
Amazon Mechanical Turk

Best Practices Guide



Amazon Mechanical Turk: Best Practices Guide

Copyright © 2008, 2009 Amazon Web Services LLC or its affiliates.

Guidelines for Planning, Designing and Testing Your HITs

Topics

- [Planning Stage \(p. 1\)](#)
- [HIT Content \(p. 4\)](#)
- [Get work done faster \(p. 6\)](#)
- [Preflight Checklist \(p. 9\)](#)

Amazon Mechanical Turk is a web service that uses human thinking and consensus as resources. The service provides a virtual marketplace where requesters advertise micro tasks and human workers complete them. Amazon Mechanical Turk is based on the idea that there are still many things that human beings do more efficiently than computers, such as giving opinions, making judgments, and extracting facts. Amazon Mechanical Turk gives businesses immediate access to a diverse, geographically dispersed, on-demand, scalable workforce and gives workers a selection of thousands of tasks to complete whenever and wherever it's convenient.

We recognize that there is an art to using Amazon Mechanical Turk effectively. This guide helps you learn this art more quickly. Use this guide as you plan, design, and test your Amazon Mechanical Turk HITs.

Planning Stage

Topics

- [Break your work into HITs \(p. 2\)](#)
- [Keep HITs simple \(p. 2\)](#)
- [Tailor the workforce \(p. 2\)](#)
- [Make the most of the results \(p. 3\)](#)
- [Leverage the knowledge of the community \(p. 3\)](#)

- [Test your HITs \(p. 3\)](#)
- [Designate an Owner \(p. 3\)](#)

This section presents topics to complete before creating your HITs.

Break your work into HITs

In the planning stage, you must define the goal of your HITs and specify exactly what the Workers will do to help you achieve your goal. Each Worker task typically becomes one HIT. For example, to create a restaurant guide you could request that Workers categorize a list of restaurants. A better approach, however, would be to break each task into a separate HIT. For example, you could present the name of only one restaurant in a HIT and ask the Worker to categorize it.



Tip

Think of each HIT as a row in a spreadsheet. Each answer in a HIT is a column in the row. Multiple answers in a single HIT correspond to multiple columns in a single row.

Keep HITs simple

A single HIT can ask for more than one answer. For example, you might present four photos for a Worker to tag in one HIT. In this case, however, the Worker should be looking for the same thing in all of the photos. If you find that you are asking the Workers to perform multiple tasks in a single HIT, such as asking the Worker to look for different things in the four photos, break the tasks into multiple HITs. For example, if your HIT asks Workers to specify if there are restaurants and historic sites in photos, break the HITs into two HITs: one about tagging photos showing restaurants, another about tagging photos showing historic sites.



Tip

Remember that you pay the same amount for each HIT so keep the work across HITs consistent.

Tailor the workforce

Amazon Mechanical Turk gives you the opportunity to qualify your workforce. For example, if your HIT requires special knowledge, such as the ability to accurately classify artwork, you can create a qualification test and select only Workers that categorize artwork correctly. You can create qualification tests with answer keys that automatically score the tests or create qualification tests that require you to review the tests and approve the Workers.

Amazon Mechanical Turk automatically maintains some statistics on Workers, such as the percentage of assignments approved and the Worker's country. You can use these statistics to qualify Workers. For a complete list of system Qualifications or to create your own, go to "Qualification Requirement" or the command line tool, respectively, in the *Amazon Mechanical Turk Developer Guide* (<http://docs.amazonwebservices.com/AWSMechTurk/latest/AWSMechanicalTurkRequester/>).



Tip

To get the best selection of Workers, we suggest using Workers that have an approval rating of 95% or higher.

If you conduct a survey and need Workers to respond only once, create HITs with multiple assignments. Amazon Mechanical Turk ensures that each Worker can accept only a single assignment in a HIT.

Make the most of the results

Amazon Mechanical Turk presents HIT results in several ways depending on the way in which you create the HITs. To review the results in Excel, use the web interface to create your HITs and download the results in CSV format. To get XML results, use the API to create HITs and download results. For information about the API, go to the [Amazon Mechanical Turk Developer Guide](#).



Tip

If you use the web interface to create your HITs, you can review and approve or reject work offline. To do so, export your results, mark the answers you want to reject, and upload the file back to the web interface.

Leverage the knowledge of the community

Some Workers like to share their thoughts about your HITs. Often times, Workers have great insight into ways to improve your HITs. Providing a comment box like the following on each HIT lets Workers provide thoughts on each task they complete.

Please provide any comments you may have below, we appreciate your input!



Tip

Many Workers communicate with each other through an online message board, for example, www.turkernation.com. Consider introducing yourself in the message board and asking for feedback and suggestions.

Test your HITs

Testing your HITs helps you discover technical errors, inconsistencies in instructions, and potential improvements to your HIT. Before publishing 10,000 new HITs, consider publishing 10. This gives you an opportunity to revise your HIT based on feedback from Workers. You can also see if the format of the answer is acceptable. Publish the full set of HITs once you are comfortable with your HIT design and test results.



Tip

After publishing your test HITs you may not be able to find them on the website. Often, this is because Workers have already completed or have accepted your HITs. If you would like to see your HITs as Workers see them, use the sandbox (<https://requestersandbox.mturk.com>) for testing. You won't be charged and there aren't any Workers. To view your HIT as a Worker, go to <https://workersandbox.mturk.com> and search for your HIT. Test your HITs in all major browsers (Internet Explorer and Firefox) to make sure your HIT has the same functionality.

Designate an Owner

Our most effective and successful Requesters have one thing in common: an Amazon Mechanical Turk administrator. The administrator interacts with the Turker community, receives feedback about

the HITs, responds to questions, designs new HITs, and organizes the results. Investing in an administrator pays off because Workers become more involved, do the work quicker and go the extra mile for these administrators.



Tip

If you use the Amazon Mechanical Turk API, plan to dedicate approximately 1 hour for every 10,000 published HITs. Using the web interface often involves more time. Test your HITs in small batches to get an accurate estimate of how much time it takes to manage your work.

HIT Content

Topics

- [Write a clear and compelling title and description \(p. 4\)](#)
- [Write effective instructions \(p. 4\)](#)

This section gives tips about creating effective HIT content.

Write a clear and compelling title and description

The HIT title and description are the first items a Worker sees when searching for HITs. Take time to make the title and description concise, engaging, and descriptive. For example, if your HIT asks workers to give opinions about movies, a vague title would be, "Survey." A better title would be "A Quick Movie Survey." Follow it with a clear description, such as "Give us your opinion on two questions about a movie based on its summary."

Write effective instructions

The answers you receive from Workers are only as good as the instructions you provide. The following tips help you create effective instructions.

Make instructions simple and concise

Short sentences that contain a single idea increase their readability and help you get the answers you want. Work to streamline your instructions without losing clarity.

The following is an example of long, complicated, and poorly conceived instructions:

Your job is to go to any resource you like and find restaurant reviewers in the specified area, include their names, addresses, and any other contact information you can find, and then specify whether they work for a particular publisher and, if so, for how long.

The following instructions are better:

1. Go to this link to view an article in the Food and Dining section of the NY Times.
2. Copy the name of the reviewer from the by-line and paste it into the following box.
3. Click on the name of the reviewer to see a list of his or her reviews.

4. Copy the URL of the list and paste the URL into the following box.



Tip

Think of your instructions as a recipe that leads Workers step by step through the tasks they must perform. Number the instructions to show your Workers how many steps they must complete for each HIT and the exact order in which to complete them.

Include sample answers

Include a sample answer for each task. Examples help make your expectations clear. Here's an example: "Tag each photo with the name of the city, state and country where the photo was taken, for example, Seattle, Washington, USA." In this example, if you expect to see the state abbreviation ("WA") instead, be explicit in your instructions.

Identify things you do not want in the answers

Specify things you do not want in the answers to clarify what you do want. For example, "Do not put the name of the structure in the photo in the response, for example, Eiffel Tower, Paris, France." Use negative examples sparingly. Using one negative example to prevent a common mistake is helpful. Using lots of negative examples often confuses Workers.

Examples of improved instructions

The following examples show a poor and well written set of instructions.

The following shows a poor instruction set.

Transcribe Newspaper Article

Type the newspaper article below, including capitalization, spaces and punctuation.

The following is a better instruction set.

Transcribe a Short Section of a Newspaper Article

Type the snippet of newspaper article below, accurately reproducing capitalization, spaces and punctuation. Click the Instructions link to expand the list of instruction.

Instructions

Join words that start at the end of one line and finish at the start of the next line.

- 1.Type each paragraph on a single line. Start a new paragraph on a new line.
2. Do not enter any blank lines.
3. Put a space between initials: e.g., N.Y. -> N. Y.
4. Use -- for a long dash, e.g., Jan 21. -- New York. Also, 6-2 should be 6 -- 2.
5. Use - for a hyphen, e.g., ex-member.
6. Return or skip the HIT if you cannot read the text.

Payment

Answers are graded automatically as follows

1. Each article is typed by three different workers.
2. We automatically compare the answers against each other using an "edit distance" algorithm.

3. Only the two answers closest to each other will be paid. If there is a tie, all answers will be paid.

Get work done faster

Workers appreciate Requesters who help them be more efficient. Making Workers more efficient effectively increases their pay rate. This section discusses ways to make your HITs more efficient and appealing to Workers.

Use relevant keywords

You should associate keywords with your HITs. Keywords help Workers find your HITs faster when searching for work. Keywords should be relevant to the HIT. Use your Requester name as a keyword.

The following is an example of poorly chosen keywords.

Where is your favorite place to vacation? [View a HIT in this group](#)

Requester: Amazon Requester Inc.	HIT Expiration Date: Feb 6, 2009 (6 days 23 hours)	Reward: \$0.05
	Time Allotted: 60 minutes	HITs Available: 1

Description: Your favorite spot

Keywords: [easy](#), [favorite](#)

Qualifications Required:
HIT approval rate (%) is not less than 95

The following is an example of well chosen keywords.

Where is your favorite place to vacation? [View a HIT in this group](#)

Requester: Amazon Requester Inc.	HIT Expiration Date: Apr 28, 2009 (6 days 23 hours)	Reward: \$0.05
	Time Allotted: 60 minutes	HITs Available: 1

Description: We want to know your favorite vacation destination and why

Keywords: [travel](#), [vacation](#), [favorites](#), [opinion](#)

Qualifications Required:
HIT approval rate (%) is not less than 95

Avoid redundancy

Make sure you ask a question only once.



Tip

To increase Worker efficiency, use radio buttons instead of drop down menus to present the answers.

The following example is poor because the questions are redundant.

Does this picture include a car?

YES NO

What type of car?

4 Door

2 Door

To remove the redundancy, eliminate the first set of radio buttons and include a “No car in the picture” option.

What type of car is in the picture?

4 Door

2 Door

No car in the picture

Don't make Workers scroll

Scrolling reduces Worker efficiency and effectively reduces their pay rate. Find a way to make your HIT fit on one page. Make long instructions collapsible and keep them collapsed by default.

Open web pages in a new window

If a Worker must open a new window to complete your HIT, open it in a new browser window so that the Worker can continue working on the HIT in the current window while referring to the data in the new window.



Tip

To make a hyperlink open in a new window, include the target attribute in the HTML. For example: `Amazon Mechanical Turk`

Position answer fields strategically

If you want to associate an answer with content on the page, such as an image or text, place the answer field close to the content. By placing answer fields adjacent to the images they're associated with, you reduce the possibility of Workers entering answers in incorrect answer fields.

If your HIT does not have images or other content associated with each answer field, consider placing the answer fields and Submit button in close proximity to reduce the amount of mousing required thereby increasing a Worker's efficiency.

Provide separate answer fields for each answer

When using the web interface, Amazon Mechanical Turk stores each answer as a separate column in an Excel spreadsheet (a CSV file). By breaking Worker's answers across multiple columns you can more easily review results. Below are examples of using separate answer fields versus a single answer field.



Tip

If you use the web interface to create HITs, remember to give each answer field a unique name. Amazon Mechanical Turk merges fields with the same name into a single column in the results.

Example of separate answer fields using different column names

WorkTimeInSeconds	Answer.Q1FirstName	Answer.Q2LastName	Answer.Q3Employer
25	Larry	Jones	Widgets, Inc.
615	Roger	Evans	Widgets, Incorporated
73	Jane	Smith	Widgets

WorkTimeInSeconds	Answer.Q1FirstName	Answer.Q2LastName	Answer.Q3Employer
104	Joan	White	Gadgets, LLC
104	Bill	Bradley	Gadgets

Example of single answer field

WorkTimeInSeconds	Answer.Q1Name
25	Larry Jones Widgets, Inc
615	Roger Evans Widgets, Incorporated
73	Jane Smith Widgets
104	Joan White Gadgets, LLC
104	Bill Bradley Gadgets

Give Workers enough time to do the work

The HIT property, *Time Allotted*, is not the time you think it should take to do the work; it is the time a Worker can hold onto a HIT before it is taken away so another Worker can work on it. Workers get frustrated if an assignment is taken away from them because the *Time Allotted* value is impractically short.



Tip

If you need results in near real-time, set *Time Allotted* to 2 or 3 minutes. In general, use a 30 or 60 minute setting for quick (but not real time) results.

Set equal work for equal pay

All HITs within a group of HITs that pay the same should require approximately the same amount of work. If a HIT in a group pays one cent to answer one question, another HIT in that group should not ask the Worker to answer four questions for the same reward. Following this guideline, don't make the amount of work dependent on an answer in the HIT, as shown in the following example.

Does this picture include a car?

YES NO

If yes, please enter the make, model and year of the car below.

Make

Model

Year

Instead, create two separate HITs: one that asks if there is a car in the picture and a second HIT for pictures that include cars in order to collect the additional information.

Pay fairly

We highly recommend that you look at HITs similar to yours to determine an appropriate reward. Remember that Amazon Mechanical Turk is a marketplace and that Workers decide what work to do

based on what other work is available. Don't expect a Worker to complete a 1 hour video transcription for \$0.07 when 5 minute audio transcriptions offer \$0.05. Think of the reward as an hourly rate instead of a reward per HIT.



Tip

Use the following formula to determine your reward per HIT based on a target hourly rate:

$$\text{(Minutes to complete HIT)} \times \left(\frac{\text{Rate per hour}}{60 \text{ minutes}} \right) = \text{Reward per HIT}$$

For example, to pay \$8/hr for a HIT that takes approximately 2 minutes, the reward should be 2 minutes X \$8 / 60 = \$0.25 per HIT. Remember to include in your time estimates per HIT the amount of time it takes to load the page and navigate through your HITs.

Be clear whether you want an opinion, judgment, or fact

In general, there are three kinds of questions you can ask of Workers:

- **Opinion**—For example "Did you enjoy the film The Godfather?"
When asking for a Worker's opinion, as long as the answer is relevant to the question, the Worker should get paid.
- **Judgment**—For example, "Is the film The Godfather worthy of an oscar?"
If you are asking a judgment question and use multiple assignments (plurality) to help adjudicate results do not automatically reject work that does not meet the majority. Automatic rejection frustrates Workers who may not have agreed with others. Consider granting bonuses to Workers who agree with the majority.
- **Fact**—For example, "What is the MPAA rating (X, R, PG, G) for the film The Godfather?"
When you are asking for a fact, make clear that you do not give rewards for inaccurate answers.



Tip

Provide clear guidelines to Workers for getting paid. Advise Workers if you will use plurality or multiple answers to review results.

Preflight Checklist

Use the tips in this section as a final checklist before launching your HITs.

Make sure your instructions are clear and concise

Before publishing your HITs, give your instructions, without explanation, to a fellow worker and make them follow them. See if the worker proceeds in the way you intend. Make sure your instructions are clear, specific, and unambiguous.

Check your links

If you have links to pictures or videos in your HITs, make sure the links work. The best way to check the links is to actually complete some of your HITs using the Amazon Mechanical Turk sandbox.

Use the Amazon Mechanical Turk sandbox to test your HITs

Test your HITs yourself end-to-end. Use the Amazon Mechanical Turk sandbox to test the workflow of your HIT, complete some assignments, and view the results. Consider whether the results enable you to achieve the overall goal of your project.

Do a trial run!

Completing HITs yourself is good but having real Workers complete HITs can reveal subtle errors. Publish ten HITs, review the results, and increase the number of published HITs as your confidence in the results increases. Don't launch with hundreds or thousands of HITs without testing. Testing can be a real time and cost-saver!

Guidelines for Managing Your HITs

Topics

- [Worker Based Acceptance \(p. 11\)](#)
- [Response Speed Acceptance \(p. 11\)](#)
- [Building a Group of Loyal Workers \(p. 12\)](#)

After you publish your batch of HITs, you, as the Requester, should monitor the results for a number of reasons, including taking action to reject work that fails to meet your standards before that work is automatically approved and thus the Worker is compensated. This guide gives you insights for managing your HITs.

Worker Based Acceptance

After inspecting many results completed by one Worker you can form an opinion about the quality of their work. You might come to know the Worker through the current batch of HITs or from a previous batch.

When you know a Worker typically does good work, the user interface enables you to approve all of their assignments at once. Conversely, when you find a Worker who typically submits substandard results, you can reject all of their assignments at once. You can also block that Worker from working on any more of your HITs.

You can guard against Workers submitting substandard work by setting up a qualification requirement. For example, you can set the qualification that only Workers that have a HIT approval rate greater than 95% can view your HITs.

Response Speed Acceptance

One of your most important tasks is to determine whether the work completed is good or not. One way to do that is to check the speed with which an assignment is completed. The generalization is that assignments completed quickly were either done by a machine or by a Worker who didn't care enough

to enter correct answers. The *Amazon Mechanical Turk User Interface Guide* explains how to filter results based on response speed, called Velocity, in the **Filter Results** page, as shown.

Filter Results

Review only the assignments that meet the following criteria:

Velocity Filter -- only show assignments submitted in seconds or less (average submit time was 36 seconds).

Status Filter -- only show assignments that are in status.

There are no hard and fast rules when it comes to determining how fast is too fast. Some Workers are fast and good. Some are fast because they're entering substandard results. Look at the work performed by Workers who respond quickly to make sure their work is satisfactory.

Building a Group of Loyal Workers

Building a group of Workers loyal to you benefits you, the Requester, in several ways. First, Workers who like working with you will accept your HITs readily and thus you won't have to pay more to encourage Workers to complete your HITs. Second, you can count on good Workers to give you good results. But how do you build a group of loyal Workers?

Use the following series of tips to build a positive working relationship with Workers.

- **Pay quickly.**—Workers appreciate getting paid in a timely manner for completing work.
- **Give bonuses**—Workers deserve bonuses when they accurately complete a great amount of work in a good amount of time. If the worker is someone you'd want to work with again, consider giving them a bonus.
- **Give feedback**—When a worker does a good job tell them. When a worker doesn't do a good job, explain to them what was wrong with the work so they can do better in the future.
- **Tell your workers what the goal of the batch of HITs is**—When workers feel like they are participating in an important project they take more pride in their work and are more likely to want to work with you again. Include, in your HIT description, the goal of your project.