

The Dungeon Master's Dashboard

Ryan Smith, rsmith0914@gmail.com

Ahmed Abu Zuraiq, abuzreg@gmail.com

Helena de Castro Alvarenga, h.alvarenga.bsb@gmail.com

1. Introduction

The fantasy tabletop role-playing game Dungeons and Dragons (D&D), has been a fun pastime since its publication in 1974 (Michaud, 2015). The gameplay (Heller, 2019) involves two sides, the player group (the party), and the Dungeon Master (DM). The DM creates a semi-structured open fantasy campaign that places the party in a world where they must complete a task (quest). These campaigns are typically held over the course of many sessions which allow the individuals to progress through the story as well as level-up their character's attributes and skills. Each party member creates a character that has a set of statistics (stats) connected to various attributes like strength and intelligence. The higher the stat, the more likely a character is able to successfully complete a related action, such as landing a blow in battle or convincing a guard to let them into a city. By using this combination of dice rolling and improvisational interactions, the party hopes to complete their campaign after several sessions.

During the campaign, it is important for the DM to track information about the party and the world. The DM must keep reference literature on hand, like a monster manual for selecting and creating monsters for the party to battle (See Appendix A), as well as each party member's character sheet, which tracks the individual character stats (See Appendix B), in order to ensure that every interaction in the fantasy world follows the established rules. Many of the interactions between characters and the world are a result of their character stats combined with their dice rolls. With this knowledge, the DM may choose to increase or decrease the difficulty of certain tasks, as well as place obstacles that force the party to work together and rely on each other's strengths. This dynamic creates a complex interplay between freedom of expression and established game norms that generate the core D&D gameplay loop.

Although the game has grown in popularity (Wieland, 2021), the player interface remains rooted in its tabletop 1974 beginnings. The DM may feel this antiquated process more than the party as they not only track each player's character but also the entire campaign since they are the ones who design and mediate the world. Constantly having to check various manuals, scan character sheets and reference their own notes, slows the pace of the game for the party and can lead to error by the DM. Our team believes that we can help increase the efficiency of a DM by creating visual aides that allow them to access and assess the game's information in a more readable format.

The DM dashboard is a very interesting problem because of the eclectic nature of the data and the amount of information that a DM needs to reference at any given point. The visualizations need to be flexible and quick to fit the improvisational nature of

the story, while still providing structure and direction for the DM to follow the rules of the world, creating an interesting dichotomy. In order to tackle these problems, we are leveraging our own experience with D&D which varies from casual play to having DMed campaigns for our friends. This experience provides insight into some of the tools that a DM may want, as well as the obstacles they face while monitoring their campaign. To help supplement our own experiences, interviewed our respective D&D groups and gathered additional insights. Although this visualization problem is not as serious as others, it provides an interesting opportunity to explore the field and improve our data visualization skills.

2. Related Work

The related work can be divided into three categories. The first is gameplay augmentation tools, whether digital or tangible/physical, that are aimed at supporting the gameplay experience. The goal can be to increase immersion or offload unwanted tasks (Hartelius et al., 2012, Plijaer et al., 2020). Ideally without taking away from the player's agency or hindering the gameplay.

The second category includes attempts at visualizing D&D historical gameplay data (e.g., a [dataset of character sheets created by past players](#)) with the purpose of understanding something about the choices people make (Soderlund, 2018, Berbesi 2019).

The third category is systems/websites aimed at supporting the DM over the course of a campaign, rather than being marketed at the players directly. Some of those are online platforms that can be used to run the games Roll20 or Fantasy Grounds, or maintain textual records of gameplay-related information (e.g., AdventureCodex) or provide the necessary material to run a campaign like the game rules and pre-made content (e.g., DnDBeyond). While others are more specific tools that DMs can use to generate maps (e.g. WorldAnvil and Inkarnate), model in-game combat parameters like damage per round (Esker 2020), or select the monsters appropriate for a given game session (Mais 2020).

In this project we are focusing on the third aforementioned category, in particular, tools that utilize some form of data visualization in support of DMs as they prepare for a campaign, or during a gaming session. Finally, we use the design study methodology as a guideline for our work (Sedlmair et al., 2012).

3. Data and Task Abstractions

To arrive at the following domain abstraction, we conducted an interview with 9 DMs to learn about the tools and processes they employed. We also relied on our own experience as players and DMs. You can find the questions and the raw interview data in the accompanying document.

3.1. Dice Module:

There are seven sets of dice that players and DMs usually use, ranging in side number from four (d4) to twenty (d20). The dice rolls are a discrete quantitative data set, with 20 typically being the “best” roll you can have in most situations. Many players are also very superstitious about their dice and tend to favor certain ones in certain situations based on which ones they deem “lucky”, as confirmed in the interviews. While we don’t have an exact number for how many sets of dice a player will have in rotation, we believe that it typically won’t be more than 5, with most players having 2-3. Since some players like to collect dice sets, we want to ensure that this module can scale. Considering both DMs and players regularly use dice, this module has a broader user pool and can be extended for player’s use also. Since most actions are performed with the d20 and most players tend to base their entire set’s luck on the d20’s rolls, we will only be tracking this dice. This decision helps streamline the process and prevents unnecessary tracking downtime.

To figure out which dice are their luckiest or best, the person will want to compare two different attributes. First, they will want to compare the distribution of their dice’s rolls, e.g. did they roll more 20s on this dice or more 1s? Second, they will want to compare this d20’s distribution to the distribution of their other d20s to see if there are noticeable differences. It should be noted that traditionally we would expect these dice to all have the same even distribution, but real world factors can influence these rates. Additionally, the interviews show that the players emphasize the recent rolling patterns of their dice. Sometimes a dice can go on an “unlucky” streak and will be replaced by another set for the session.

3.2. Party Tracker:

The party tracker will consist of multiple sub visualizations. The overarching task is to allow the DM to strategize their encounters by seeing what are the strengths and weaknesses of each individual and the party, as a whole, which will require look-up, filter, and comparing tasks. Additionally, parties typically have a low cardinality of 3-5 players per party, which provides an opportunity for less summation analytics on the dashboard and more raw data translations.

Each character has six ability scores, strength, wisdom, intelligence, dexterity, constitution, and charisma. Each of these scores ranges from 1-30 and is sorted into ability modifier bins, based on how high the number is. Ability modifiers are a conversion that turns these stats into increased or decreased modifications on d20 rolls (see appendix C for full list). For example, a character with a strength ability score of 16 will add 3 to any d20 roll that requires their character to use their strength ability. If they originally rolled an 18, they score a 21 for the roll with the 3 extra points. The DM or official rules will dictate which ability scores are used for each instance.

The DM will need a way to query each of these six ability scores and compare them to other party members so that they can choose how they would like to plan an encounter. If they want to increase the difficulty of an event, they may see that the party has low strength and put a boulder in their path that requires a high strength roll

(check) to pass. A filter function can be useful here as a means to quickly see who would likely have difficulty in not completing a task. It should be noted that the ability score does not dictate the outcome of an event. Someone who has low strength can still roll a high enough number to move the boulder, but it will be harder since they do not receive the extra points of someone with a high strength ability score. This discrete quantitative data set is unique, as the ability scores and modifier scores are intertwined and should be represented as such in the visualization. Ultimately, the modifier is the more important of the two pieces of information as it is the number that alters the interaction, but we believe the ability score should also be displayed as it increases while your character levels-up.

Additionally, the dashboard will include health data for each party member. This is also discrete quantitative data and will allow the DM to track how many hit points (HP) the party members have left. When a party member passes out, they are at risk of dying, thus potentially losing their character forever. Therefore, the DM will want to look-up each party member's health to see who is close to passing out and who still has a lot of HP. While no true range can be given, as HP can range from 0 (passing out) to as high as the game allows it, most player characters will have a maximum HP within 20 of each other. Similar to the other party tracker visualizations, this data can help dictate how the DM wants to decide on the difficulty of an encounter. Even though the HP visualization is a simple feature of the party tracker, it will be a nice addition to the tools which further allows the DM to control the flow of each conflict.

An experience points (EXP) visualization will also be included. The needed information will be discrete quantitative data with a range between 0 and an adjustable maximum that depends on character level. This mechanic would be reflected in the visualization, allowing the DM to know how long until the party levels-up and gets stronger, which means they can adjust their campaign accordingly.

Finally, the walking speed, armor class and passive perception of characters will be shown as a tabular data section. These values generally range from 15 ft to 40 ft (in increments of 5ft) for speed. Passive perception and armor class have a similar cardinality to stats . The task of comparing this data across characters can be easily executed with the support of a simple table.

3.3. Monsters Shopper

An encounter is a sequence of events where the players engage in an activity. An important type of activity is fighting monsters. When designing a fighting encounter DMs consider multiple criteria as they choose monsters, those include:

1. the goal of the encounter
2. the number of players
3. players' level vs. challenge level*
4. consistency with story
5. consistency with setting (e.g., no fire monsters in an arctic region, unless the surprise is intended)
6. what is enjoyable to DMs (DMs play too!)

* To calculate the challenge level of an encounter we need the Level of each character and the Challenge Rating of each monster. From the [DM Basic Rules](#): *“A monster’s challenge rating tells you how great a threat the monster is. An appropriately equipped and well-rested party of four adventurers should be able to defeat a monster that has a challenge rating equal to its level without suffering any deaths. For example, a party of four 3rd-level characters should find a monster with a challenge rating of 3 to be a worthy challenge, but not a deadly one”*

To accomplish that they use different tools but mainly:

1. [The official D&D Monster Manual](#): where monsters are sorted alphabetically by their name but an element of surprise is achieved by randomly browsing the book.
2. [DnDBeyond](#), which has a tool for building encounters that include some filtering widgets and a shopping cart of monsters.
3. Existing campaigns made by other DMs.
4. Their own knowledge about monsters from series, books, ..etc. As a start then looks for it in the above sources.
5. Opinions of experienced DMs
6. Other websites:
 - a. <http://dnd5e.wikidot.com/>
 - b. <https://www.themonstersknow.com/>

To determine if a monster fits the above criteria, DMs look at the monster's stats, along with images of it and possibly some description. Refer to Appendix A for an example. Based on the taxonomy in the official Monsters Manual, each monster has ability scores which are quantitative dimensions indicating strength, dexterity, constitution, intelligence, wisdom and charisma. Quantitative dimensions on monsters also include their armor class, health points and speed. All the quant dimensions are positive with no upper limit for monsters (even though limits exist for players’ stats). Monsters have multiple categorical and ordinal dimensions that include their name, size (6 levels), type (14 levels), movement mode (4 levels, but can have no value too so 5) levels but can have any combination or none) alignment (9 levels), skills (includes a skill type and a modifier for that skill e.g, stealth +3), and the challenge level (a number between 0 and 30), amongst others. The dataset we have available includes 763 monsters, and includes URLs to description webpages for half of them (which we can use to get images of those monsters)

The monster they find might be used as it is, or the stats might be kept (to ensure it is a balanced monster stats-wise) but its background and appearance might be changed to fit the story or encounter setting. When designing an encounter multiple monsters might be picked, and they can have a boss/minion relationship (e.g., dragons with kobolds) with some more challenging than others. The tasks for this module include discovering monsters by browsing or exploring (based on how clear is their

criteria) to identify a monster that fits the DM's need well and exploring the options. The targets are one monster or more, identified by name.

3.4. Relationships Tracker

In D&D character relationships are important information to track, specially for roleplaying and storytelling purposes. To maintain verisimilitude the DM should remember if a non-playing character (NPC) dislikes a player character (PC), for example, to narrate their actions accordingly. Our initial interviews showed that out of 7 or 9 interviewees reported tracking in some way how characters feel about another. After the interviews were conducted, it was noticed that the data of interest is usually checked by one on one interactions. Another insight gained is that this tracker could be useful not only for DMs, but for players also. Some interviewees mentioned they prefer asking for the players themselves to remember this type of information, which was not anticipated in our initial considerations of the project.

Five DMs reported classifying the relationships they keep track of as good/neutral/bad and 3 out of them quantify it in a scale of positive/neutral/negative using the Strixhaven relationship score (Annex D). With this in mind, we can view the relationships between characters as mutable tabular data of a score for a character to character relationship on a scale of [-2, -1, 0, +1, +2]. Even considering the DMs who do not quantify this data in numbers, we can transform their binning system of bad/neutral/good to the scale. Going forward in the project we may find it useful for the DMs to add some form of classification for the relationship as well, such as familiar, amorous, work, etc. The tasks to be executed in this module include locating the character of interest and identifying the relationship that character has with another. Possibly also summarizing/presenting how that character is seen by groups of people by analyzing the tendency of positive or negative relationships centered in the character of interest.

4. Solution

In this project we will develop a dashboard with tabs for each one of the four pillar modules of the project: dice, party and relationship trackers and monster shopper.

4.1. Encodings

Dice Tracker:

The dice tracker will use two modes of view side by side: overlapping histograms, mainly to visualize the distribution of rolls of a few dice, and a line chart to compare how multiple sets of dice perform overtime. We will make use of color channels to encode the individual dice in conjunction with a mean line to further encode the statistical values. Interactions are also added to facilitate visualization and will be further discussed in the results.

Since some dice will be used more often than others, normalizing the values could mislead a user to think a die is unlucky when it possibly just has not been rolled enough to be representative. Therefore, we will use the raw data and derive the mean. We are setting a limit of 5 dice sets for input to facilitate the visualization of the data.

Party Tracker:

The ability score portion of the party tracker is still in the process of design set-up definition and will possibly be a collection of bar graphs side by side, with each chart corresponding to a different attribute score, with each bar of the graph corresponding to the ability score modifier of a character or superimposed bar charts like the dice histogram. This would allow the DM to see the scores in relationship to the modifiers. This function can quickly show the peaks and valleys of the team's stats.

The health of the players would be recorded in a horizontal stacked bar chart (progress bar) that shows the current health of the players, as well as their total life amounts. Additionally, these charts will be stacked above each other for quick comparisons. The EXP chart would follow a similar design for similar reasons.

The "tabular data" section of the tracker with speed, armour class and passive perception will be listed on the dashboard and require no formal visualization. We hope to make the information accessible without cluttering the tracker and distracting from the main components.

Monster Shopper:

The Monster Shopper module is an interface for faceted exploration of monsters, starting from the aspects more important to DMs (story → stats, or stats→story) with pictures and descriptions for each on demand or in proximity. It will include a shopping cart panel where selected monsters are shown and a panel with bar charts and text indicating the current difficulty of the encounter (based on the combined challenge rating). Other panels will show the current player's levels (can be pulled from or synced with the Party Tracker) and the details of a single monster, which can also allow for that monster to be customized (edited) and added to the collection. The main idioms might include lists, widgets, and possibly simple bar charts if deemed necessary. As for interactions, the different panels are going to be coordinated in some fashion.

Relationships Tracker:

The character relationships data can be viewed either as horizontal bar charts centered in one character. The tracker is derived from a master table of n-characters by n-characters and their respective relationship score. We use the color channel to encode the quality of the relationship (blue for positive and red for negative). If time allows, we could also have the type of relationship (familiar, amorous, work) as tags to filter data. An interaction could be added to change the center of the relationships from one character to another by clicking on their icon.

4.2. Implementation

Technology: The website as a whole will be built using Vue.js which is a framework for building reactive websites, meaning any changes in data (model) will be reflected automatically in the visualizations (views). Then the trackers will be built as individual Vegalite components, while the monster shopper will use javascript and Vue.js. The data will be allocated into JSON and CSV files or stored in a MongoDB database

Data sources:

For some of our dashboard components, we are relying on live data. This live data will come from an existing commonly used tool for playing D&D (Roll20) to increase the dashboard's utility and the likelihood of its adoption by end users. We experimented with getting gameplay data out of Roll20 and found it feasible to do so through the scripting system that Roll20 supports along with a Chrome extension. We will aim to feed that data into our dashboard and also provide manual ways for data entry if Roll20 is not used. For the trackers we will manually input realistic sample data as a start using our experience playing the game.

For other dashboard components (namely the Monsters Shopper) the data will not be live and will be retrieved from the following resource:


https://data.world/datafam-con/june-2022-dungeons-and-dragons-collab-with-gamesnightviz/workspace/file?filename=dnd_monsters.csv

4.3. Preliminary Results (Scenarios of use and Sketches)

Party Tracker


The party is coming up to an ancient door that is protected by some foreboding stone golems. The DM quickly checks to see how the party is looking. Their health bars indicate that they are still pretty battered from their goblin ambush, with only one character having most of their HP still. Feeling pity for their friend's characters, who have been with them through this campaign for the last two years, the DM decides that maybe only one golem will attack the party, hopefully allowing the brave explorers to squeak by again. The DM also wants to know the best way to present the entrance of the tomb. Should they highlight the arcane sigils that line the door, so a high wisdom character can read them for a clue to enter? Or should they mention the faulty stonework, allowing anyone with high strength to brute force their way in? Quickly reviewing the party's stats, the DM notices that the new player to join the campaign has a character with high dexterity. The DM also notices that none of the characters have above average strength or intelligence. That decides it, the DM wants to give the new player a chance to shine and thus weaves in a series of ropes that lead through an upper hole, into their description of the door, perfect for a dexterous character to

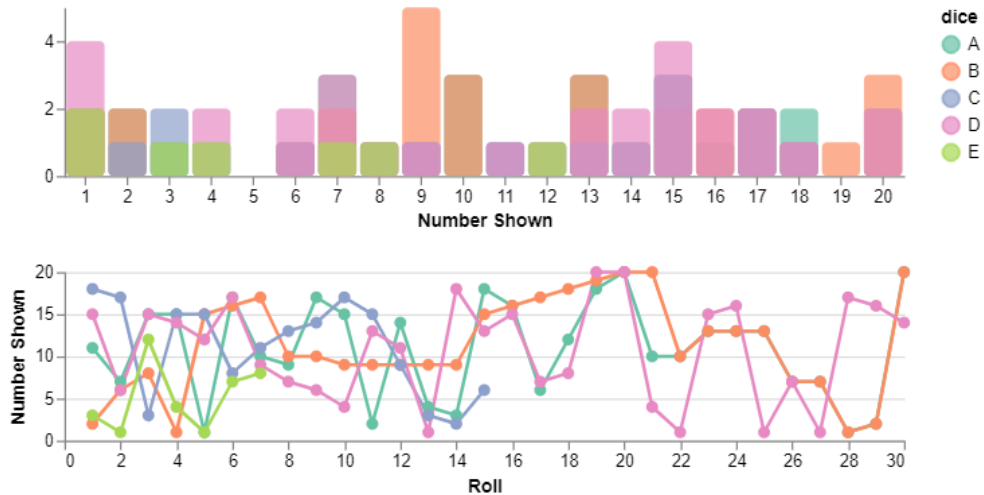
traverse. The party quickly deals with the golem and the new player is able to show-off their skills by climbing the ropes. All of these decisions were performed quickly and efficiently by the DM, without the party realizing that the event was tailored for them on the fly.



Dice Tracker

Like many D&D players, the DM is very superstitious about their dice. They have been rolling poorly recently and want to know if the dice are worth keeping. Before they throw their dice back into their bag they want to check the distribution of the rolls to see if it actually matches their perception or if a couple of bad rolls just stood out. Additionally, they know the campaign boss is coming up and they want to use their luckiest dice to give the party a run for their money. The DM pulls up the dice tracker and notices that the distribution of their “unlucky” d20 is actually pretty average and that they just remember the stand out rolls. They select a few other sets of dice to compare against this current set. The histograms overlay each other with an obvious winner. They see that their gold embossed dice roll 15 or higher more often than their other dice. Looks like the party is in for a tough fight.





Monster Shopper

The DM wants a monster that fits the cave-like environment where they are going to place the next encounter for their players. The players heard from townfolks about a dragon hoarding treasures that live in the nearby cave so the DM wants to find a dragon monster that is of the appropriate challenge rating given the player's experience. They enter the player's current levels in a dedicated panel first to be a reference against which they will judge the encounter difficulty. The DM first filters monsters by type using a filtering widget and picks "dragon" then looks at a table listing those dragons, filters any dragon with a challenge rating (cr) above 9, sorts the results and picks one of the top 3 that looks the best to them (using their images and description) using the Monsters Details panel. That dragon is added to a shopping cart, and a challenge rating for the encounter as a whole is updated once the monster is added to the cart. Seeing that the general challenge is low they might add a few lesser monsters like goblins following the same process but aiming for Humanoids as a monster type and a lower challenge rating until the encounter's challenge fits their desired level. The initial sketch of the shopper is inspired by [Dnd5e Monster Database](#) and [D&D Beyond's Encounter Builder](#).

Monster Shopper

MONSTER FILTERS

REMEMBER MY FILTERS

Environment (1) ✕ ▼

Source ▼

Alignment ▼

CHALLENGE RATING

1/2 ————— 19

0 1/2 5 10 15 20 25 30

SIZE

Tiny ————— Large

● ● ● ● ● ● ● ● ● ●

[CLEAR ALL FILTERS](#)

[SHOW ADVANCED FILTERS](#)

Monster type

Aberration

Beast

Celestial

Construct

Dragon

Elemental

Fey

Fiend

Giant

Humanoid

Monstrosity

Ooze

Plant

Undead

PC

Level ✕ ▼

Level ✕ ▼

Level ✕ ▼

Level ✕ ▼

NAME	TYPE	SIZE	CR	DIFFICULTY
Young Sapphire Dragon Fiend's Treasury: Dragons	Dragon	Large	9	●●●●
Young Solar Dragon Spelljammer: Adventures in Space	Dragon	Large	9	●●●●
Young Blue Dragon Basic Rules	Dragon	Large	9	●●●●
Young Amethyst Dragon Fiend's Treasury: Dragons	Dragon	Large	9	●●●●
Young Silver Dragon Basic Rules	Dragon	Large	9	●●●●
Young Moonstone Dragon Fiend's Treasury: Dragons	Dragon	Large	8	●●●●
Young Bronze Dragon Basic Rules	Dragon	Large	8	●●●●

Monster Details

YOUNG BRONZE DRAGON
Large Dragon, unaligned

Armor Class 14 (natural armor)
HP 140 (14 HD × 10)
Speed 40 ft., fly 50 ft.

STR	DEX	CON	INT	WIS	CHA
20	12	10	10	10	10

Saving Throws: DEX +4, CON +3, INT +3, WIS +3, CHA +3
Skills: Acrobatics +7, Perception +7, Stealth +3
Damage Immunities: lightning
Senses: blindsight 60 ft., darkvision 120 ft., hearing distance 120 ft.
Languages: Common, Draconic
Challenge: 8 (2,500 XP) Proficiency Bonus: +3

Amphibious. The dragon can breathe air and swim.

Multiattack. The dragon makes three attacks on its turn: one with its claw, one with its dragon attack, and one with its tail. One large, 40-ft (12.8m) piercing damage.
Claw: Melee Weapon Attack: +4 to hit, reach 5 ft., one target; 1d6 (3d6) slashing damage.
Dragon's Breath (Recharge 5): The dragon uses one of the following breath weapons.
Lightning Breath. The dragon exhales lightning in a 30-foot line. Each creature in that line must make a DC 13 Dexterity saving throw, taking 55 (16d6) lightning damage on a failed save, or half as much damage on a successful one.
Dragon's Breath. The dragon exhales dragonfire energy in a 30-foot cone. Each creature in that cone must succeed on a DC 13 Dexterity saving throw. On a failed save, the creature is pushed 40 feet away from the dragon.

ENCOUNTER SUMMARY

DIFFICULTY
Hard

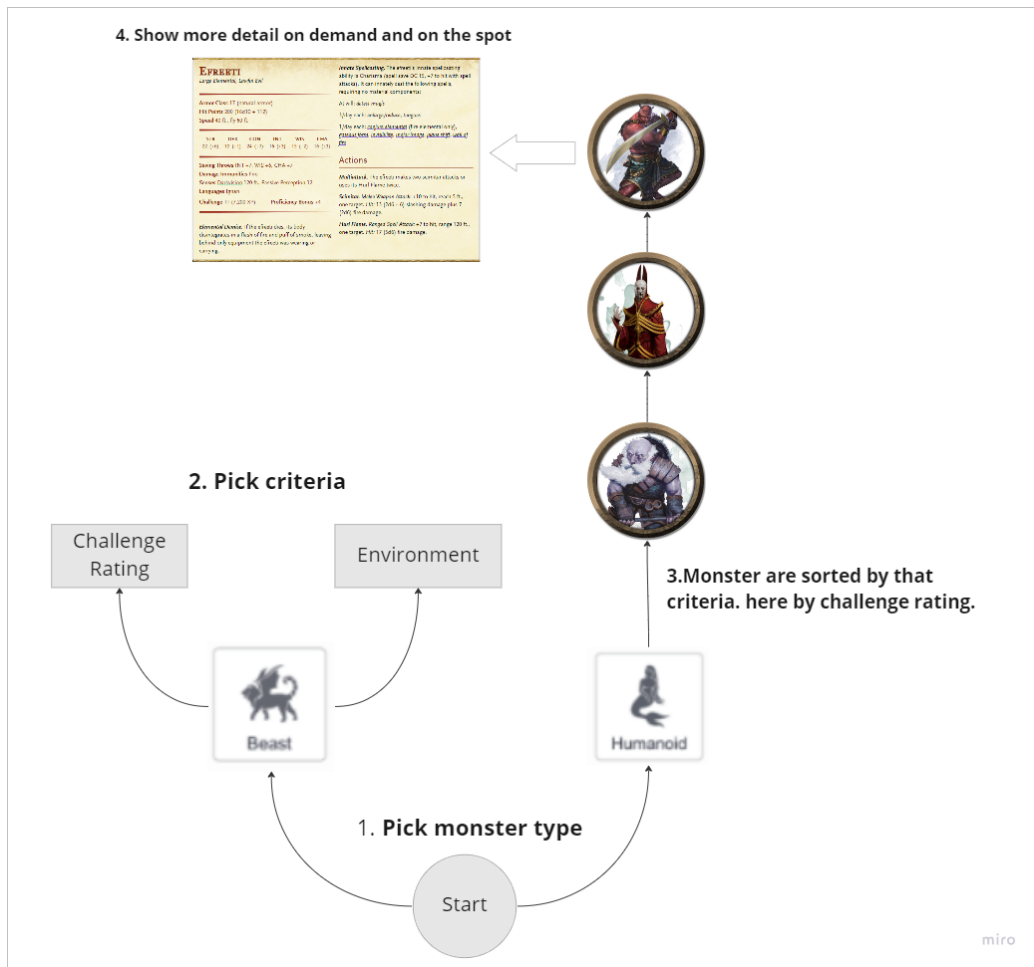
TOTAL XP
3900

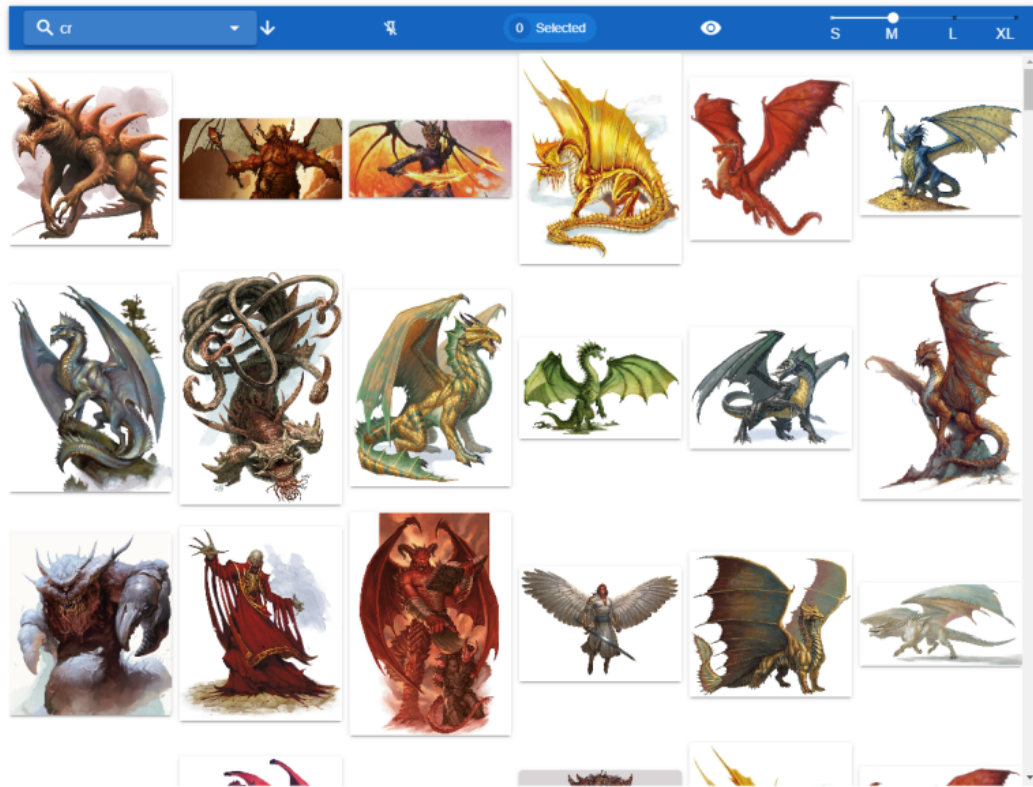
ADJUSTED XP
3900 XP (x1)

EASY: 1000 XP
MEDIUM: 2000 XP
HARD: 3000 XP
DEADLY: 4400 XP

DAILY BUDGET: 14000 XP

YOUNG BRONZE DRAG...
Large Dragon
CR: 8
XP: 3900






Relationship Tracker

The DM is having a reappearance of a NPC the party met in the last city they visited in a previous session. They narrate the encounter and the NPC is being friendly with the player. "Wait! Wasn't that NPC that hated my character because I stole bread from their bakery?" This may be intentional, the NPC may be actively faking their friendliness, or the DM could have just gotten confused. To avoid this confusion, the DM can log an editable master table with all the necessary characters and their relationships to the Relationships Tracker of the dashboard. Then, when the PCs enter a bakery and ask for the NPC, the DM just needs to search the name in the dashboard and see their relationship with the party members to guide their storytelling. If an interaction goes really bad one day, the DM can edit the table to update a decreased score.

Character	A	B	C	D	E	F	G
A		+2	0	-1	+1	-2	+1
B	+2		-2	+1	+1	-1	+2
C	0	-2		0	-1	+1	0
D	-1	+1	0		-1	+1	+1
E	+1	+1	-1	-1		0	-1
F	-2	-1	+1	+1	0		-2
G	+1	+2	0	+1	-1	-2	



5. Milestones

Previous steps

Domain understanding and abstraction (initial draft 21/Oct, ongoing afterwards):

- Helena: conducted informal interviews with DMs to learn about their current processes with the goal of identifying their needs
- Ahmed: Monsters Shopper
- Ryan: Party Tracker and Dice Tracker
- Helena: Relationship Tracker
- All: giving feedback to others

Design (3 weeks, by 11th Nov):

- All members

Writing (Throughout the term, final report by Dec 14th):

- All members

Next steps (started)

Definition of last design details (1 week, by Nov 18th):

- All members (1-2h)

Implementation (4 weeks, by Dec 9th):

- Ahmed - Monster shop module, relationship tracker module + getting gameplay data out of Roll20 into the dashboard (15-20h)
- Ryan - Dice tracker module, party tracker tracker module (15-20h)

Testing (1 week, by 13th):

- Helena (5-13h)

Writing final report (3 weeks by Dec 14th):

- Helena (10-12h)

Final Presentation preparation (1 week, by Dec 14th):

- All members (1-3h)

6. Discussion, Future Work, Conclusions

N/A

7. Bibliography

- Berbesi, L. (2019, November 24) DnD Exploratory Analysis using R. Towards Data Science. Retrieved October 19, 2022, from: <https://towardsdatascience.com/dnd-exploratory-analysis-using-r-d13757b09768>
- Esker. (2020, May 29). New Interactive DPR Visualization Tool. Morris Unofficial Tabletop RPG News. Retrieved October 19, 2022, from: <https://www.enworld.org/threads/new-interactive-dpr-visualization-tool.672422/>
- Hartelius, U, and Fröhlander, J, and Björk S. (2012). Tisch digital tools supporting board games. In Proceedings of the International Conference on the Foundations of Digital Games (FDG '12). Association for Computing Machinery, New York, NY, USA, 196–203. <https://doi.org/10.1145/2282338.2282376>
- Heller, E. (2019, February 21). D&D beginner's guide: how to get started with Dungeons & Dragons. Polygon. Retrieved October 14, 2022, from <https://www.polygon.com/deals/21294556/dnd-how-to-play-dungeons-dragons-5e-guide-spells-dice-character-sheets-dm>
- Michaud, J. (2015, November 2). The tangled cultural roots of Dungeons & Dragons. The New Yorker. Retrieved October 14, 2022, from <https://www.newyorker.com/books/page-turner/the-tangled-cultural-roots-of-dungeons-dragons>
- Plijaer BA, and O'Neill D, and Kompier E, and Wallner G, and Bernhaupt R. (2020). Truesight Battle Grid - Enhancing the Game Experience of Tabletop Role-Playing through Tangible Data Visualization. In Extended Abstracts of the 2020 Annual Symposium on Computer-Human Interaction in Play (CHI PLAY '20). Association for Computing Machinery, New York, NY, USA, 103–107. <https://doi.org/10.1145/3383668.3419931>
- Sedlmair, M, and Meyer, M and Munzner, T. (2012, Dec) Design Study Methodology: Reflections from the Trenches and the Stacks, in *IEEE Transactions on Visualization and Computer Graphics*, vol. 18, no. 12, pp. 2431-2440,, doi: <https://doi.org/10.1109/TVCG.2012.213>
- Soderlund, J. (2018, October 4) Dungeons and Data Science: What Can Data Tell Us About D&D? Displayr. Retrieved October 19, 2022, from: <https://www.displayr.com/dungeons-and-data-science>
- Wieland, R. (2021, May 19). 2020 Was The Best Year Ever For Dungeons & Dragons. Forbes. Retrieved October 14, 2022, from <https://www.forbes.com/sites/robwieland/2021/05/19/2020-was-the-best-year-ever-for-dungeons--dragons/?sh=6a2107b94f37>

Appendix

Appendix A:

Chokers

Chokers are ravenous, greedy hunters. They lurk in the darkness of caves and dungeons, snatching any prey they can get their grasping tentacles on. Easily bribed and hard to escape from, they can at best be controlled and at worst, be consumed by.

Alien Physique. A typical choker appears to be a mottled brown, semi-humanoid creature with long tendril-like limbs ending with four fingers. A choker's skull, spine, and rib cage are bony, but its limbs are tentacles with joints of cartilage. Their limbs have a peculiar movement and seem almost fluid. Their hands and feet are covered in a spiny pad that helps them grip almost anything with ease.

Basement Dwellers Chokers live in dark, solitary environments with plenty of places it can use to ambush potential prey. As such you will usually find chokers living in caves, dungeons, city slums and, in some instances, forests and swamps.

Ravenous Appetite. A choker's hunger can almost never be satisfied. It's build however, does not allow it to attack prey head on without facing difficulty. Because of this, chokers prefer to ambush individuals, using their unique abilities to hide in unusual places such as on ceilings and above doorframes. A choker will do almost anything for food, a smart adventurer can use their knowledge of undercommons to bribe the choker into leaving you alone or even helping them achieve their goals.

Tinkler Collectors. Chokers are of limited intelligence and have little society of their own. Despite this, they are often fascinated by the products of more advanced cultures, and have been known to collect small tokens from their prey. Chokers are socially advanced enough to identify themselves and others with names and occasionally form rudimentary friendships or alliances with less savvy humans.



VARIANT: VINE CHOKERS
Some chokers have adapted to survive and blend into forested environments. These choker's skins have adapted to resemble vines and are covered in forest growth. This growth has a toughness similar to bark and is generated by poisonous spore pods, providing an additional layer of protection. A vine choker has a challenge rating of 2 (450 XP) and has the following traits.

Condition Immunities. poisoned

Innate Spellcasting (1/Day). The vine choker can innately cast *barbed*, requiring no material components. It's innate spellcasting ability is Wisdom.

Spore pods. Poisonous spore pods on the vine choker's body explode when the vine choker takes damage. Creatures within 5 feet of the vine choker when it takes damage must make a DC 17 Constitution saving throw or take 4 (1d6) poison damage and become poisoned for 1 hour. The vine choker can trigger this ability as a bonus action 3 times a day.

CHOKER
Small aberration, chaotic evil

Armor Class 15
Hit Points 13 (3d6 + 3)
Speed 20ft

STR	DEX	CON	INT	WIS	CHA
16 (+3)	14 (+2)	13 (+1)	4 (-3)	13 (+1)	7 (-2)

Skills. Athletics +3, Stealth +8
Senses. Darkvision 60 ft., passive Perception 12
Languages. Undercommon
Challenge. 1 (200 XP)

Climbing. The choker can climb difficult surfaces, including ceilings, without suffering a speed penalty or needing to make an ability check.

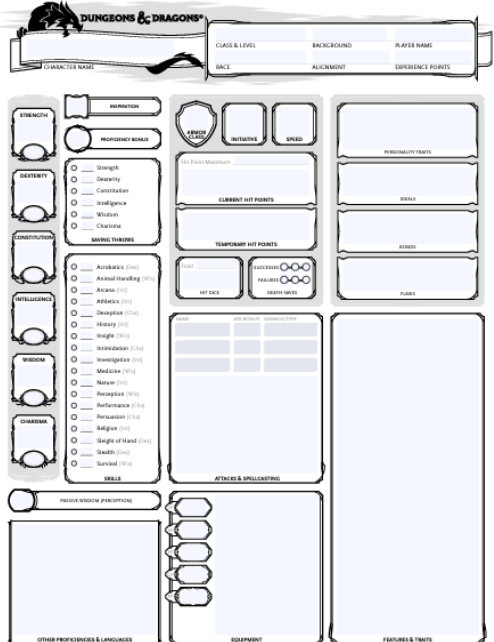
Quickness. Although not particularly dexterous, a choker is so communally quick, it can take an extra standard action or move action during its turn each round.

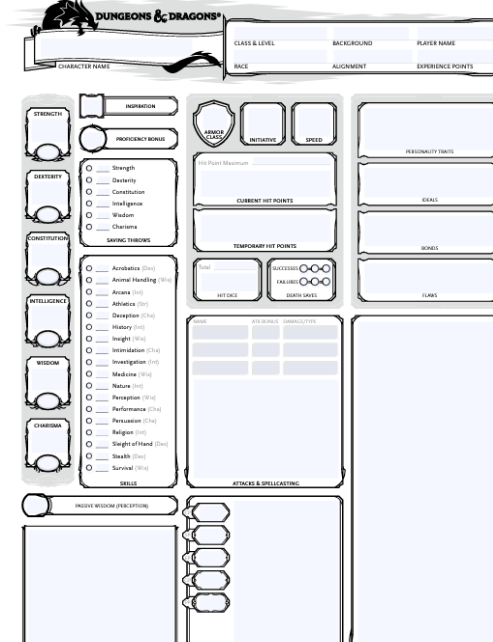
ACTIONS


Constrict. Melee Weapon Attack: +4 to hit, reach 10 ft., one target. Hit: 5 (1d6 + 3) bludgeoning damage and the target is grappled (escape DC 14). Until this grapple ends, the creature is restrained and the choker can't constrict another target.

Tentacle. Melee Weapon Attack: +4 to hit, reach 5 ft., one target. Hit: 6 (1d8 + 2) piercing damage.

Appendix B:







Appendix C:

ABILITY SCORES AND MODIFIERS			
Score	Modifier	Score	Modifier
1	-5	16-17	+3
2-3	-4	18-19	+4
4-5	-3	20-21	+5
6-7	-2	22-23	+6
8-9	-1	24-25	+7
10-11	+0	26-27	+8
12-13	+1	28-29	+9
14-15	+2	30	+10

Appendix D:

Strixhaven Relationship Points

Player characters will have the option to gain Relationship Points over the course of the adventure. These points reflect the relationship status a character has with another NPC.

All relationships start at 0, and either increase or decrease by 1 depending on the nature of the characters interactions. A friendly response will earn a positive relationship point, while a rivalrous response will earn a negative relationship point.

-A player can indicate that they would like to attempt to gain a positive Relationship point with an NPC, but this is at the discretion of the DM

-Conversely a player can take a negative relationship point with an NPC at any point at their own discretion, once per-character, per-day.

*Relationship Points and Status:

+/- 0 Neutral

+1 Friendly: While they are still getting to know you, the NPC is positively inclined towards you

-1 Antagonistic: Something about you just ticks this NPC off

+2 Friend: This NPC considers you a personal friend, and will go out of their way to assist you as best they can

-2 Rival: This NPC holds a great disdain for you, and will make your life harder any chance they get

+3 Beloved: You have a deep and meaningful connection with this NPC. While not inherently romantic, you are deeply bonded, so much so that you benefit from this NPC's spirit, even when they aren't physically near you (You gain this NPC's Beloved Boon)

+3 Enemy: This NPC hates you with every fiber of their being. They are extremely motivated to go out of their way to thwart your endeavors, and may even go so far as to physically harm you. Their hatred of you is so elemental it begins to affect your everyday life, even when this NPC isn't physically present (You suffer from this NPC's Enemy Bane)