

University of Notre Dame  
Lucy Family Institute for Data and Society

---

UNITED AGAINST SOCIAL MEDIA  
COMMENT ANNOTATION GUIDELINES

---



Questions? [uoa.annotation.study@gmail.com](mailto:uoa.annotation.study@gmail.com)

June 10, 2026

# Contents

<b>1 Quick Start</b>	<b>2</b>
<b>2 Task A — Marking entities</b>	<b>2</b>
2.1 Entity heads and mentions . . . . .	3
2.2 Explicit vs. implicit — and where the head goes . . . . .	3
2.3 Entity type & lookup . . . . .	3
<b>3 Task B — Classifying entities</b>	<b>4</b>
3.1 Abuse vs. criticism (the key distinction) . . . . .	4
3.2 Rationale spans . . . . .	5
3.3 The labels (reference) . . . . .	5
<b>4 Worked example (both tasks)</b>	<b>6</b>
<b>5 Golden rules</b>	<b>7</b>
<b>6 Stage B — interface screenshots</b>	<b>7</b>

# 1 Quick Start

This task supports a University of Notre Dame research project on online abuse in college-sports discussions. You read short social-media comments and mark *who* is being discussed and *how*. You will do **two kinds of short tasks**, mixed together within this HIT. The portal shows you one comment at a time and tells you which task you are on:

- **Task A — Mark the entities:** highlight each player/team/etc. mentioned, link repeated references to the same one, and set its type. (*No questions, no rationale.*)
- **Task B — Classify the entities:** the entities are already marked for you; answer a few yes/no questions about each and highlight the text that justifies your answers.

**Before you start, watch the 3–5 minute walkthrough:** [Tutorial video](https://www.youtube.com/watch?v=ch_iTN-TIeg) ([https://www.youtube.com/watch?v=ch\\_iTN-TIeg](https://www.youtube.com/watch?v=ch_iTN-TIeg)).

**Three rules that apply to everything:**

1. You annotate the **target comment**, using the **parent** and **submission/root** (when shown) as context to understand who/what is meant. **One exception:** if an entity is referred to in the target only *implicitly* (a pronoun/description) and is named only in the parent or submission/root, select its *entity head* on that name in the parent/submission (and mark it recoverable from history). Everything else — mentions, rationales — stays in the target comment.
2. Annotate an entity only if the target comment actually discusses, praises, blames, or attacks it (not every name that appears).
3. Keep highlighted spans **minimal but sufficient**.

You can work with the mouse, or use the keyboard: ↑/↓ move between rows and ←/→ pick an option.

## 2 Task A — Marking entities

In a Task-A comment you identify the entities and how they refer to each other. You do *not* answer any abuse/label questions here — that happens in Task B.

**Important:** the spans come pre-filled with a **model suggestion** — it is only a *starting point*. Add the spans/links it missed and remove the ones it got wrong. **Do not just submit the suggestion unchanged;** rubber-stamping it is flagged for review.

**Annotate ONLY the entities the comment actually discusses or attacks — not every name it mentions.** A player or team that is merely *named in passing* (listed, referenced as background, or mentioned only to set the scene) and is **not itself talked about, praised, criticised, or attacked** must be **skipped:** no head, no mention, no type. When in doubt, ask “*does the comment say anything about this person/team?*” — if not, don’t annotate them.

*Example:* “She carried us. Madison Booker is straight trash, hope she blows out her knee, even with Mabrey and Prince out there.” → annotate **Olivia Miles** (via *She*) and **Madison Booker**; **do NOT** annotate *Mabrey* or *Prince* — they are only named as being on the court and nothing is said about them.

## 2.1 Entity heads and mentions

- **Entity head:** the main reference to an entity (usually its name) — a player, team, coach, referee, fanbase, etc. Each distinct entity gets exactly *one* head per comment.
- **Entity mention:** any other reference to the *same* entity — a pronoun or description such as *she, they, that guard, the refs*.
- **Coreference link:** connect each mention to its head (mention → head).

Example: *Olivia Miles* (head) ← *she, her* (mentions); *UConn* (head) ← *they, that team*.

## 2.2 Explicit vs. implicit — and where the head goes

The entity head is normally in the target comment, but **not always** — read these three cases carefully, they are the most common mistake:

- **Explicit (named in the target):** select that name in the target as the head. E.g. *Lauren Betts, UCLA*.
- **Implicit but resolvable from context:** the target uses only a pronoun/description, and the entity is named in the **parent** or **submission/root**. Select the head *on that name in the parent/submission*, link the target's pronoun to it, and mark *target recoverable from history = parent* (or *root\_submission*).  
E.g. **Parent:** “Olivia Miles controlled the whole game.” **Target:** “She was incredible.” → head *Olivia Miles* (in the parent), *She* → *Olivia Miles*.
- **Implicit and not resolvable:** a pronoun/description that nothing in the target, parent, or submission lets you identify. Select the implicit reference *itself* in the target as the head and mark *recoverable from history = no*.  
E.g. “He sold the game.” → head *He*.

## 2.3 Entity type & lookup

For each head, set its **type**: *player NCAA, player non-NCAA, coach, fan(s), team, referee, or other*. Then:

- **player NCAA** → fill the **Player Lookup** (search panel) to link the canonical NCAA player.
- **team / fan(s)** → fill the **Team Lookup** (required for teams; optional but must be valid for fans).
- **Can't uniquely identify it?** If the comment and the shown context don't let you pin down *which* specific player or team it is, do **not** guess — leave the lookup empty and tick “**Cannot uniquely identify this NCAA player**” (or “**...this team**”). The type still gets recorded; only the exact ID is left unresolved.

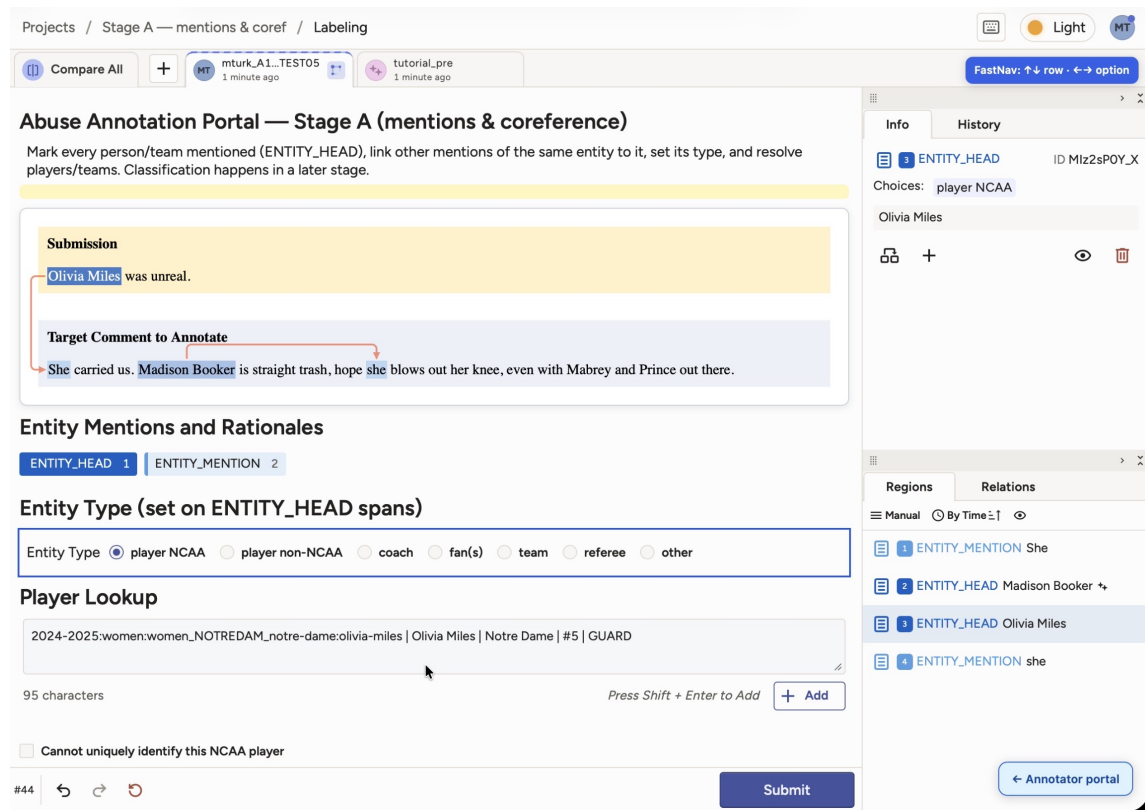


Figure 1: Task A (Stage A) on the worked example. *Olivia Miles* is marked as an ENTITY\_HEAD in the submission, and the mention *She* in the target is linked to it with a coreference relation; *Madison Booker* is a second head. The selected head’s **Entity Type** is set to *player NCAA*, and the **Player Lookup** is resolved to the canonical player (*Olivia Miles, Notre Dame*).

### 3 Task B — Classifying entities

In a Task-B comment the entity heads, mentions, and their links are **already marked for you**. **Do not delete, move, or change any entity head/mention span or any of their relations** — leave them exactly as they are. Your job is only to: (1) answer a short set of yes/no questions for *each* entity, and (2) add **rationale** spans and link each one to the entity head or mention it supports. All labels are applied **independently per entity**.

#### 3.1 Abuse vs. criticism (the key distinction)

Not every negative comment is abusive. Mark **Abusive = YES** only for insulting, degrading, threatening, hateful, or clearly hostile language toward the entity. Criticism of performance or outcome is **not** abuse.

- Non-abusive criticism: “She missed too many shots tonight.”
- Abusive: “She was straight trash tonight.”

- Abusive + Wish of Harm: “I hope he gets hurt.”

### 3.2 Rationale spans

A **rationale** is the minimal text that justifies your labels for an entity. Highlight the smallest span(s) that clearly support your decision (don’t select whole sentences unless needed), and **link each rationale to the entity it supports**. An entity may have several rationales; a span may support more than one entity only if it genuinely justifies each.

### 3.3 The labels (reference)

**Entity core.** Abusive (YES/NO); Severe Abuse (only if Abusive = YES). If Abusive = NO, all abuse-type labels are NO (and stay hidden).

**Abuse types** (only when Abusive = YES):

Label	Choices	Description
Race Hate	YES / NO	Abuse targeting race or ethnicity.
Gender Hate	YES / NO	Abuse targeting gender.
Religious Hate	YES / NO	Abuse targeting religion.
Wish of Harm	YES / NO	Expresses desire for harm or injury.
Threat	YES / NO	Threatening language toward the entity.
Insult / Derogation	YES / NO	Direct insults or degrading language.
Sexual Harassment	YES / NO	Sexualized or inappropriate harassment.

#### Form / pragmatics:

Label	Choices	Description
Profanity / Obscene	YES / NO	Presence of profanity or explicit language.
Sarcasm	YES / NO	The statement is sarcastic or ironic.
Humor	YES / NO	The statement is intended as a joke.
Quotation / Reported Speech	YES / NO	Includes quoted or reported content.

**Attribution / motive** (what the discussion is about):

Label	Choices	Description
Performance	YES / NO	Play in the game.
Officiating	YES / NO	Referees or calls.
Gambling	YES / NO	Betting or gambling.
Fandom / Rivalry	YES / NO	Fanbases or rivalries.
Appearance	YES / NO	Physical appearance.
Personal Life	YES / NO	Personal life.
Injury	YES / NO	Injury or health.
Coaching	YES / NO	Coaching decisions.

#### History / resolution:

Label	Choices	Description
implicit target	yes / no	Entity referred to indirectly (pronoun/description) rather than by name.
target recoverable from history	parent / root_submission / no	Only if implicit: where the entity can be identified, or <i>no</i> if it cannot.
interpretation needs history	no / parent / root_submission	Whether understanding the comment requires the parent/submission.

*The Stage-B interface screenshots are at the end of this guide (see “Stage B — interface screenshots”) — read the rules above first, then refer to them.*

## 4 Worked example (both tasks)

**Submission/root:** “Olivia Miles was unreal.”

**Target comment:** “She carried us. Madison Booker is straight trash, hope she blows out her knee, even with Mabrey and Prince out there.”

Three names appear in the target — *Madison Booker*, *Mabrey*, *Prince* — but only two entities are actually discussed. *Mabrey* and *Prince* are just named as being on the court; they are not evaluated, so we do **not** annotate them.

#### Task A — mark entities:

- **Olivia Miles** — the target says only “She”, so take the head from the **submission/root** and link *She* → *Olivia Miles*. Type: player NCAA.
- **Madison Booker** — head in the target; link the mention *she* → *Madison Booker*. Type: player NCAA.

#### Task B — classify each entity:

- *Olivia Miles* (non-abusive): rationale *carried us* → the entity. Abusive = NO; Motive = Performance; implicit target = YES; target recoverable from history = root\_submission; interpretation needs history = no.
- *Madison Booker* (abusive): **two rationales for the same entity** — *straight trash* → the **head** (*Madison Booker*), and *hope she blows out her knee* → the **mention** (*she*). Abusive

= YES; Insult/Derogation = YES; Wish of Harm = YES; Threat = NO; Motive = Performance; implicit target = NO.

## 5 Golden rules

- Annotate only what the text and shown context support — don't add outside assumptions.
- Exactly one head per real-world entity per comment; link every mention to its head.
- In Task B, every label is per entity, and every rationale links to the entity it supports.
- **Annotate only entities the comment discusses or attacks — never every name.** Skip players/teams just named in passing (background, not evaluated).
- Be consistent across similar cases.

## 6 Stage B — interface screenshots

These show the Stage B card from top to bottom on the worked example. The entity spans are pre-loaded and **locked**; you classify one entity at a time with the < *prev* / *next* > pager, and add RATIONALE spans linked to the entity they justify.

Projects / Stage B — classification / Labeling

Compare All + mt 2 r < prev Entity 1 of 2 — Olivia Miles (player NCAA) next >

### Abuse Annotation Portal — Stage B (classify the marked entities)

The entities are already marked and LOCKED (from Stage A, shown in blue) — you cannot change them. One entity is shown at a time: set its labels (YES / NO / UNSURE), then use the < > arrows below the card to move to the next entity. To add evidence: pick RATIONALE, select the text, then link it to the entity it justifies.

**Submission**

Olivia Miles was unreal.

**Target Comment to Annotate**

She carried us, Madison Booker is straight trash, hope she blows out her knee, even with Mabrey and Prince out there.

ENTITY\_HEAD 1 ENTITY\_MENTION 2 RATIONALE 5

Figure 2: Stage B overview. Heads/mentions are pre-loaded and locked (blue); rationale spans (red) are linked to the entity they justify — *carried us* for Olivia Miles, and *straight trash / hope she blows out her knee* for Madison Booker.

Projects / Stage B — classification / Labeling

Compare All + MT < prev Entity 2 of 2 — Madison Booker (player NCAA) next > FastNav: ↑↓ row · ←→ option

**Submission**

Olivia Miles was unreal.

**Target Comment to Annotate**

She carried us. Madison Booker is straight trash, hope she blows out her knee, even with Mabrey and Prince out there.

Link a RATIONALE span to the entity it justifies

**Entity Core**

Abusive  YES  NO  UNSURE

Severe Abuse  YES  NO  UNSURE

**Abuse Type**

Race Hate  YES  NO  UNSURE

Gender Hate (misogyny+misandry)  YES  NO  UNSURE

Religious Hate  YES  NO  UNSURE

Wish of Harm  YES  NO  UNSURE

Threat  YES  NO  UNSURE

#45 < prev Entity 2 of 2 — Madison Booker (player NCAA) next > Submit < Annotator portal

Info History

View region details  
Select a region to view its properties, metadata and available actions

Regions Relations

Manual By Time: 1

- ENTITY\_HEAD Olivia Miles
- ENTITY\_MENTION She
- ENTITY\_HEAD Madison Booker
- ENTITY\_MENTION she

Figure 3: Entity Core. For *Madison Booker* the comment is **Abusive = YES** (and Severe Abuse = YES), which reveals the abuse-type questions below.

Projects / Stage B — classification / Labeling

Compare All + HT < prev Entity 2 of 2 — Madison Booker (player NCAA) next >

**Submission**

Olivia Miles was unreal.

**Target Comment to Annotate**

She carried us. Madison Booker is straight trash, hope she blows out her knee, even with Mabrey and Prince out there.

Severe Abuse  YES  NO  UNSURE

**Abuse Type**

Race Hate  YES  NO  UNSURE

Gender Hate (misogyny+misandry)  YES  NO  UNSURE

Religious Hate  YES  NO  UNSURE

Wish of Harm  YES  NO  UNSURE

Threat  YES  NO  UNSURE

Insult/Derogation  YES  NO  UNSURE

Sexual Harassment  YES  NO  UNSURE

**Form / Pragmatics**

#45 ↶ ↷ ↻ < prev Entity 2 of 2 — Madison Booker (player NCAA) next > Submit

Figure 4: Abuse Type (shown only when Abusive = YES). For the running example Insult/Derogation and Wish of Harm apply, while Race / Gender / Religious Hate, Threat and Sexual Harassment do not.

Projects / Stage B — classification / Labeling

Compare All + MT < prev Entity 2 of 2 — Madison Booker (player NCAA) next >

**Submission**  
Olivia Miles was unreal.

**Target Comment to Annotate**  
She carried us. Madison Booker is straight trash, hope she blows out her knee, even with Mabrey and Prince out there.

**Form / Pragmatics**

Profanity/Obscene  YES  NO  UNSURE

Sarcasm  YES  NO  UNSURE

Humor  YES  NO  UNSURE

Quotation/Reported Speech  YES  NO  UNSURE

Figure 5: Form / Pragmatics, asked for *every* entity regardless of abuse: Profanity/Obscene, Sarcasm, Humor, Quotation/Reported Speech.

Projects / Stage B — classification / Labeling

Compare All + MT mt sec < prev Entity 1 of 2 — Olivia Miles (player NCAA) next >

Profanity/Obscene YES

**Submission**  
Olivia Miles was unreal.

**Target Comment to Annotate**  
She carried us. Madison Booker is straight trash, hope she blows out her knee, even with Mabrey and Prince out there.

**Attribution Focus (Motive)**

Performance  YES  NO  UNSURE

Officiating  YES  NO  UNSURE

Gambling  YES  NO  UNSURE

Fandom/Rivalry  YES  NO  UNSURE

Appearance  YES  NO  UNSURE

Personal life  YES  NO  UNSURE

Injury  YES  NO  UNSURE

Coaching  YES  NO  UNSURE

Figure 6: Attribution Focus (Motive): what the discussion is about. Here **Performance = YES**.

Projects / Stage B — classification / Labeling

Compare All + MT mt 1.1

< prev Entity 1 of 2 — Olivia Miles (player NCAA) next >

**Submission**

Olivia Miles was unreal.

**Target Comment to Annotate**

She carried us. Madison Booker is straight trash, hope she blows out her knee, even with Mabrey and Prince out there.

Random/rivalry  YES  NO  UNSURE

Appearance  YES  NO  UNSURE

Personal life  YES  NO  UNSURE

Injury  YES  NO  UNSURE

Coaching  YES  NO  UNSURE

**History / Resolution**

Implicit target  YES  NO  UNSURE

Target recoverable from history  PARENT  ROOT\_SUBMISSION  NO

Interpretation needs history  NO  PARENT  ROOT\_SUBMISSION

#45 ↩ ↪ ↻

< prev Entity 1 of 2 — Olivia Miles (player NCAA) next >

Submit

Figure 7: History / Resolution. *Olivia Miles* is named only as *She* in the target, so **implicit target = YES** and **target recoverable from history = ROOT\_SUBMISSION**; the comment is understandable on its own, so **interpretation needs history = NO**.