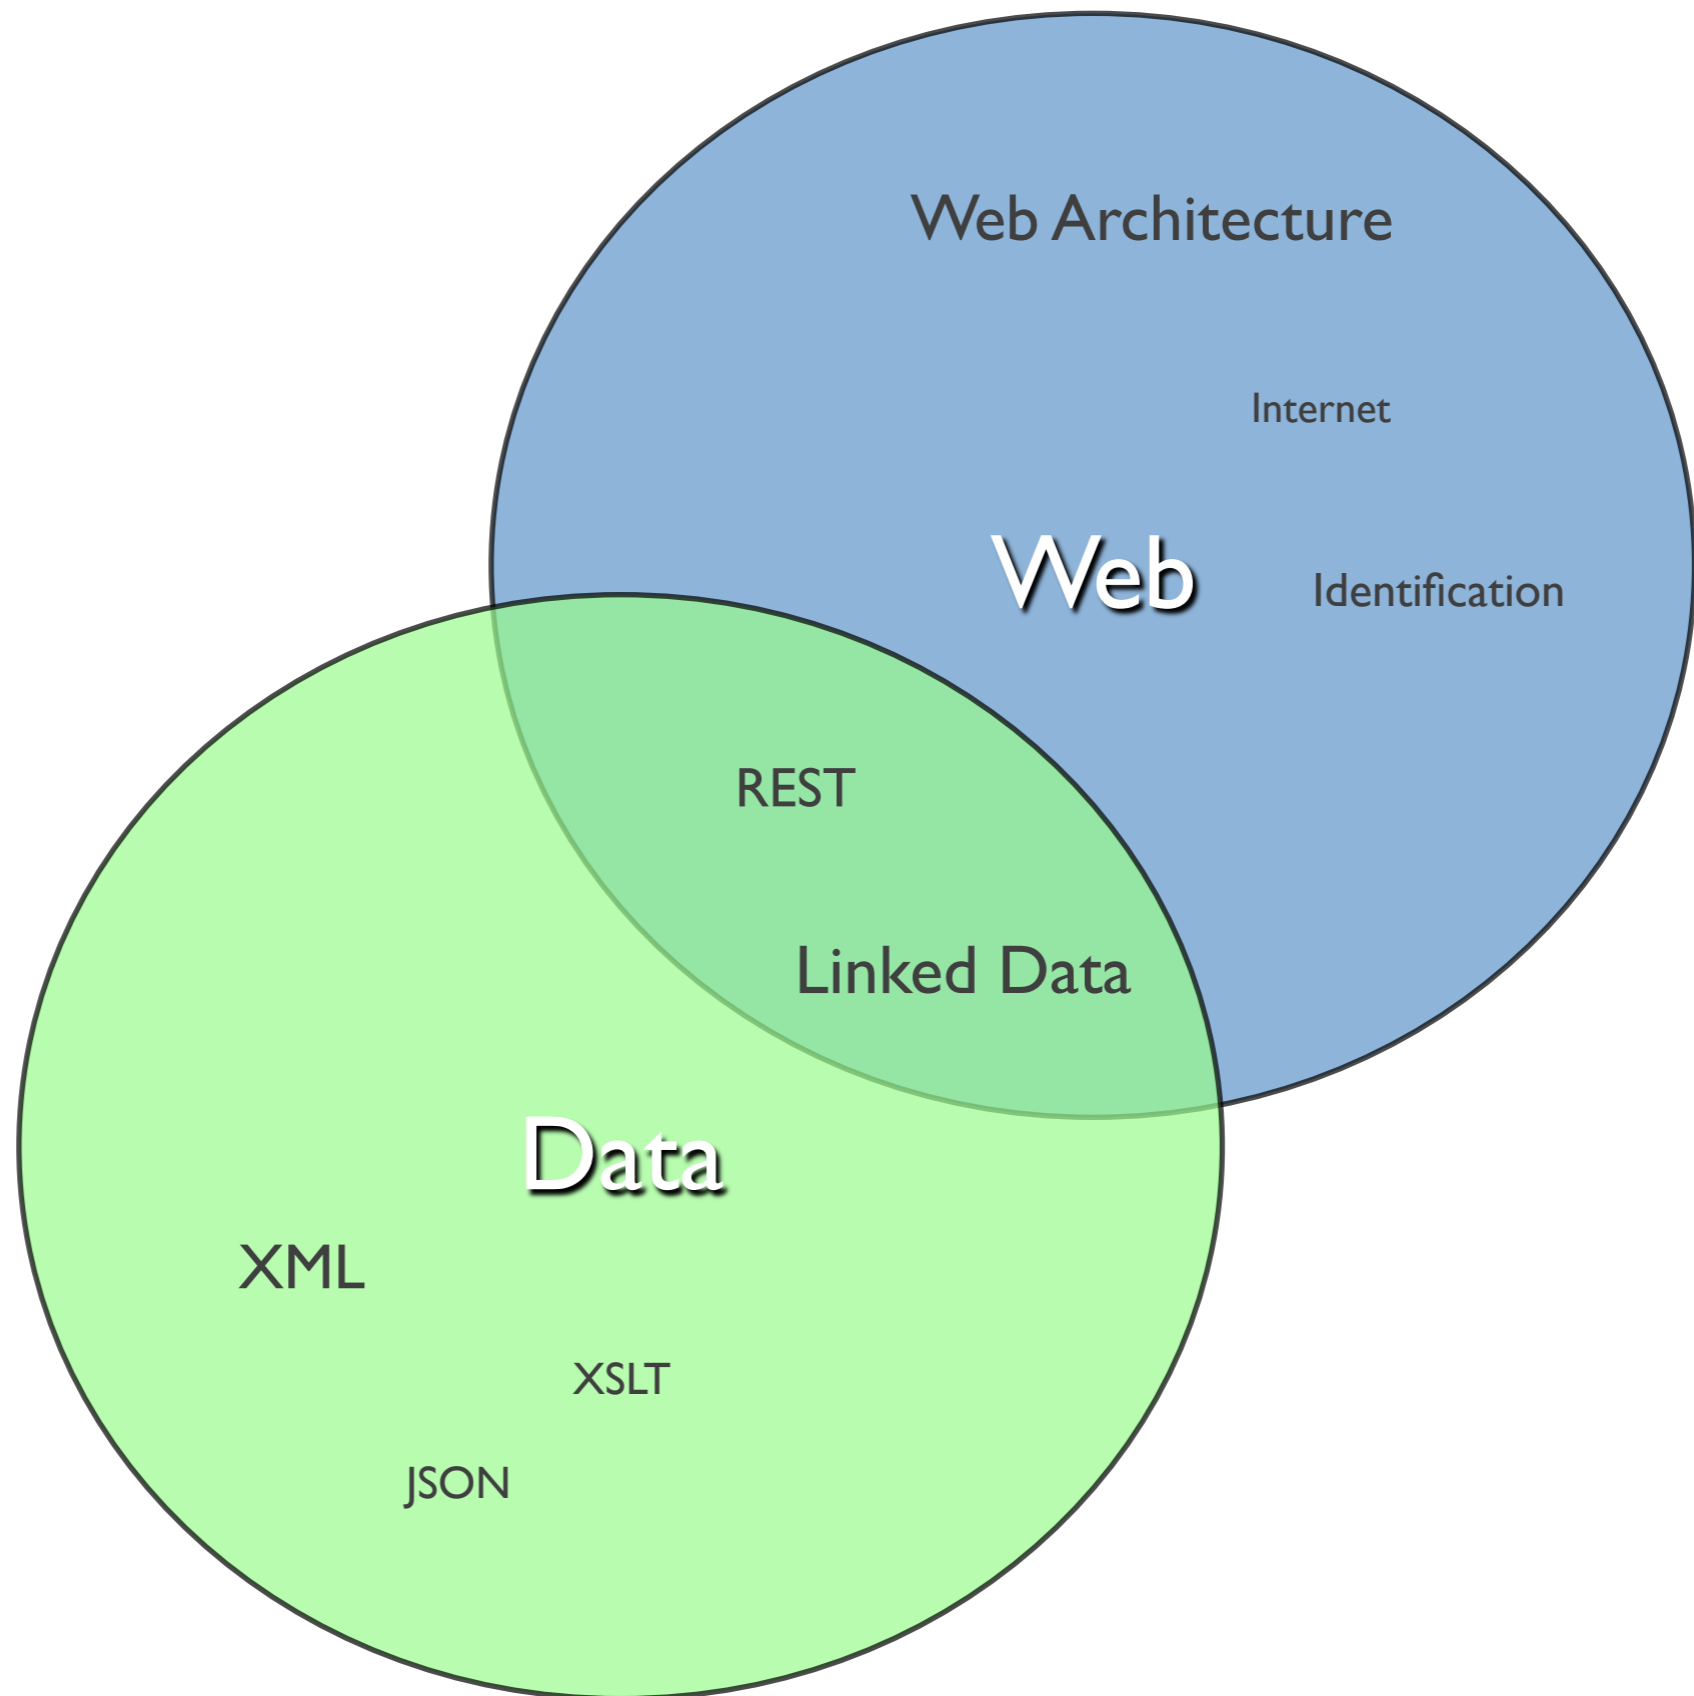
INFO/CS 4302
Web Information Systems

FT 2012

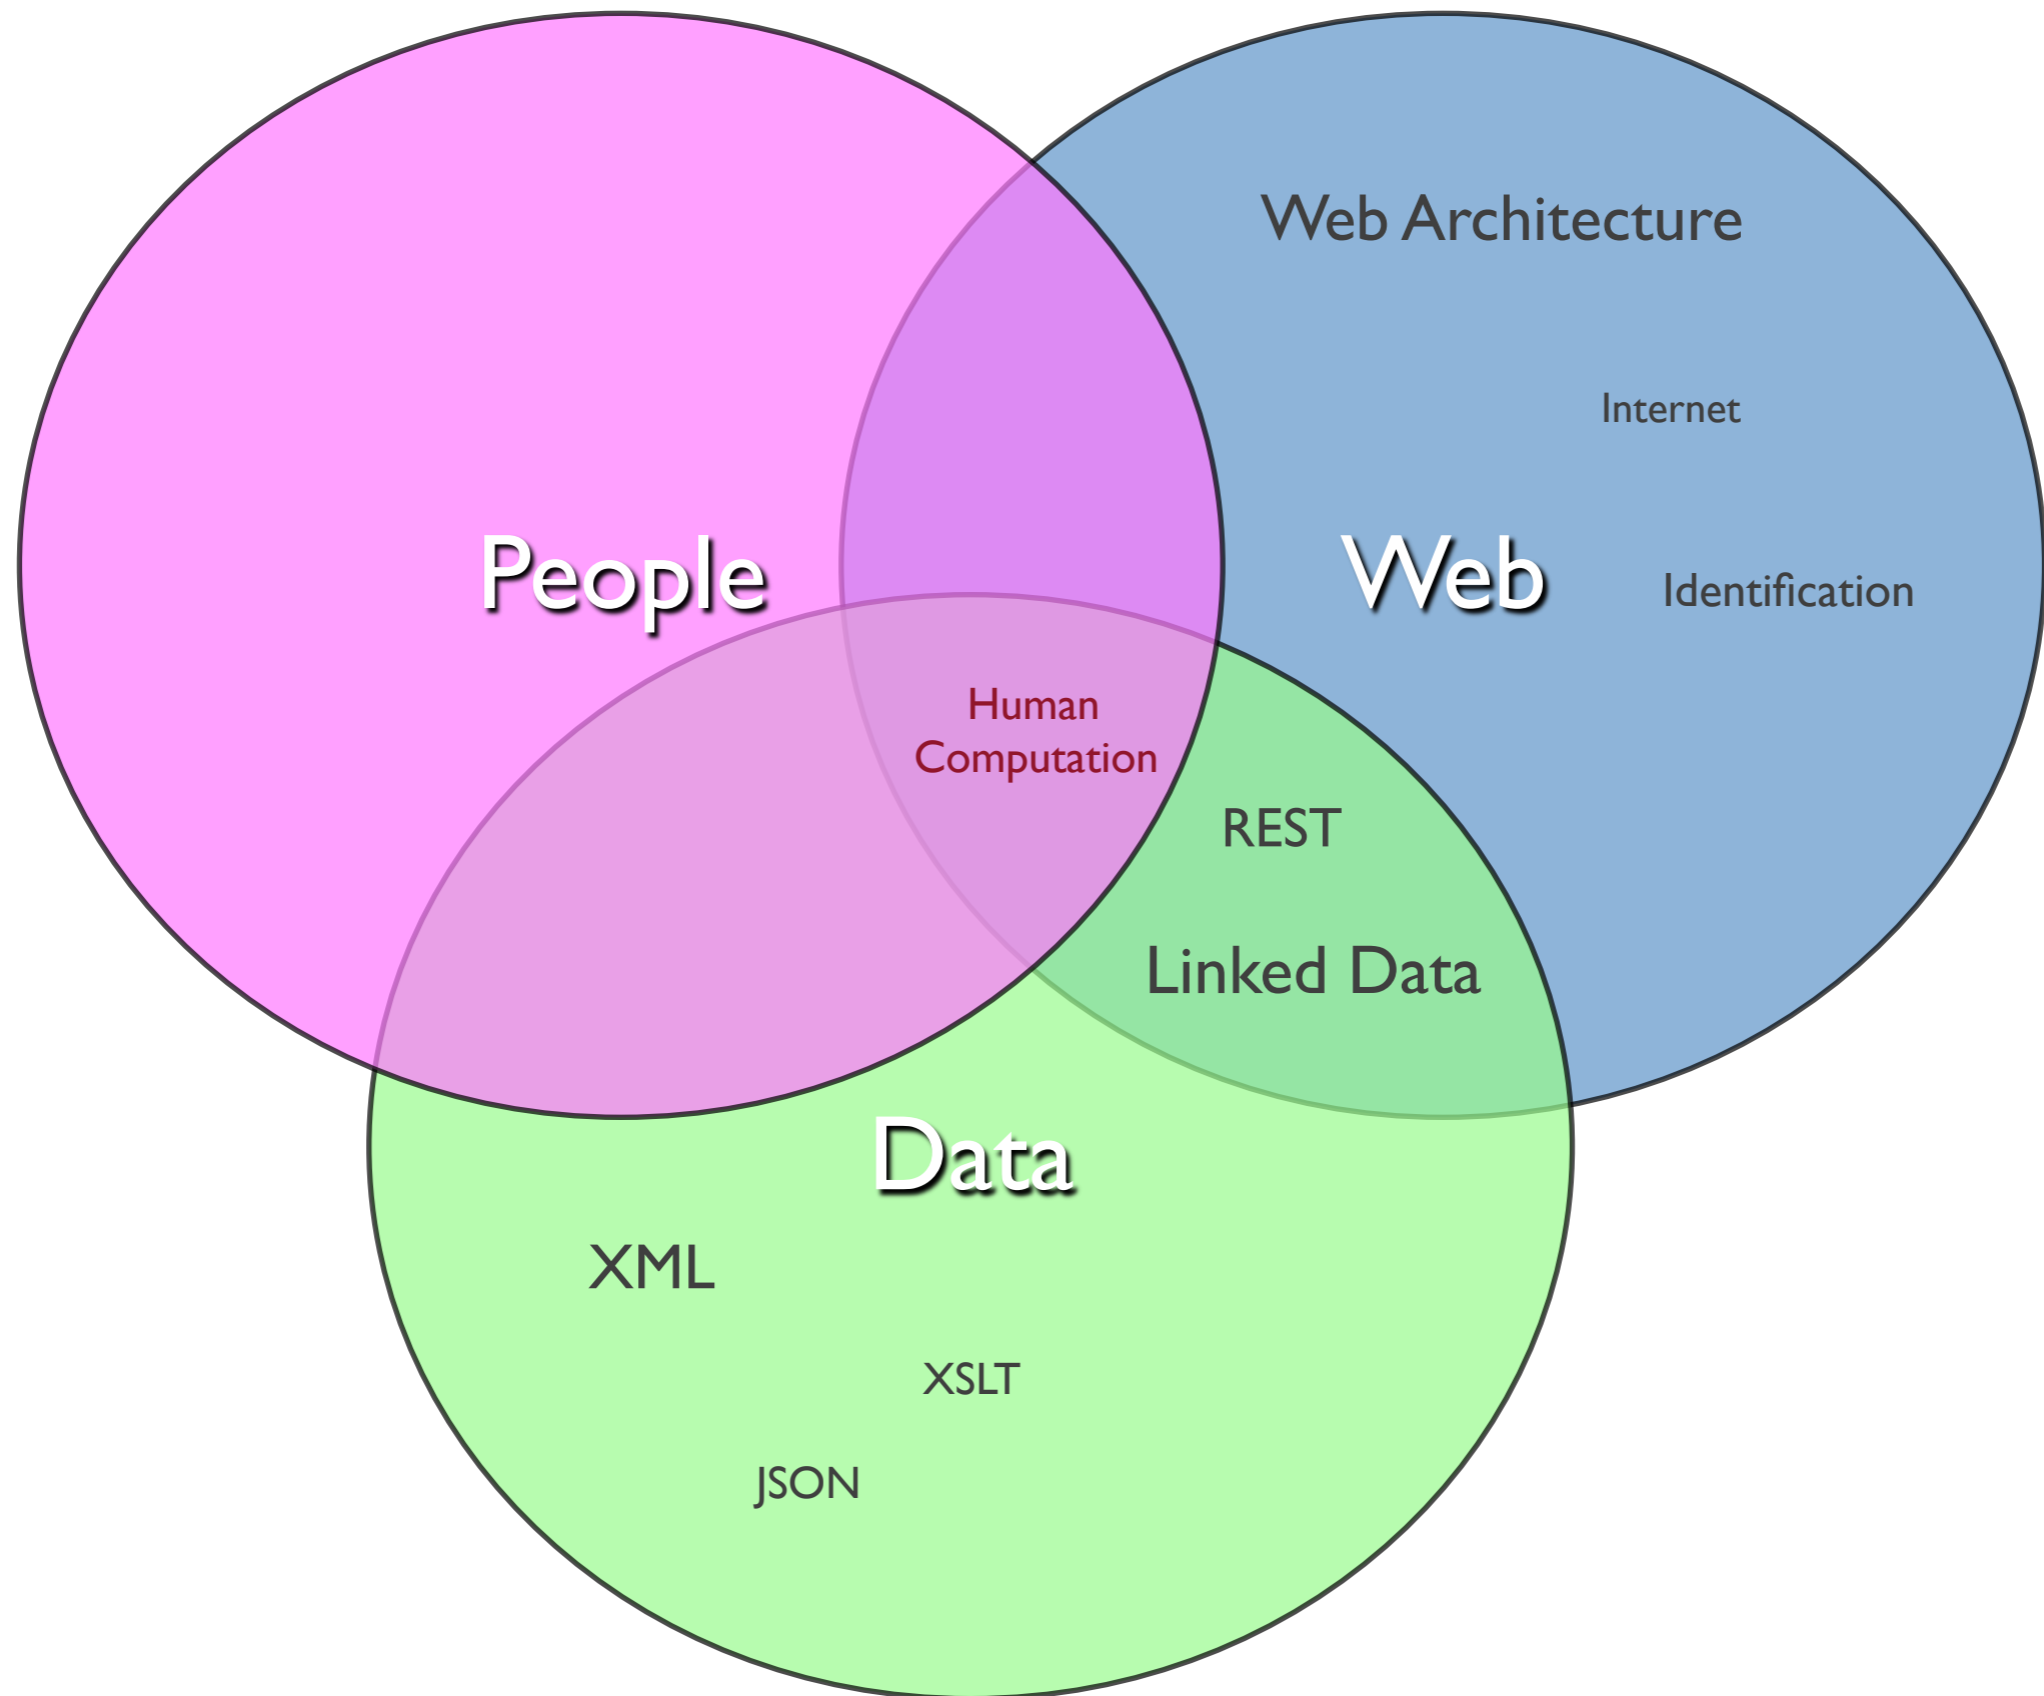
Week 13: Human Computation

- Bernhard Haslhofer -

This course so far..



Today we add one more dimension....



We talk about...

- **Problems** that are hard for AI algorithms
- How **Human Computation** can help
- Human Computation **examples**
 - reCAPTCHA
 - ESP Game
 - Crowdsourcing marketplaces (Amazon Mechanical Turk)
 - Citizen Science projects (Galaxy Zoo, eBird)

...and play some **games**.

Human Computation: Core R x

www.humancomputation.com/Tutorial.html

The 3rd Human Computation Workshop (HCOMP 2011)

Welcome About the Workshop Program Tutorial Previous Workshops

Human Computation: Core Research Questions and State of the Art

There will also be a 4-hour tutorial, given by Luis von Ahn and Edith Law, at AAAI on **August 7 (2:00 PM - 6:00 PM)** which will give newcomers and current researchers a bird's eye view of the research landscape of human computation. The tutorial will be based on materials from a new book called "**Human Computation**" (published by **Morgan & Claypool Synthesis Lectures on Artificial Intelligence and Machine Learning**), which will be distributed to all tutorial attendees. Click [here](#) to register for the tutorial.

Download the slides for [Part 1](#) and [Part 2](#) of the tutorial. In addition, if your institution has subscription to the series, you can download the book [here](#). Here is a rough description of what will be covered in the tutorial.

- Introduction**
 - Computation: Now and Then
 - What is Human Computation
 - Tackling AI Problems: From Vision to Biology
 - Tutorial Overview
- Human Computation Algorithms**
 - A Definition of Algorithms
 - Building Blocks of Algorithms (operations, controls, program synthesis)
 - Programming Frameworks
 - Evaluating Human Computation Algorithms (correctness, efficiency)
- Aggregating Outputs**
 - Objective versus Cultural Truth
 - Classification (latent class models, learning from imperfect data)
 - Beyond Classification (ranking/voting, clustering, structured outputs, beliefs)
- Task Routing**
 - Push versus Pull Approaches
 - Push Approach (allocation, matching, inference)
 - Pull Approach (search and visualization, task recommendation, peer routing)
- Understanding Workers and Requestors**
 - Human Computation Markets (paid crowdourcing, security and access, games)

MORGAN & CLAYPOOL PUBLISHERS

Human Computation

Edith Law
Luis von Ahn

SYNTHESIS LECTURES ON ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING


Ronald J. Brachman, William Cohen, and Thomas G. Dietterich, *Series Editors*

ARTIFICIAL INTELLIGENCE (AI) PROBLEMS

The Breckinridge and Lane Democrats, having taken courage at the recent eastern advices, are organizing energetically for the campaign. Several prominent Democrats who at first favored DOUGLAS, are coming out for the other side, apparently under the pressure of Federal influence. An address to the National Democracy of California, urging the party to support BRECKINRIDGE, has recently been published, which manifestly has strengthened that side of the question. It is signed by 65 Democrats, many of whom occupy respectable and prominent positions in the party, 22 of them are Federal office-holders, eight more are recipients of Federal patronage, and the others represent a mass of politicians giving the document most weight. The Douglas Democrats are also active. The Irish and German vote will mostly go with that branch of the party, but it is difficult to estimate which wing is the stronger. Thus far 17 Democratic newspapers have declared for DOUGLAS, 13 for BRECKINRIDGE, and 9 remain non-committal, with even chances of going either way. Under these circumstances the Republicans entertain not unjustifiable hopes that the Democratic divisions may be so equally balanced as to give the State to LINCOLN. Some very respectable Bell and Everett meetings have been held in different parts of the State, but thus far that party does not exhibit much rank and file strength.

The New-York State Yacht Squadron, on its annual cruise to Newport, came into the harbor yesterday afternoon. The following are the names of the boats that came to anchor here: *Jessie*, *Geraldine*, *Evelyn*, *Annie*, *Mannering*, *Julia*, *Bonita*, *Magie*, *Widgon*, *Rambler*, *Fleur-de-Lis*, *Henrietta*, *Sea-Drift* and *Maria*, with the steamer *America* as a tender. On anchoring, each boat fired a gun, according to custom. The reports were heard distinctly in the city, causing considerable inquiry as to "what was up," and quite a number of sanguine individuals came into our office to inquire if the guns were not annunciatory signals of the successful laying of the Atlantic Cable. We invariably replied in the negative. The squadron will leave to-day for Newport. The yachts *Washington* and *Rambler*, of this city, start with it, with parties of New-Haven people.





Artificial Intelligence (AI) Problems

- There are many problems that are
 - easy to solve for humans
 - but difficult for even the most sophisticated computer algorithms

Artificial Intelligence (AI) Problems

- For certain problems no AI algorithm can exceed human performance
 - **perceptual tasks** (object recognition, music classification)
 - **natural language analysis** (sentiment analysis, language translation)
 - **cognitive tasks** (planning and reasoning)

HUMAN COMPUTATION

What is Human Computation?

- Computation = the process of mapping **input** to **output** using an explicit, finite set of instructions (**algorithm**)
- Human Computation = computation carried out by a **human**

EXAMPLES - RECAPTCHA

CAPTCHA

following finding

<http://en.wikipedia.org/wiki/File:Modern-captcha.jpg>

- CAPTCHAs
 - prevent automated programs from abusing online services (buy one million tickets online)
 - prove that you are a human and not a computer
- > 200 M CAPTACHAs are solved per day
 - each taking approx. 10 seconds of human effort
 - 500.000 hours per day in total

reCAPTCHA

- reCAPTCHA channels effort into a “useful” purpose: **transcribing scanned books**
- Books written before the computer age are currently being digitized using OCR
 - Google Books
 - Internet Archive
- In older prints approx. 30% of words cannot be recognized by OCR

reCAPTCHA

The Norwich line steamboat train, from New-London for Boston, this **morning** ran off the track seven miles north of New-London.




- Old printed material is being transcribed word by word by people typing CAPTCHAs on the Web
- Displays words from scanned texts that OCR could not decipher

reCAPTCHA

The Norwich line steamboat train, from New-London for Boston, this morning ran off the track seven miles north of New-London.

morning



- User input needs to be **verified**
- reCAPTCHA gives the user two words
 - one for which the answer is not known
 - a **control word** for which the answer is known
- If the user correctly types the control word the system assumes that
 - they are human
 - gains confidence that the unknown word is typed correctly

reCAPTCHA

- 750 million people (10 % of humanity) have helped digitized at least one word
- Current transcription rate / day: > 160 books
- Experiment with 50 scanned NY Times Articles
 - Standard OCR accuracy: 83.5 %
 - reCAPTCHA accuracy: 99.1 % (216 errors)
 - professional manual transcription: 189 errors

OCR Transcription

The Breckinridge and Lane Democrats, having taken courage at the recent eastern advices, are [xxxxxxxxxx] energetically for the campaign: Several prominent Democrats who at first favored DonoLea, are coming out for the other aide, apparently under the [xxxxxxxx] of Federal [xxxxxxxx]. An address to the National Democracy of California, urging the party to support HaeslipsDas, has recently been published, which manifestly bss strengthened that aide of the [xxxxxxxx]: It is signed by 65 Democrats, many of whom occupy respectable and prominent positions in the party, 22 of them are Federal office-holders, [xxxxx] more are recipients of Federal patronage, and the others represent a mass of politicians giving the document [xxxx] [xxxxxx] mTheDouglas Democrats are also active The Irish and German vote will mostly go with this# branch of the party, but it is [xxxxxxxx] to [xxxxxxxx] [xxxxxx] [xxxx] [xx] the stronger. Thus far 17 IT newspapers have declared for DonGres, 13 for Base\$- laaIDGS and 9 remain non-committal, with even chances of going either way. Under these circumstances the Republicans entertain not unjustifiable hopes that the Democratic divisions may be so equal, - ly balanced as to give the State [xx] LiaCOLV. Same very [xxxxxxx] Bell and Everett meetings have been held in different parts of the State, bat thus far that party does not exhibit much rank sad ale air en.

reCAPTCHA Transcription

The Breckinridge and Lane Democrats, having taken courage at the recent eastern advices, are organizing energetically for the campaign. Several prominent Democrats who at first favored Douglas, are coming out for the other side, apparently under the pressure of Federal influence. An address to the National Democracy of California, urging the party to support Breckinridge has recently been published, which manifestly has strengthened that side of the question. It is signed by 65 Democrats, many of whom occupy respectable and prominent positions in the party, 22 of them are Federal office-holders, eight more are recipients of Federal patronage, and the others represent a mass of politicians giving the document most weight. The Douglas Democrats are also active The Irish and German vote will mostly go with that branch of the party, but it is difficult to estimate which wing is the stronger. Thus far 17 Democratic newspapers have declared for Douglas, 13 for Breckinridge and 9 remain non-committal, with even chances of going either way. Under these circumstances the Republicans entertain not unjustifiable hopes that the Democratic divisions may be so equally balanced as to give the State to Lincoln. Some very respectable Bell and Everett meetings have been held in different parts of the State, but thus far that party does not exhibit much rank and file strength.

EXAMPLES - ESP GAME

ESP Game - Motivation

- Images on the Web present a major technical challenge
 - millions of them are available
 - textual descriptions are rare
 - computer vision techniques are not accurate enough
- The only method for obtaining precise image descriptions is **manual labeling**
- 5000 people can assign labels to all images indexed by Google in 31 days (at least back in 2003)

ESP Game - Motivation

- It is estimated that over 200 million users play online games every week



ESP Game - Motivation

- By the age of 21 the average American spent more than **10,000 hours** playing such games
- Equivalent to **five years** working on a **full-time job** 40 hour per week
- **Games With A Purpose** (GWAP) channel this effort toward solving computational problems and training AI algorithms
- The **ESP Game** is one GWAP instance, applied to the problem of labeling Web images

ESP Game - Basic Idea

- Engage **pairs of players** in a simple game
- Let them **tag images** independently
- **Reward** them when their tags agree

11.1.4.1 System Description

- Partners are randomly assigned from among all people playing the game
- The goal of the ESP game is to guess what their partner is typing for each image
- Once both players have typed the same string ("agreed on an image"), they move on to the next image



Fig. 11-1: Partners agreeing on an image. Neither of them can see the other's guesses.

- Partners get (bonus) points if they agree on an image
- Agreed-upon strings are typically good labels for the image

ESP Game - Output Accuracy

- For one input (image) there can be multiple outputs (tags)
- **Taboo words**
 - are words that the players are not allowed to enter
 - are commonly guessed words related to the image
 - are determined dynamically
 - force players to enter **more specific tags**
 - guarantee that each image will get many **different labels**

Time to play some games....

<http://www.gwap.com/>

GWAP Design Considerations

- GWAPs implement **casual games**
 - low entry barrier (e.g., easy online access)
 - few simple controls, extremely **easy to learn**
 - non-pushing (multiple scoring opportunities)
 - can be consumed within short periods of time; 5-20 min during work (breaks)
 - inclusive, gender-neutral, little violent content
- Human Computation games must also be **fun**

Existing Human Computation Games

Table 6.1: Survey of human computation games.

Game	Description	Mechanism	AI Problem
The ESP Game [334]	two players match on a tag for the same image	output-agreement	object recognition
Peekaboorn [337]	player 1 reveals parts of the image associated with a secret word, player 2 must guess the secret word	function computation (problem inversion)	object location
Verbosity [338]	player 1 describes the properties of the entity associated with a secret word, player 2 must guess the secret word	function computation (problem inversion)	knowledge extraction
TagATune [190]	two players exchange tags and determine if two music clips are the same or different	function computation (input agreement)	music classification
FoldIt [69]	players fold protein structures to minimize total energy	function computation (optimization)	protein folding
HerdIt [24]	players select tags that describe the music	output-agreement	music classification
Categorilla [329]	players name an entity that fits a template (e.g., Things that fly)	output-agreement	natural language processing
MoodSwings [174]	players click on a 2-dimensional grid to indicate the valence and intensity of the mood of a music clip	output-agreement	music mood classification
Phrase Detective [57]	players identify relationships between words and phrases in a short piece of text	output-agreement	natural language processing
Phylo [255]	players align colored blocks by moving them horizontally and inserting gaps	output-agreement	genome alignment

CROWDSOURCING MARKET PLACES

Mechanical Turk is a marketplace for work.

We give businesses and developers access to an on-demand, scalable workforce.

Workers select from thousands of tasks and work whenever it's convenient.

185,774 HITs available. [View them now.](#)

Make Money by working on HITs

HITs - *Human Intelligence Tasks* - are individual tasks that you work on. [Find HITs now.](#)

As a Mechanical Turk Worker you:

- Can work from home
- Choose your own work hours
- Get paid for doing good work



[Find HITs Now](#)


or [learn more about being a Worker](#)

Get Results from Mechanical Turk Workers

Ask workers to complete HITs - *Human Intelligence Tasks* - and get results using Mechanical Turk. [Register Now](#)


As a Mechanical Turk Requester you:

- Have access to a global, on-demand, 24 x 7 workforce
- Get thousands of HITs completed in minutes
- Pay only when you're satisfied with the results



[Get Started](#)

Mechanical Turk



Amazon Mechanical Turk

- A platform for **requesters** to post **tasks** to **workers** to perform in return for monetary payment
- Tasks are called **HITS** “Human Intelligence Tasks.”
- Approx. 400,000 workers by 2010

AMT Tasks

- Tasks are typically small
- 90% of the HITS have rewards < 10 cent
- Typical tasks
 - classification (images, music, documents, ...)
 - transcription
 - creation of original content (reviews, stories, blog posts)
- Can be created programmatically (**API**)

AMT Workers Demographics

- Survey 2010 / 1000 Workers
- From 60 countries, with majority (~80%) from the US and India
- Workers' characteristics
 - 2:1 female/male ratio in the US (reverse in India)
 - on average younger and lower income than general population
 - secondary (US) versus primary (India) source of income
 - tend to have higher education

EXAMPLES - CITIZEN SCIENCE

What is Citizen Science?

- Science is **data-intensive**
 - climate patterns
 - species distributions
 - trajectories of stars
- Tedious, time-consuming, and sometimes impossible for a few scientists
- Idea: engage **non-scientists** in the **collection** and **interpretation** of data

Galaxy Zoo

- Millions of **galaxies** can be seen in images taken by the Hubble space telescope and other telescopes on Earth
- Galaxies can easily be **classified** by their **shape**
 - hard task for automated algorithms
 - easy task for humans

CLASSIFY

SCIENCE

STORY



ASTRONOMERS

DISCUSS

PROFILE



Classify



SDSS



Invert

Help

Restart

SHAPE

Is the galaxy simply smooth and rounded, with no sign of a disk?



Smooth



Features or disk



Star or artifact

Galaxy Zoo - Facts

- Largest astronomical collaboration in history
 - 200,000 participants from 113 countries
 - > 100 M classifications of galaxies
 - Resulted
 - in new discoveries
 - in an expanded project
- <https://www.zooniverse.org/>

Time to classify some galaxies....

<http://www.galaxyzoo.org/>

Galaxy Zoo - Data Collection

- Some users made $> 100,000$ classifications, most only around 30
- Bogus data
 - occur if users have browser problems or someone actively cheats
 - 0.05% of all users were found unreliable in a first experiment
- Classifications are weighted

eBird



- [Home](#)
- [About eBird](#)
- [Submit Observations](#)
- [View and Explore Data](#)
- [My eBird](#)

[Sign In](#) | [Register as a New User](#)

Translate to: [English](#) | [Español](#) | [Français](#)

Welcome to eBird

eBird is as birder does!

Support eBird



eBird News and Features

[eBirder of the Month -- Mary Gustafson](#)

November 12, 2011

The [Rio Grande Valley Birding Festival](#) is taking place this weekend, which means that our current November eBirder of the Month will be buzzing around Harlingen corralling teams of leaders, making sure field trips get off on time, and generally making sure that the 2011 festival comes off as another huge success. Sadly, all this means that Mary won't be getting out in the field as much as she would like. So it seems like an appropriate time to recognize her not only for her hard work as field trip chair for this festival, but also for her unflagging commitment to eBird, both as a user and a reviewer. Moreover, Mary has solidified the Rio Grande Valley Birding festival's commitment to eBird by ensuring that all field trip data are entered into the system, and that all festival participants become aware of the project. This adds up to hundreds of new eBirders each year thanks to this festival! Thanks Mary, and be sure to get out (e)Birding again when things settle down on Monday!




[Possible hybrid Harlan's Red-tailed Hawk x Rough-legged Hawk in Kansas](#)

November 07, 2011

In North America, Red-tailed and Rough-legged Hawks share highly variable plumages, often causing



eBird Rarity Photos Pool



eBird

- Engages a global network of **bird watchers** to report their **bird observations** to a centralized database
- Anyone, anywhere, anytime
- 50,000 individuals from > 200 countries
 - volunteered > 4 million hours
 - collected > 70 million bird observations

References / Readings

- Law, E. and Luis von Ahn: [Human Computation](#). Synthesis Lectures on Artificial Intelligence and Machine Learning, 2011. 5(3): p. 1-121
- Luis von Ahn et al.:
[reCAPTCHA: Human-Based Character Recognition via Web Security Measures](#). Science, 2008, 321 (5895)
- Luis von Ahn and Laura Dabbish: [Designing games with a purpose](#). Communications of the ACM, 2008. 51(8).
- Lintott, C.J., et al., Galaxy Zoo: morphologies derived from visual inspection of galaxies from the Sloan Digital Sky Survey. Monthly Notices of the Royal Astronomical Society, 2008. 389(3): p. 1179-1189.
- Sullivan, B.L., et al., eBird: A citizen-based bird observation network in the biological sciences. Biological Conservation, 2009. 142(10): p. 2282-2292.