

Using games to solve challenging multimedia problems



Oge Marques, PhD

Professor

College of Engineering and Computer Science Florida Atlantic University – Boca Raton, FL (USA)







Association for Computing Machinery

Advancing Computing as a Science & Profession

For additional information, please visit http://dsp.acm.org/



About ACM



ACM, the Association for Computing Machinery is the world's largest educational and scientific computing society, uniting educators, researchers and professionals to inspire dialogue, share resources and address the field's challenges.

ACM strengthens the computing profession's collective voice through strong leadership, promotion of the highest standards, and recognition of technical excellence.

ACM supports the professional growth of its members by providing opportunities for life-long learning, career development, and professional networking.

With over 100,000 members from over 100 countries, ACM works to advance computing as a science and a profession. <u>www.acm.org</u>

Take-home message

Solutions to several challenging problems in multimedia can be improved using **human computation**, particularly through properly designed games.



https://youtu.be/vUH-eZTSTfs

Questions

• What is the problem?

• What is the proposed solution?

• Does it work (as advertised)?

How does it compare to other possible solutions?

The rest of the talk

- Examples of games (and the problems they help solve)
- Our work
- Challenges
- Insights from the psychology of games and gamers
- Advice for (young) researchers
- Final reflection
- Q&A

What types of problems?

- *Research questions* that can be mapped to *tasks* that:
 - Are easy for humans and hard for computers
 - Require intensive labor
 - Enable noble scientific pursuits
 - Improve human life



Seung Lab

Computational Neuroscience @ Princeton



By joining Eyewire, you can help map the connectome, starting with connections between retinal neurons. Eyewire gameplay advances neuroscience by helping researchers discover how neurons connect and network to process information. You also help develop advanced artificial intelligence and computational technologies for mapping the connectome.







Foldit

Foldit is a revolutionary new game, in which you play to solve puzzles, and we test your solutions to work on curing cancer, AIDS, and a host of diseases.



Gamification poured into every lesson.



Read, Listen, Speak

Each lesson includes a variety of speaking, listening, translation, and multiple choice challenges.

In-Lesson Grading

Instantly see which answers you get correct. When you miss a challenge, we'll quickly show you how to improve.

	•		
		2 • •	
 •) • 🗖		
	_		



Duolingo motivates you to stay on track by recording how many days in a row you spend learning a language.

Hearts

Hearts keep your lessons alive! You lose them when you answer incorrectly. When you're out of hearts, start over and try again.

GWAP(.com)



GWAP(.com)

FRIDAY, AUGUST 19, 2011

GWAP: 2008-2011

GWAP (for Games With a Purpose) was an academic project at CMU that explored the idea of Human Computation games to solve problems that computers cannot solve.

Over 200,000 people played games here, but unfortunately the creators decided to move on to other projects.

POSTED BY LUIS VON AHN AT 6:02 PM NO COMMENTS:

Selected examples



von Ahn, L., Dabbish, L.: Esp: Labeling images with a computer game. In: AAAI Spring Symposium: Knowledge Collection from Volunteer Contributors. (2005) 91–98

Selected examples

Artigo

Web-based platform

(<u>http://www.artigo.org</u>) with 6 artwork annotation games + artwork search engine in English, French, and German.



- Project funded by the German Research
 Foundation (DFG)
- 2008-2013:
 - over 7 million tags (mostly in German)
 - 180,000 players (about a tenth of whom are registered)
 - 150 players per day, in average.

GAMES

- → ARTIGO GAME [®]
- → ARTIGO TABOO [®]
- → KARIDO [®]
- → ARTIGO-QUIZ BETA
- → TAG A TAG BETA
- → COMBINO BETA

Selected examples

- Smorball (<u>http://www.tiltfactor.org/game/smorball/</u>)
 - Dartmouth College's *Tiltfactor* interdisciplinary studio that "designs and studies games for social impact"
 - Addresses limitations in full-text searching of digitized material due to poor output from OCR software.
 - Gameplay: As the coach of the Eugene Melonballers, you must stop the opposing smorbots from getting your team's endzone.
 - As an opponent approaches, type the phrase in the corresponding row and press "enter" in order to command your athlete to tackle the opponent.







Problems

- Our design process is often reversed:
 - 1. start from a problem
 - 2. think of a crowdsourcing solution
 - 3. create a tool
 - 4. make it look like a game

Problems

- Our terminology is not exactly inspiring:
 - "Serious games"
 - "Games with a purpose" (GWAP)
 - "Human-based computation games"
 - "Non-entertainment focused games"



Our work: Ask'nSeek

Ask'nSeek

[Carlier et al. (ECCV 2012), Salvador et al. (CrowdMM 2013)]

- A two-player, web-based, game that asks users to guess the location of a hidden region within an image with the help of semantic and topological clues.
 - One player (*master*) hides a rectangular region somewhere within a randomly chosen image.
 - The second player (*seeker*) tries to guess the location of the hidden region through a series of successive guesses, expressed by clicking at some point in the image.







Master and *seeker*



At the end of a game

654 Ast 1	Ask nSeek	
You are playing as SEEKER	Indications	s:
<image/>	On the left of the red car Above the white car Above the road	

Ask'nSeek: collected data



Ask'nSeek: collected data

Spatial information

the cat

Points "on the right" of the cat



Ask'nSeek: collected data

Spatial information

"On" points



"Partially on" points



Ask'nSeek: model



Ask'nSeek: model

$$\begin{aligned} \mathcal{L} &= [\mathcal{L}_1, \dots \mathcal{L}_K] \cup \phi = \{\mathcal{L}_l\}_{l=1\dots K+1} \\ \mathcal{X} &= X_{cv} \cup X_{on} \cup X_{pon} \\ \mathcal{R} &= \{R_{i,l,r}\}_{x_i \in \mathcal{X}, r=1\dots 4, l=1\dots K} \end{aligned}$$
 Points Relations

Content Analysis

Users

• Examples of results







 Live prototype at: <u>http://tinyurl.com/asknseek</u>



- Original Goal
 - Object detection and labeling

Ask' nSeek





Modified Goal
 Object segmentation







Our work: Click'n'Cut

Click'n'Cut

[Carlier et al. (CrowdMM 2014), (MTAP 2015)]

- Interactive Segmentation Tool where users are asked to produce *foreground* and *background* clicks to perform a segmentation of the object that is indicated in a provided description.
 - Every time a user produces a click, the segmentation result is updated and displayed over the image with an alpha value of 0.5.
 - This segmentation is computed using an algorithm based on object candidates (Arbelaez et al., 2014) and aims at guiding the user to provide information (i.e., meaningful clicks) that will help improve the quality of the final segmentation result.





UNIVERSITAT POLITÈCNICA DE CATALUNYA



Click'n'Cut (1/105)

Extract the fish.



Left click on the Foreground Right click on the background To reset your clicks, please click "Clear Points" Click on any point to remove it Use the slider to modify the mask transparency Once you are satisfied with the mask, click 'Done' to go to the next task



• Experiments

- Click'n'Cut Experts: 15 computer vision researchers from academia, both students and professors.
- Click'n'Cut Workers: 20 paid workers from the platform microworkers.com
 - Each worker was paid 4 USD for annotating 105 images.
- Ask'nSeek Players: 162 players (mostly students) played the Ask'nSeek game on any number of images they wanted to.
- 100 images, 105 tasks
- 5 Gold Standard tasks, to control for errors

- Segmentation results
 - Average Jaccard Index

	Click'n'Cut	Click'n'Cut	Ask'nSeek
	Experts	Paid workers	Players
All users	0.9	0.14	0.44
Users with less than 50% errors on Gold Standard	0.9	0.63	0.43
Users with less than 20% errors on Gold Standard	0.9	0.82	0.40

- Crowdsourcing loss
 - a loss induced by having a task performed by workers instead of experts



Gamification loss

- Q: Why are the segmentation results obtained with the Ask'nSeek game poorer than the ones obtained with the Click'n'Cut interactive segmentation tool?
- A: Fewer (and spatially limited) clicks, usually on salient areas.





Beyond Click'n'Cut

Pull the Plug? Predicting If Computers or Humans Should Segment Images

Danna Gurari

Suyog Dutt Jain

Margrit Betke

Kristen Grauman

(CVPR 2016, to appear)

Tackles the question of when to "pull the plug" on computer and human annotators in the context of **foreground object segmentations** by implementing two systems that automatically decide, for a batch of images, when to replace:

- 1. humans with computers to create coarse segmentations required to initialize segmentation tools and
- 2. computers with humans to create final, fine-grained segmentations.

Experiments demonstrate that a mix of human and computer efforts is better than relying on either resource alone.

Our work: Guess That Face!

• Guess That Face!

[Marques, Snyder, and Lux (CHI 2013)]

 a face recognition game that reverse engineers the human biological threshold for accurately recognizing blurred faces of celebrities under time-varying conditions.









Oge Marques, Justyn Snyder Florida Atlantic University - Boca Raton, FL - USA





Mathias Lux Klagenfurt University - Klagenfurt, Austria

- Motivation:
 - Human vision: we are remarkably good at recognizing (severely blurred) famous faces
 - *Computer vision*: face recognition still not mature
 - Machine learning: what if we train algorithms with blurry faces instead?
 - Game play: SongPop



Oge Marques, Justyn Snyder Florida Atlantic University - Boca Raton, FL - USA





Mathias Lux Klagenfurt University - Klagenfurt, Austria

• The game:

- Player must analyze a series of randomly-generated images of celebrities while the images transition from an initially severely-blurred state to their original state over a constant interval of time (~ 8 seconds).
- While the image is being progressively de-blurred on the screen, players are prompted to select the name of the celebrity who they believe is correct, which they do once they have confidence in their answer.





Oge Marques, Justyn Snyder Florida Atlantic University - Boca Raton, FL - USA





Mathias Lux Klagenfurt University - Klagenfurt, Austria

- Datasets:
 - Original: 48 popular celebrities + 2 politicians
 - Hollywood stars
 - Sports celebrities
 - Each image in the dataset also has two variations: de-saturated and horizontally-flipped
- Deblurring example:





Oge Marques, Justyn Snyder Florida Atlantic University - Boca Raton, FL - USA





Mathias Lux Klagenfurt University - Klagenfurt, Austria

• Results:



Guess That Face

 Live prototype at: <u>http://tinyurl.com/guessthatface</u>



Summary / Lessons learned

- Not yet at "viral" / "addictive" level
- Higher response rates for Ask'nSeek came as a result of bonus points for a class
- For GTF, a different reward system gave extra points to students who placed in the 'High scores' table. As a result:
 - Many more games than expected
 - Some players "memorized" the dataset
- For Click'n'Cut the quality of the work(ers) is the bottleneck

Challenges

• "Gamification" as a buzzword (or "cure-all")

DILBERT

BY SCOTT ADAMS



Challenges

- Games should offer *intrinsic* motivation
 - Competition
 - Self-drive
 - Self-improvement
 - Pleasure

• Games should be inherently *fun*!

Recommended reading

"Getting Gamers : The Psychology of Video Games and Their Impact on the People who Play Them"

(http://www.psychologyofgames.com/book/)



The Psychology of Video Games and Their Impact on the People who Play Them



- Why do normal people become raving lunatics online?
- Why do players cheat and peek at strategy guides?
- Why are fanboys and fangirls so ready for a fight?
- Why do we get nostalgic about good old games?
- How do games get us to keep score and compete?
- How do games get us to grind & chase achievements?
- How do games keep us excited for new loot?
- How do games immerse us in imaginary worlds?
- Why do we go crazy for digital game sales?
- How do mobile games get you with in-app purchases?
- How do games keep players paying and playing?
- How do games get players to market to each other?
- Do we shape our game avatars or do they shape us?
- Is video game violence a problem? Why do we like it?
- Do video games make you smarter?

Disclaimer





















 Psychology is the key to making games that are more fun, that get played for longer before being shelved or traded in, and that get talked about more.



The Psychology of Video Games and Their Impact on the People who Play Them



JAMIE MADIGAN

- On scoring and competing against others:
 - Everything we do is compared to or related to other players. Pure numbers [...] have limited meaning to us on their own.
 - [...] it's whom you're comparing yourself against that matters the most.
 - Leaderboards, achievements, scores, and replays built into modern games provide a great context in which to apply *social comparison theory* (*Leon Festinger, 1954*) and the way that developers can use it to keep you playing.

- On scoring and competing against others:
 - [...] strange thing about human psychology in the context of competition: not all rungs on a tournament ladder are equally spaced.
 - [...] the difference between first and second is bigger than anything, except maybe the distance between last place and second to last.

- On scoring and competing against others:
 - The "big fish, little pond effect" describes how we feel better about our performance if it puts us near the top of a low-performing group than if the same performance puts us near the bottom of a highperforming group.
 - [...] developers should provide opportunities for being at the top of different rankings (wealth, gear, etc.).
 - [...] any non-arbitrary goalpost will do (e.g., the concept of *par* in golf)

- On estimating our abilities:
 - The Dunning-Kruger effect (1999) describes how people who are bad at something overestimate their performance, and those who are experts underestimate.
 - Single-player campaigns and tutorials may highlight this effect by making novice players think they're better prepared for competitive play than they really are.

- On what keeps players playing:
 - "Self-determination theory (SDT)" (Przybylski, Rigby, and Ryan, 2010): a framework for understanding why people are motivated to keep playing games.
 - We do it to satisfy needs for:
 - -**Competence:** relates to feeling like we're doing well and getting better.
 - –Autonomy: is satisfied when we feel we are given meaningful decisions to make.
 - -**Relatedness:** happens when we feel that we are important to others.

Case study: SongPop

- Song-guessing game
 - Award-winning
 - Multiplatform
 - Highly addictive
- Latest version (SongPop 2)
 - 100,000+ song clips
 - 1,000+ playlists
 - Practice mode (with Melody, the mascot)

Case study: SongPop

- Extremely well-designed!
- Fast, sleek, intuitive, addictive
- Taps onto emotional aspects evoked by music
- Several related games
 - SongPop Party (for AppleTV)
 - Rock On (spin-off with 3,000+ rock songs)
 - TravelPop (images)
 - MoviePop (video clips)

Case study: SongPop

• Very mindful of psychology of gamers

https://youtu.be/-uriVSmd-uk

- Worst advice I could (but will not) give:
 - Gamify everything!

Prof. Dr. Mathias Lux ITEC - Klagenfurt University

- Reasonable pieces of advice:
 - Don't try to gamify if you're not a gamer.
 - People work better when things are *fun* for them.

Select multimedia problems worth researching

• Engage in research on game effectiveness: e.g., how to find out if people are having fun?

- Be mindful of new devices and technologies
 - VR kits
 - Sensors
 - Wearable gadgets

 Challenge the design workflow (and *turn it upside-down*!)

- Immerse in the gaming world
 - Psychology of games
 - Game design
 - Marketing and commerce

• "Recipe for success"

- Games that are well-designed (look like games, have a certain addictive component, make you want to share with friends, have a long shelf life)
- Games that are fun to play
- Game logs that convey useful information that would not be easily obtained otherwise
- Meaningful (open) multimedia problems
- Sound machine learning strategies to leverage the knowledge acquired through game logs.

Pop Quiz

- Based on everything you heard from me, which of these 2 options you believe I'd be more likely to choose:
 - 1. Gamify Click'n'Cut

or

2. Partner with FreshPlanet for creating a new game (FacePop, anyone?)

Final reflection: SWOT analysis

Final reflection: SWOT analysis

• Strengths

- There are problems to be solved
- People love games
- Weaknesses
 - Poorly designed games turn people away
- Opportunities
 - Multiple game platforms
 - Growing interest in games
- Threats
 - Better solutions without using games

Let's get to work!

 Which multimedia problem would you like to solve using games and/or crowdsourcing?
 – Contact me with ideas!

Oge Marques, Ph.D. Professor

Computer & Electrical Engineering and Computer Science (CEECS) 777 Glades Road Boca Raton, FL 33431-0991 tel: 561.297.3857 fax: 561.297.2800 email: omarques@fau.edu Skype: ProfessorOge