

The Dungeon Master's Dashboard

Ahmed Abu Zuraiq, Helena de C. Alvarenga and Ryan Smith

Abstract—This paper introduces The Dungeon Master's Dashboard, a suite of tools for the interactive visualization of Dungeons&Dragons data. The suite is implemented with Vega-Lite, D3.js, OpenseaDragon, and customized AI tools. It consists of four modules proposed to facilitate a set of diverse tasks the narrator of the game (dungeon master) has to execute. We conducted pre-design interviews with domain experts to understand their needs. The Dice Tracker module allows the user to compare dice sets performance and the Party Tracker gives a clear and objective overview of the group's attributes for easy comparison between player characters. Additionally, for a swift role-play improvisation, the Relationship Tracker displays whether the relationship of a character with others is positive or negative. Furthermore, this work describes an innovative tool for browsing for creatures, the Monster Shopper. This final module balances data filtering and visualization enjoyment.

Index Terms—Role playing games, dungeons and dragons, dice, monsters, character, visualization, interactive, design study.

1 INTRODUCTION

Since its launch in 1974, the fantasy tabletop role-playing game Dungeons and Dragons (D&D) [11] has conquered popular culture and is referenced in a variety of media, including movies [6, 18, 19, 33], books [9], web and television series [2, 7, 26, 27, 41], podcasts [24, 44] and music [15, 21]. The gameplay [13] involves two sides, the player group (the party), and the Dungeon Master (DM). The DM is the narrator that creates a semi-structured open fantasy adventure (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.

Although the game has grown in popularity [42], the player interface remains rooted in its tabletop 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 errors by the DM. Our team believes that we can help increase the efficiency of a DM by creating data visualizations that allow them to access and assess the game's information in a more readable format.

Visualizing D&D data 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. This problem provides an interesting opportunity to explore the field and improve our data visualization skills.

-
- Ahmed Abu Zuraiq is with Simon Fraser University's School of Interactive Arts and Technology. E-mail: abuzreq@gmail.com.
 - Helena de C. Alvarenga is with the University of British Columbia's Bioinformatics Department. E-mail: h.alvarenga.bsb@gmail.com.
 - Ryan Smith is with the University of British Columbia's Computer Science Department. E-mail: rsmiith0914@gmail.com.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org.
Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxxx

2 DOMAIN BACKGROUND

During the campaign, it is important for the DM to track information about the party and the fantasy world. The DM must keep reference literature on hand, such as a monster manual for selecting and creating monsters for the party to battle, as well as each party member's character sheet, which tracks the individual character stats, 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. 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.

2.1 Rolling Dice

There are seven types of dice that players and DMs usually use, ranging in side number from four (d4) to twenty (d20). The d20 is used for most of the game roll checks, may it be trying to attack a dragon, stealthily pass by a guard or deceive a witch. The player rolls their d20 and then modifies it according to their stats and compares it to a roll by the DM to see if the action was successful. The needed roll depends on the difficulty of the action and other game mechanics that are not going to be discussed in this work. Typically, rolls closer to 20 are better and correspond to success and, symmetrically, rolls closer to 1 are worse and correspond to failure.

The extreme numbers, often called "natural" 1 or 20, are special. Some DMs consider those rolls for more intricate role-play. For instance, if a player rolls a 2 for a check to see if their character can find any secret doors in a room, the DM might narrate that they can not find it. Yet if the player rolls a 1 in that check, the DM might narrate that the character does not find a door, but while looking for it they trip on a rock and fall. A roll of 20 when attacking a monster in battle, for example, also allows the player to deal extra damage to the creature.

2.2 Characters

In D&D each player creates a character (PC) that 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 A for the full table). 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 will score a 21 for the roll with the 3 extra points. The DM or official rules will dictate which ability scores are used for which instance.

The PCs have levels at which they have access to different skills and special abilities. To level up, they need to gain *experience points* (XP), usually by completing combat encounters. The game also has other mechanics, namely the armour class, speed and passive perception. The *armour class* (AC) is a value that expresses how difficult it is to hit an opponent, while the speed is the distance in feet that a creature (character or monster) can move per turn and finally *passive perception* (PP) is a score that indicates how much a character can perceive about their surroundings without actively looking for something, such as signs of an ambush for example. The DM also needs to track how many *hit points* (HP), or health, the party members have left. When a party member passes out, they are at risk of dying, thus potentially losing their character forever.

The DM also creates extra *non-playing characters* (NPCs) for the players to interact with. To maintain verisimilitude the DM should remember if an NPC dislikes a PC which can alter an interaction in favour of or against the party. Therefore, character relationships are important information to track, especially for role-playing and storytelling purposes.

2.3 Encounters

An encounter is a sequence of events where the players engage in an activity, such as fighting monsters. Similar to player characters, monsters have ability scores and modifiers, speed, and AC. To balance how difficult a combat encounter is, the number of characters, their levels and the *challenge rating* (CR) of each monster is taken into consideration. *“A monster’s challenge rating tells you how difficult the monster will be to defeat. An appropriately equipped and well-rested party of four adventurers should be able to defeat a monster that has a challenge rating equal to their 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 battle, but not a deadly one.”* [35]. In addition to the CR, monsters have types (beast, celestial, fey, humanoid, undead and others) which are also taken into consideration when a DM is creating a combat encounter.

3 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 of these tools is to increase immersion or offload unwanted tasks [12], 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 (for instance, a data set of character sheets created by past players [23]) with the purpose of understanding something about their choices, such as weapon type [3, 32].

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. These can be online platforms that are used to run the games like Roll20 [38] or Fantasy Grounds [31], or maintain textual records of gameplay-related information (AdventureCodex [29]) or provide the necessary material to run a campaign like game rules and pre-made content (D&D Beyond [34]). Additionally, there are more specific tools that DMs can use to generate maps (World Anvil [10] and Inkarnate [20]), model in-game combat parameters like damage per round [8], or select the monsters appropriate for a given game session [22].

In this project we focus on the third aforementioned category, in particular, tools that utilize some form of data visualization in support of DMs as they (1) prepare for a campaign, or (2) as they are running a gaming session. For the first, the closest related work is the monster selection tools by Mais [22] and the Encounter Builder tool on DnD-Beyond [34] for monster selection. Though both tools rely heavily on selecting by data attributes and do not incorporate images of monsters

as part of the selection. For the second category of tools used during a game session, we do not find examples of tools explicitly tailored for visualizing D&D data that are available for public use. Instead, virtual RPG tools like Roll20 usually host most of the gameplay data we are focusing on, albeit not visualized. Finally, we use the design study methodology [30] as a guideline for our work.

4 PRE-DESIGN INTERVIEWS

A round of pre-design informal interviews with nine domain experts was conducted to learn about the tools and processes they employ when planning and running a D&D game. All of the interviewees were asked the same questions individually, and the interviews lasted an average of 30 minutes. The volunteers were experienced DMs between 20 and 30 years old, 8 males and one non-binary person, all currently pursuing or finished post secondary education. It should be noted that the participants we had access to mainly spoke Portuguese, as the nine interviewees were Brazilian. A transcript of the translated interviews is available in the Appendix section. Even with a limited sample of domain experts, it is noticeable that the style and approach that DMs take when designing and running the adventures differs greatly. The main themes arising from the interviews will be summarized next. In later sections, we present a suite of modules each pertaining to different visualization needs identified through interviews.

4.1 Monsters Selection for Encounters

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 the story.
5. Consistency with setting (for instance, no fire monsters in an arctic region, unless the surprise is intended).
6. What is enjoyable to the DM themselves.

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. There are several sources the DMs use to find monster information, such as the official D&D Monsters Manual [35], where monsters are sorted alphabetically by their name and an element of surprise is achieved by randomly browsing the book. Another popular source is DnD Beyond [34], which has a tool for building encounters that include some filtering widgets and a shopping cart of monsters. Many DMs also take inspiration from campaigns run by other DMs as well as their own knowledge about monsters from series, books, etc.

4.2 Party Stats

The DMs expressed that the most important factors to keep track of, during the game, are the stats of the party. Constantly having to ask players for their current AC can break the immersion. The DM uses this information for individual checks as well as comparing party members. For instance, if a PC is close to dying, the DM might want to change the monster’s target to a character with more health. For that, it is essential to know which characters have higher HP.

4.3 Party Relationships

The relationship between characters was confirmed as information DMs take note of as well. Seven of the interviewees reported tracking how characters feel about one another. DMs generally like to keep the world they create as alive and realistic as possible and maintaining character consistency is an essential part of it. Five DMs reported classifying the relationships they track as good/neutral/bad and 3 out of them quantify it on a scale of positive/neutral/negative using the Strixhaven relationship score [37].

Table 1. An example of a party's stats, inspired by real D&D data. Each character is represented in a line and has a set of attributes. Strength (Str), Dexterity (Dex), Constitution (Con), Intelligence (Int), Wisdom (Wis) and Charisma (Cha) each correspond to a different ability modifier. The HP current is the temporary hit points a character has and can be healed back to full capacity (HP max) with spells or potions. Each character also has their individual Armour Class (AC), Passive Perception (PP) and Experience (XP) values.

Character	Str	Dex	Con	Int	Wis	Cha	AC	HP Max	HP Current	PP	XP
Aquila	-2	+1	+1	+2	+1	+1	15	20	13	11	250
Dunk	+1	+1	-1	+1	+1	+2	13	19	3	13	250
Mednawne	+3	+1	+2	+1	-1	+1	18	24	1	9	250
Minerva	+1	+3	+1	+2	-1	+1	15	24	24	10	250
Royce	-1	-1	+1	+2	+1	+1	15	22	15	11	250

4.4 Dice Superstition

Our initial hypothesis that many people are superstitious about their dice and tend to favour certain ones they deem “lucky” was confirmed in the interviews. The interviewees that expressed being superstitious about their dice mentioned that they take into consideration both rolls of extreme numbers (1 and 20) and rolls that are above or below 10.

5 PARTY TRACKER MODULE

The party data consists of all the member's stats, which is discrete quantitative data. This data can be understood as a table of characters and the values for each of their six ability modifiers, current HP, maximum HP, AC, PP and XP. It is important to note that ultimately, the ability modifiers (refer to Sect. 2.2) are more important than the ability scores as it is the number that alters the interaction and therefore will be the value considered for the visualization.

The ability modifiers range from -5 to 10 while the other listed stats have a minimum value of 0. Because there is a multitude of factors that contribute to the stat values, it is hard to conclude an exact maximum value as they potentially have no upper limit. Realistically AC and PP will rarely go over 30 and most characters have XP and HP that are within 25 points from each other and rarely go over 100. Parties typically have a low cardinality with an average of 5 players, which provides an opportunity for fewer summation stats on the dashboard and more raw data translations. Table 1 shows an example of party data based on a real D&D party.

The overarching goal is to allow the DM to strategize their encounters by seeing what are the strengths and weaknesses of each individual and how they compare to the party as a whole. This requires the module to optimize look-up, filter and compare tasks.

5.1 Solution

The Party Tracker (Figure 1) is designed to facilitate the comparing and tracking of the party's stats for the DM. The horizontal bar charts, at the top of the module, are specific stats from the character sheets that we believed were relevant through personal experience and pre-design interviews. Some stats that you would find on the character sheet, like walking speed, have been omitted since they are rarely needed and would add unnecessary clutter to the module. The health and experience bars include a current and maximum stacked bar to show how much health the character is missing or how much experience is needed for them to level. The most important stats, passive perception and health, were placed in the upper left-hand section in accordance with western norms of reading pages from the top left. We do not believe that varying the size of the charts by importance would dramatically impact the comparison and search tasks and chose to keep them all the same size for aesthetic reasons. The horizontal design was selected to contrast the vertically grouped bar chart and prevent visual confusion when viewing the module. The DM should be able to quickly compare and track the party's information through these graphs.



Fig. 1. Example of a Party Tracker with a toy data set. (a) Shows a general overview of the module, while (b) illustrates the interactions allowed in this visualization. Hovering over a bar displays a description and value for it and clicking the legend at the top right corner allows for filtering characters. The data for this image was generated using random number generators to show the robustness of the visualization and how it can manage extreme parties.

The grouped bar chart displays the ability modifiers of each team member. By grouping the chart by ability modifier, instead of character, the DM can quickly view the strengths and weaknesses of the whole group. A diverging bar graph allows the DM to see if the characters are going to receive negative or positive modifiers on their roles, which traditionally correlates to the proficiency of that attribute. This idiom allows the DM to see an overview of the party's ability modifiers, compare the members, understand their strengths and weaknesses, and construct a scenario to increase or decrease the difficulty of the campaign.

Colour is used as a channel to code each character so the DM can easily distinguish which bar represents which player. Additionally, the DM can filter by character to focus on individuals and small groups, an interaction that is especially useful when the group breaks apart. Sliders are provided for the stress test data set and cover the stats and constitution attributes. These do not show all possible options because they are meant to be proof of concept for future work. The sliders allow the DM to directly manipulate the chart to track changes throughout the campaign, creating a flexible but structured module.

5.2 Implementation

Vega-Lite [28] is a high-level grammar used in generating interactive graphs through declarative JSON syntax. Vega-Lite API [1], a Javascript library, was used for the implementation of the visualization. Vega-Lite API supports creating visualizations through chained function calls in Javascript which are then translated to Vega-Lite JSON specifications. All interactive functionality was created with the



Fig. 2. Party Tracker visualization of the example party presented in Table 1.

existing Vega-Lite API toolset, including pan and zoom. While the Vega-Lite API proved to be a quick and efficient way to generate idioms, the relatively small online community meant that many solutions needed to be translated from Vega-Lite to Vega-Lite API which caused initial issues that were eventually surmounted.

The low cardinality of party members meant that there were no issues with run time, nor do we expect there to be any in parties with more members. Two JSON files were manually generated, one that included “real-world” data and another that was designed to stress test the idiom. The real-world data was taken from the ongoing campaign of one of the team members and is meant to depict the module in a realistic setting. The stress test data included the full range of possibilities that could occur in a game. This was an important distinction because of the freedom that D&D provides, every DM will run their party with different rules. Additionally, the data sets needed to be rearranged to allow the API to access the correct attributes in the most efficient manner, for coding purposes.

The code was then compiled in Observable [4] for public display and the HTML Observable sliders were generated for interactivity. Given the scope of the project, we only created sliders for one of the characters to act as a proof of concept for future work.

Link to Observable Page:

<https://observablehq.com/d/f6c2e85a64d1bcf9#pack>

5.3 Scenario of use

The party is coming up to an ancient door that separates them from some foreboding monsters. The DM checks their stats and sees that their health bars indicate they are low on health from their goblin ambush, especially Mednawne who decides to pre-emptively uses a potion to increase his constitution. The DM uses the slider to alter Mednawne’s stats accordingly. Half of the party decides to take a short rest while the other half, Aquila, Mednawne and Dunk, quickly scout ahead. The DM can now filter the interaction by the individuals who are participating. The DM wants to construct an event that requires the scouting team to work together so he scans and compares their attributes to see that no one individual has a high wisdom modifier. Using this realization, he decides to place a protection sigil on the door that would require a high wisdom roll. Individually, they could not break this spell and will have to work together. He checks the passive perception of the scouting party and sees that Dunk has a high enough passive perception to notice some etchings on the wall that can help them with the barrier. Using the Party Tracker module, the DM was able to compare and track the party’s stats in a quick and efficient manner that allowed the DM to create a custom-tailored encounter, in

real-time, without the players realizing it.

6 CHARACTER RELATIONSHIP MODULE

The relationships between characters can be seen as mutable divergent data scores with a character-to-character relationship scale of [-2, -1, 0, +1, +2]. Even considering the DMs who do not quantify this data in integers, their binning system of bad/neutral/good can be binned to fit the scale.

According to the interviews, relationship data is usually only pertinent to one on one interactions between the PCs and the NPCs. Therefore the tasks to be executed in this module include locating the character of interest and identifying the relationship they have with others.

6.1 Solution

Following the abstractions, we designed a simple diverging horizontal bar chart centred on one character, derived from a master table of n-characters by n-characters and their respective relationship score. The Relationship Tracker has a dropdown menu interaction to switch between characters (Figure 3). The colour channel was chosen to encode the quality of the relationship (blue for positive and orange for negative).

6.2 Implementation

The Relationship Tracker was created using the JavaScript Library Vega-Lite. The data was manually generated using parameters that reflected a realistic set of relationships within the game. JS code was used to manipulate the data, while Vega-Lite constructed the visualization.

We used Observable to host our visualization in a public format as well as compile the code and provide HTML buttons for minor interactivity. Even though a world can contain many NPCs, we do not expect it to impair the performance of the visualization as the cardinality is still relatively low in comparison to other data sets that can be used with Vega-Lite in Observable.

6.3 Scenario of use

At some point in the campaign. the party finds itself at a bakery that they have previously visited and the DM needs to narrate the interaction with the baker. The DM quickly checks the DM Dashboard, searching for the baker Conrad Brandywood. His relationship with the party is not good!

They remember that the party was in a tough spot a few sessions ago. They did not have any food and Minerva decided to steal bread from that bakery. Unfortunately, at that time her roll to see if the baker

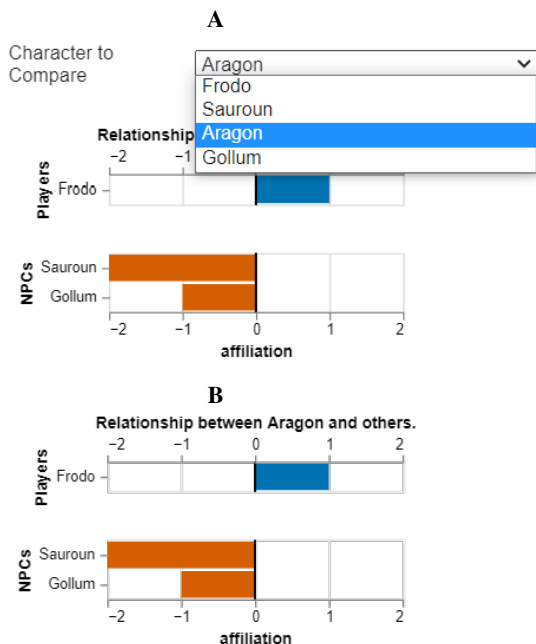


Fig. 3. After selecting a character from the dropdown menu (A), the Relationship Tracker shows the relationship score of the character with others (B). The tracker is organized by type of character: player or non-player (NPC).

noticed was very low. Although the baker was not fast enough to catch Minerva he remembers her face very clearly.

The DM describes the re-encounter interaction as unfriendly and narrates that Conrad calls the city guards. That maintains the NPC's actions consistent and the world feels more alive with players having consequences for their decisions.

7 MONSTERS SHOPPER MODULE

Based on how monsters are modelled 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 armour class, hit points and speed. All the quantitative 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, for example, stealth +3), and the challenge level (a number between 0 and 30), amongst others. The data set we have available houses 400 monsters and includes URLs to web pages containing descriptions and images for each monster. We scraped these web pages and obtained an image for each monster. The 400 monsters can all be found on the official Monster Manual book [35].

The tasks for this module involve locating and identifying monsters with the desired challenge rating and monster type as well as serendipitous visual exploration for interesting monsters. The module also supports browsing when the criteria for the desired monster is not clear *a priori*. Finally, producing annotation is also supported as means of keeping track of progress accrued during exploration.

7.1 Solution

DMs can have serendipitous encounters with monsters when browsing the Monster Manual, the large images of monsters on each page are often inspiring. However, monsters are sorted alphabetically and it is hard to quickly find monsters with known CR or type. On the other hand, the Encounter Builder on DnDBeyond is quite effective for filtering monsters but the monster's images take a very small portion of the screen.

We wanted to design the Monster Shopper to enable serendipitous and inspiring encounters while also being effective at filtering monsters down by data attributes. We explored multiple solutions around the concept of serendipity, inspired by the work by Thudt et al. [39]. The exploration showed a need for laying out monster images in a way that invited DMs to explore and ignite their curiosity. The solution had to be visually dominant and it needed to structure the images of the monster in a meaningful way without wasting space or incurring too much structure (e.g. sort by data attributes) so it becomes replaceable by data filters. With that in mind, we found an open-source solution that fits most of what we are looking for and customized it to fit the task of monster shopping. At the core of that solution is data preprocessing that made it possible to create the desired spatial and visually dominant layout. We describe preprocessing part in the following section, but next, we describe our solution.

Overview: The main view of the Monster Shopper is the image view, which is a rectangular grid where monsters are spatially organized by visual similarity and DMs can zoom and pan on it (Fig. 4). Two interactive visualizations, a histogram and a treemap, can be used to select data ranges which result in highlighting the images of the monsters falling in those ranges. When at least one selection is made on either of the visualizations, the monsters falling outside the selected range are dimmed with a dark scrim (50% transparency). This gives the highlighted (not dimmed) monsters more visual pop out, while allowing them to glance at the unhighlighted ones if they desired.

Visual Encodings and Interactions: The histogram depicts the count of monsters falling under different challenge rating (CR) values. A one-dimensional brush can be used to select the desired range on the histogram and the current extremes of the brush range are shown on the brush handles. The histogram domain ranges from zero to the maximum values and extra ticks are shown for the range between zero and one because the CR for many monsters falls within that range. The second visualization is a treemap depicting the monster type, each type is mapped to a cell and each cell is labelled by the type name and the count number. The area of cells encodes the number of monsters in the corresponding monster type category. Clicking on any of the cells highlights that cell as well as the monsters belonging to it. By pressing on a modifier (SHIFT key) more than one category can be selected. Any time a category or more is selected, overlaid bars are shown on the histogram view (Fig. 4 - B). This shows the percentage of these monster types to the total monsters within the same CR (part-to-whole comparison) as well as show the distribution of that monster type across different CR values (identify distribution). The colours for both the treemap cells and the background/foreground of the histogram are based on Tableau10's colour scheme which carefully combines colours that work well together. Finally, the selection on both the histogram and the treemap can be cancelled by clicking on a Reset button.

Details-on-Demand and Annotation: Upon hovering over any monster image, a red outline surrounds it to increase its visual saliency and a tooltip shows its name, type and CR. Furthermore, when a monster is clicked, a details panel is shown. The details panel loads an embedded web page (an iframe) from a source containing more details about each monster and a bigger image¹. On the details view, DMs can click on the View Full Record button which takes them to the aforementioned details web page. They can also click on the star icon which marks that monster. Marking is supported as a way for DMs to save their shopping result (i.e., derived data through annotation), and it is encoded with a glowing outline around the image of the marked monster. The Encounter Builder in DnDBeyond uses a separate panel to keep track of monsters in a 'shopping cart' style. We find this kind of marking to be more diegetic and immersive. We picked a blue-to-purple glow as we found it most salient regardless of whether the monster and its neighbours were highlighted.

Vis Idioms Justification: On the Encounter Builder, DMs can filter monsters using widgets but they cannot learn about the distribution of monsters along the data attribute controlled by that widget. In our solution, we offer a brushable histogram which acts as a scented

¹For example, see <https://www.aidedd.org/dnd/monstres.php?vo=aaarakocra>

widget [43]. This is essential because the majority of the monsters actually fall with a CR that is below 6 (i.e. somewhat right-skewed distribution) and filtering without knowing that can lead to empty queries. We followed a similar approach of making the monsters' distribution transparent by using a treemap to visualize the count of monsters under each type. A study comparing treemaps and bar graphs for hierarchical data [17] suggests using bar graphs when the data density is low (≈ 4096 items) at any given level in the hierarchy and treemaps otherwise. We diverge from the guidelines of the study in that our data is not hierarchical, we also found the large areas of the treemap cells to be easier to localize when wanting to click on them. Furthermore, the relatively balanced distribution of monsters across types did not create cells with extreme areas or aspect ratios. Finally, it can be confusing to have a bar chart that supports clicking on the bars, when a histogram lacking that interaction is placed beside it.

7.2 Implementation

To build the Monster Shopper, we customized an open-source system by the AMNH Science Visualization Group [25]. The system consists of two parts: data preprocessing written in Python, and a data visualization part written in Javascript. The Monster Shopper can be accessed at: <https://dnd-monsters-atlas.netlify.app/#>

Preprocessing: The preprocessing part takes a collection of images and a table of properties containing data attributes for each image. The images are first compressed (encoded in Machine Learning terms) as 4096 feature vectors. This is accomplished by taking the intermediate result of running an image through an image classification machine learning model called VGG16 using the Keras [5] framework. Following that, those feature vectors are further projected into a 250 features vector using the Principal Components Analysis (PCA) [14] technique. Following that, each vector of 250 dimensions is projected to 2 dimensions using T-SNE [40]. At this point, the result is can be visualized as points on a 2D scatterplot. To visualize the result on a grid, the 2-dimensional vectors are mapped to a 2D grid using a technique developed by Klingemann [16]. Users can choose the dimensions of this grid, but if it was smaller than the number of monsters then not all monsters in the dataset will appear. After assigning 2D vectors to grid cells, a large image matrix is created by combining the images of the monsters in each of the cells in the 2D grid, where each image can have a size as small as 128x128 pixels. This large image matrix is displayed later instead of the individual monster images which reduces the amount of media to load and results in a faster and more interactive experience. Following that, the coordinates of each sub-image in the large image matrix are stored as 'tiles' for later use. Finally, the tile information and the monsters' data attributes are linked by id and a JSON file is created which combines both.

Dataviz: The data visualization part takes as input the large image matrix and the data JSON produced in the preprocessing part. The visualization part is built using Javascript and uses D3.js for creating the histogram and the treemap, and OpenSeaDragon to display the large image matrix.

Changes: The work we did as part of developing the Monster Shopper includes:

1. We ran the preprocessing part using our own monsters dataset.
2. Added support for selecting more than one monster type on the treemap.
3. Highlighted the parts in the histogram (orange foreground) that represent the currently selected monster-type categories. Done using D3.js
4. Changed the details on demand view (when a monster is clicked) to include a full embedded webpage of the monster data fetched from and a star button to mark a monster.
5. Once a monster is liked, we added a feature of highlighting it with a glowing outline to distinguish it on the grid. Done using OpenSeaDragon.
6. Reduced the strength of dimming for the unselected monsters so their details are more discernible.

7. Changed the x-axis domain for the histogram to show the values between 0 and 1.
8. Changed the colour scheme to use Tableau10 for both the treemap and the histogram.

7.3 Scenario of use

A party includes 4 players, each at level 9. The players heard from town folks 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 (A monster with a CR of 9 would be a worthy opponent as described in Sect. 2.3). The DM first filters monsters by type using the treemap, picks "dragon" and the visualization highlights the dragon images in the overall view. The DM then filters any dragon with a challenge rating (CR) above 9 using the histogram. The DM then examines some or all of the highlighted dragons through a combination of considering their looks and clicking on them for more details. When they like a particular dragon they can mark it and continue browsing. Once they have felt satisfied with their options they can either pick one of the marked dragons and unmark the others or keep them for later reference. 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. When the combination of monsters is final, they can click on the View Record button for each of their choices which opens a new browser tab to a webpage with details about that monster. The browser tabs can be used for reference later on as they continue preparing for their campaign.

8 DICE TRACKER MODULE

Since the other types of dice are less frequently used in the game, only the d20 will be considered for the DM Dashboard. Dice rolls are discrete quantitative data sets of integers ranging from 1 to 20. Because the number of dice sets a DM has can vary and the interviewees expressed not having many, the Dice Module of the DM Dashboard was created for five dice sets.

The dice rolls are typically compared to a number (the difficulty class, or DC) to check if an action has succeeded or not. The basic DCs are 10-Easy, 15-Medium and 20-Hard. Therefore, roll quality is abstracted into bad rolls that are 10 and under, okay rolls that are 10 to 15, and good rolls that are 15 and above. Given the special nature, previously mentioned, of natural 1 and 20 rolls, these were also binned into roll quality.

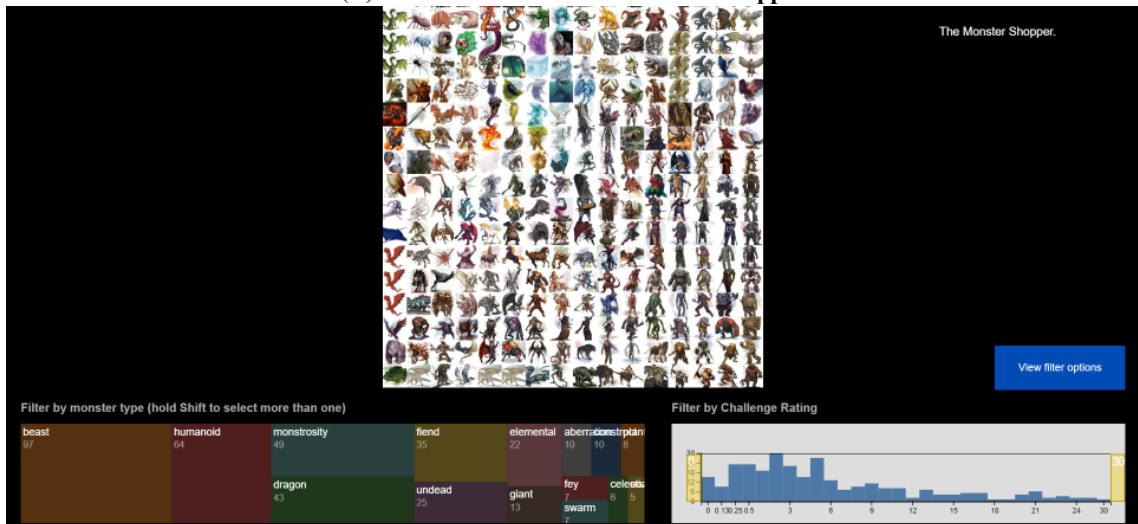
The DMs want to figure out which dice are their luckiest by analyzing the distribution of their current dice's rolls (for instance, did they roll more 20s on this dice or more 1s?) and comparing each d20's distribution to the distribution of their other sets to see if there are noticeable differences. It should be noted that statistically the rolls should be evenly distributed as the roll count increases, but real-world factors like poor dice quality can influence these rates. Additionally, the interviews show that the DMs 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.

8.1 Solution

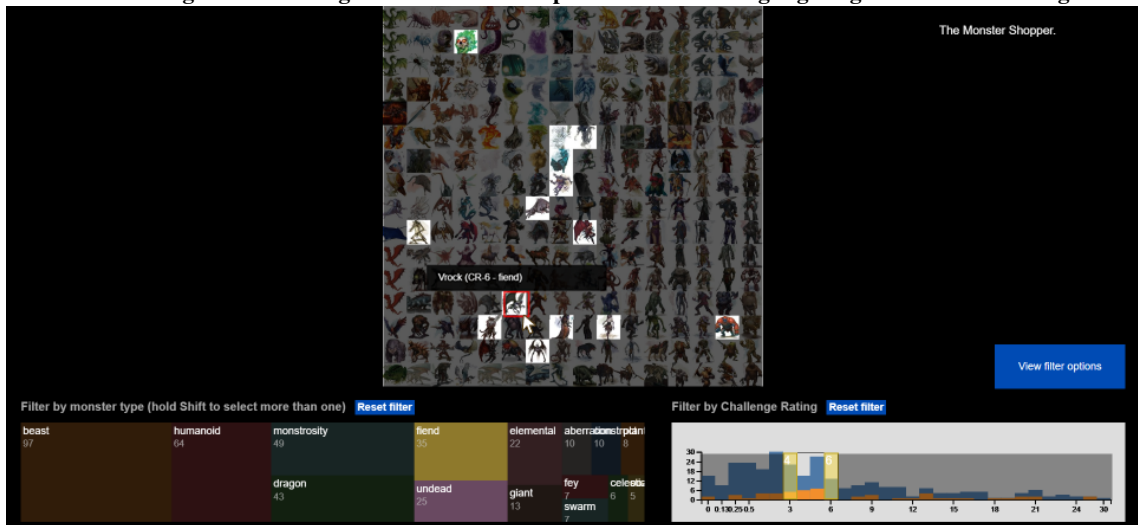
The Dice Tracker (Figure 5) was designed with three modes of view. The scatter plot timeline at the top allows the DMs to navigate through their rolls as well as observe the data in a less manipulated form. As they scroll and zoom through their timeline the other charts will filter accordingly. This overview visualization fits the task of exploration and discovery as they can see patterns for when they swapped dice as well as the various stages of the game.

Faceted histograms are used in the bottom left idiom to show the distribution of rolls for each dice set. The colours are coordinated with the same sets from the timeline to show a connection between the two idioms. The DM can quickly see the mean of each set to compare which dice have the highest average luck. The histograms are stacked above each other for the easy vertical comparison of each number on

(A): the full interface of the Monster Shopper.



(B): DMs can select data ranges on the histogram and the treemap which results in highlighting the monsters falling within that range.



(C): More details about a monster can be shown on demand. Monsters can also be starred which leaves a mark on the image grid.



Fig. 4. The Monster Shopper has a grid overview of the monsters at the top, spatially organized by visual similarity using an artificial intelligence algorithm. The general view (A) can be filtered by monster type and challenge rating (CR), which highlights the resulting monsters (B). Hovering the mouse icon over a monster will show its name, type, and CR. Clicking on a monster opens a details panel at the left corner and clicking on the star marks the monster and highlights it with a glowing outline (C). DMs can zoom and pan on the images grid.

the D20 distribution. Since, statistically, the distributions should even out with more rolls, the axis for the histogram is not fixed and will adjust according to the largest number of rolls within the active dice sets. Additionally, a count of the total number of rolls is displayed above each histogram. The histograms of the dice allow the DM to gain an overall understanding of each of the rolls and allow support them in choosing a dice set.

Stacked horizontal bar charts visualize the percentage of roll quality in the bottom right idiom. These are aligned horizontally with their dice set for a natural connection between the two idioms. The sequential colour scheme helps visualize the difference between the worse and better rolls. The roll quality was bucketed based on team decisions accompanied by research from the pre-design interviews. The stacked bar chart allows DMs to easily see how their dice rolls are being broken down, they can compare their bad, okay, good and natural 1 and 20 rolls within the same dice. Since the bar charts are stacked on top of themselves they can also compare the sets against each other. Understanding how the rolls compare between different sets as well as within the rolls of itself can help DMs decide on the luckiest dice for the campaign.

8.2 Implementation

A JSON data file was manually generated to contain all the information for the Dice Tracker. A number generator was used to mimic the random nature of dice rolls, while their chronological order was input manually to ensure there were realistic switching patterns between the sets. All data manipulation, visualization and interaction were coded in the Vega-Lite API.

The code was then compiled in Observable for public display. Observable's tooltips were also used to create sliders for module interaction. The implementation of the dice module faced similar issues as the Party Tracker where a small community of users was not always efficient in answering questions. Given the relatively low cardinality of the dice sets, issues with performance in a realistic setting are not expected. Link to Observable Page: <https://observablehq.com/d/08c9f9de6cb8579c>

8.3 Scenario of use

The DM is feeling unlucky with their rolls and it seems like they're rolling 1s too often during battles. They check the Dice Tracker (Figure 5 A) and the "Stone" dice is indeed on an unlucky streak for the last few rolls. Because a new encounter approaches they check their last battle (around roll 40-60) (Figure 5 B). The "Owl" dice set seems to be lucky on battles, so they decide to use it on the next one to gain some confidence.

9 DISCUSSION

Although we are happy with our final modules, we have reflected on the strengths, limitations and lessons learnt from the project.

9.1 Strengths

We presented a suite of modules that collectively cover different parts of a D&D campaign, the combination of which is found anywhere to the best of our knowledge. The Monster Shopper offers a balance of visually-driven and data-driven exploration of monsters that is not present in current DM support tools. It is also fast despite dealing primarily with images. The Party Tracker summarizes the most important stats for the DM, which are usually found scattered on the players' sheets or in the DM notes. The Party Tracker also uses simple and effective idioms to achieve that which makes it suitable for a wide audience.

The Dice Tracker supports multiple levels of analysis, leaving the choice of level to the DMs to choose from. First, percentage bars give a general impression about the dice quality, then the roll distribution shows a second layer that confirms the first and adds to it. Finally, the raw roll data can be used to filter down or confirm all the above. The Relationship Tracker supports the simple task of identifying relationships between a character of interest and others. It achieves that with a simple idiom as well.

9.2 Limitations

Vega-Lite proved to be a powerful language for iterating through data and quickly experimenting with different idioms but caused us to eventually hit a creative wall. Even small nuances, like having items appear only when selected, had to be abandoned due to the limitations we faced. Using a more flexible programming library, like D3, would have helped alleviate these issues and improved the complexity of the idioms. Observable also proved to have limitations with its lack of direct impute and manipulation capabilities. Inputting new dice rolls and nuanced tools to manipulate the stats would be welcomed upgrades to the modules. The Monster Shopper allows you to only filter by two attributes. A DM may want to select a monster based on alternative constraints. Additionally, the layout algorithm needs more fine-tuning as it involves multiple algorithms in sequence, each with different hyperparameters. Finally, having multiple Modules forced us to spread our time between them in every phase of the project. Refocusing our effort on fewer modules could have helped polish and improve their various aspects as well as increased their complexity.

9.3 Lessons Learned

We learned the importance of understanding scope and choosing the correct tools. While Vega-Lite and Observable seemed to be strong choices at the beginning, the ultimate limitations of these two kept us from adding specific aspects to the modules. Additionally, we had originally believed that we may have under scoped the project but found that learning a new library (Ryan with Vega-Lite) and creating a custom Monster Shopper with a visualized component was more time-consuming than originally believed, as evidenced by the project milestones (Table 2). We believe that these areas will improve with natural experience as we learn to gauge the size of a project and understand the nuances of different languages/tools.

10 FUTURE WORK

We would like to build a dashboard by combining all the modules into one website for easy access. This would also directly connect each of the modules so you can reference relationships, dice and individual players across visualizations. We would like to integrate this dashboard with online RPG tools such as Roll20 [38] to obtain live gameplay data. Furthermore, we would like to scale up all modules to support more customization in order to cater to the different styles and needs of DMs. Finally, we would like to conduct usability and utility testing with domain experts to validate our design decisions and implementations. The study can be conducted on individual modules or the combined dashboard. We would also like to explore the following for each of the modules.

The Party Tracker presents an easy to read visualization of the listed stats but can have additional tools and functionality built into the module. Having tabular data, like a character's equipment or religious alignment, would be a welcomed addition for streamlining the organizational process of DMS. Additional customization features, such as adding and subtracting characters and altering the min and max of stats would be another step forward. This would provide flexibility for customization of the module and keep it flexible for DMs while maintaining the necessary organization which makes it useful.

For the Relationship Tracker, we wish to augment our data with classification for the relationships, such as familiar, amorous, work, etc. This allows DMs to choose the kind of relationships they wish to focus on. We would also like to include image icons representing characters in our current design, which we believe will make it more relatable and enjoyable. An interaction could be added where clicking on a character's icon makes it the center of the comparisons.

We would like to augment the timelines of the roll tracker to mark important events that took place during the campaign, for example, when battles took place. This gives DMs additional context, allowing them to narrow down their analysis to those critical parts of the campaign.

At the moment, the Monster Shopper allows DMs to filter by only two data attributes, more attributes should be supported. The layout algorithm involves multiple algorithms in sequence each with different

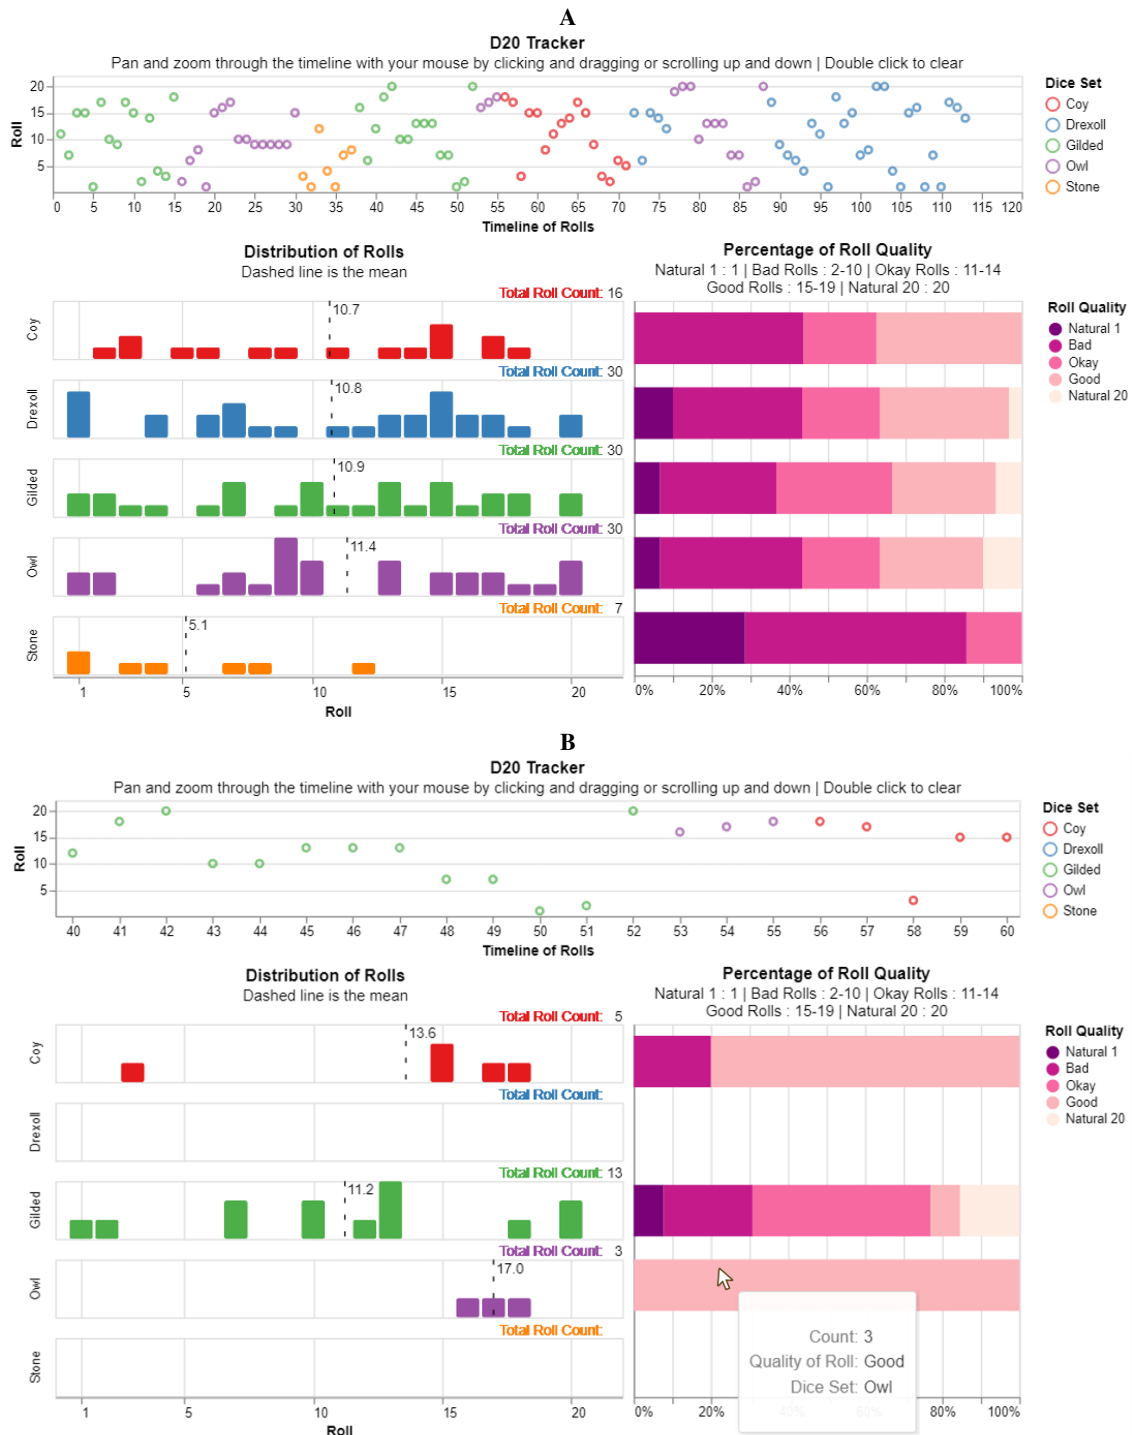


Fig. 5. The Dice Tracker shows three different charts for visualization of dice rolls (A). Zooming and dragging the timeline at the top filters the other subplots. (B) Hovering the mouse icon over a bar displays a box with the description of the roll.

hyper-parameters and so more work might be needed in tuning each step. Finally, we would like to add an encounter difficulty estimator as seen in the Encounter Builder on DnD Beyond [34]. The estimator informs DMs of the current difficulty of an encounter based on the challenge rating of the monsters marked/chosen so far.

11 CONCLUSIONS

The main objective of this work was attained, we developed a suite of four visualization tools for a variety of Dungeons&Dragons data. We conducted pre-design interviews with domain experts to understand

their needs, and post-design user tests will allow us to further investigate the utility of the Dungeon Master's Dashboard. Overall, our tool provides the DMs with the ability to choose monsters, compare party data and dice sets, and look up characters' relationships. We are excited to see the future of role playing games via visualization tools and hope to continue our work with the DM dashboard.

ACKNOWLEDGMENTS

The authors wish to thank the volunteers for the interviews and Prof. Tamara Munzner for the support and guidance.

Table 2. Expected and actual completion time of tasks throughout the course. Ahmed Abu Zuraiq (AZ) Helena de C. Alvarenga (HA) and Ryan Smith (RS) each completed 91h, 85h, and 91h of work, respectively.

Task	Estimated hours	Real hours	Deadline	Description
Pitch (x3)	AZ: 2 HA: 2 RS: 2	AZ: 2 HA: 2 RS: 2	Sep. 28	Create slides and presentation.
Pre-design interviews	HA: 3	HA: 10	Oct. 21	Conduct, transcript and translate informal interviews with domain experts.
Abstractions	AZ: 2 HA: 2 RS: 2	AZ: 5 HA: 5 RS: 5	Oct. 21	Data and task abstractions.
Proposal	AZ: 6 HA: 6 RS: 6	AZ: 6 HA: 6 RS: 6	Oct. 21	Discuss project ideas, brainstorm designs, create design drafts and write proposal.
Proposal update and review	AZ: 4 HA: 4 RS: 4	AZ: 5 HA: 10 RS: 5	Nov. 15	Describe encodings and preliminary results.
Final design	AZ: 2 HA: 2 RS: 2	AZ: 2 HA: 2 RS: 2	Nov. 30	Update the design according to feedback and discussions.
Implementation	95	110	Dec. 13	Completed suite of tools.
- Dice Tracker	RS: 25	RS: 30	Dec. 13	Create visualization of dice data, add interactions.
- Party Tracker	RS: 25	RS: 25	Dec. 13	Create visualization of party data, add interactions and sliders.
- Relationship Tracker	AZ: 15	AZ: 10	Dec. 13	Create visualization of relationship data, add dropdown menu.
- Monster Shopper	AZ: 30	AZ: 40	Dec. 13	Scrape and classify images, PCA, map to grid, add interactions and locking feature.
Final presentation	AZ: 2 HA: 2 RS: 2	AZ: 6 HA: 10 RS: 6	Dec. 13	Prepare slides, write script, record and edit video.
Report	AZ: 5 HA: 15 RS: 5	AZ: 15 HA: 40 RS: 10	Dec. 16	Finalize the report paper, prepare images of the visualization, review final writing.
Total hours	177	272		

REFERENCES

- [1] Vega-lite api. <https://vega.github.io/vega-lite-api/> Accessed: 13-12-2022.
- [2] D. 20. *The Unsleeping City*. Dropout, 2019.
- [3] L. Berbesi. Dnd exploratory analysis using r, Nov 2019. <https://towardsdatascience.com/dnd-exploratory-analysis-using-r-d13757b09768> Accessed: 19-10-2022.
- [4] M. Bostock. Observable. <https://observablehq.com/> Accessed: 13-12-2022.
- [5] F. Chollet et al. Keras, 2015. <https://github.com/fchollet/keras> Accessed: 16-12-2022.
- [6] J. F. Daley and J. Goldstein. *Dungeons & Dragons: Honor Among Thieves*. Paramount Pictures, 2023. Producers: Davis, Stephen and Latham, Jeremy.
- [7] M. Duffer and R. Duffer. Chapter one: The vanishing of will byers. In *Stranger Things*. Netflix, Jul 2016. Season 1 - Episode 1.
- [8] Esker. Dnd exploratory analysis using r, May 2020. <https://www.enworld.org/threads/new-interactive-dpr-visualization-tool.672422/> Accessed: 19-10-2022.
- [9] D. M. Ewalt. *Of dice and men: The story of dungeons & dragons and the people who play it*. Scribner, 2014.
- [10] J. Forbes and D. Havlidis. World anvil: The ultimate worldbuilding toolset & rpg campaign manager. <https://www.worldanvil.com/> Accessed: 13-12-2022.
- [11] G. Gyax and D. Arneson. *Dungeons and Dragons. Tactical Study Rules*, 1974.
- [12] U. Hartelius, J. Fröhlander, and S. Björk. Tisch digital tools supporting board games. In *In Proc. International Conference Foundations of Digital Games*, pp. 196–203, 2012.
- [13] E. Heller. A beginner’s guide to playing dungeons and dragons, Feb 2019. <https://www.polygon.com/deals/21294556/dnd-how-to-play-dungeons-dragons-5e-guide-spells-dice-character-sheets-dm> Accessed: 30-11-2022.
- [14] I. Jolliffe. *Principal Components in Regression Analysis*, pp. 129–155. Springer New York, New York, NY, 1986. doi: 10.1007/978-1-4757-1904-8_8
- [15] M. P. Jumpsteady. Dungeon master. In *Master of the Flying Guillotine*. Psychopathic Records, Feb 2005. Track 14.
- [16] M. Klingemann. Rasterfairy, 2015. <https://github.com/Quasimondo/RasterFairy> Accessed: 16-12-2022.
- [17] N. Kong, J. Heer, and M. Agrawala. Perceptual guidelines for creating rectangular treemaps. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):990–998, 2010. doi: 10.1109/TVCG.2010.186
- [18] G. Lively. *Dungeons & Dragons: Wrath of the Dragon God*. Warner Home Video, 2005. Producers: Richards, Steve and Clark-Hall, Steve and Esenwein, Wolfgang.
- [19] G. Lively. *Dungeons & Dragon 3: The Book of Vile Darkness*. IM Global, 2012. Producers: Richards, Steve and Rudnick, Brian.
- [20] I. E. LLC. Ikarnate: Create fantasy maps online. <https://inkarnate.com/> Accessed: 13-12-2022.
- [21] S. Lynch. D & d. In *Superhero*. What Are Records?, Jan 2003. Track 14.
- [22] V. Mais. Dnd 5e monster database & analysis 2.0, 2020. <https://public.tableau.com/app/profile/valerie1556/viz/Dnd5eMonsterDatabase-2.0/Explore> Accessed: 19-10-2022.
- [23] B. O. Mancarci. Dnd exploratory analysis using r, 2019. <https://github.com/oganm/dnddata> Accessed: 19-10-2022.
- [24] G. A. McElroy, J. T. McElroy, T. P. McElroy, and C. E. McElroy. *The Adventure Zone*. Maximum Fun, 2014.
- [25] A. M. of Natural History Science Visualization Group. Image collection: Visualizing amnh image collection with machine learning. <https://github.com/amnh-sciviz> Accessed: 13-12-2022.
- [26] C. R. Productions. *Critical Role*. Youtube and Twitch, 2015.
- [27] C. R. Productions. *The Legend of Vox Machina*. Amazon Prime Video, 2022.
- [28] A. Satyanarayan, D. Moritz, K. Wongsuphasawat, and J. Heer. Vega-lite: A grammar of interactive graphics. *IEEE transactions on visualization and computer graphics*, 23(1):341–350, 2016.
- [29] B. Schrader and N. Arellano. Adventurer’s codex: A d&d 5e online tool. <https://adventurerscodex.com/> Accessed: 13-12-2022.
- [30] M. Sedlmair, M. Meyer, and T. Munzner. Design study methodology: Reflections from the trenches and the stacks. *IEEE transactions on visualization and computer graphics*, 18(12):2431–2440, 2012.
- [31] L. Smiteworks USA. Fantasy grounds: a virtual tabletop application. <https://www.fantasygrounds.com/home/home.php> Accessed: 13-12-2022.
- [32] J. Soderlund. Dungeons and data science: What can data tell us about d&d?, Oct 2018. <https://www.displayr.com/dungeons-and-data-science/> Accessed: 19-10-2022.
- [33] C. Solomon. *Dungeons & Dragons*. New Line Cinema, 2000. Producers: Hammel, Thomas M. and Jam, Kia and Solomon, Courtney.
- [34] W. R. Team. D&d beyond: An official digital toolset for dungeons & dragons (d&d) fifth edition (5e). <https://www.dndbeyond.com/> Accessed: 13-12-2022.
- [35] W. R. Team. *Monster Manual*. Wizards of the Coast, 5 ed., 2014.
- [36] W. R. Team. *Player’s Handbook*. Wizards of the Coast, 5 ed., 2014.
- [37] W. R. Team. *Strixhaven: A curriculum of chaos*. Wizards of the Coast, 2021. p.48 - Relationships.
- [38] L. The Orr Group. Roll20: Online virtual tabletop for pen and paper rpgs. <https://roll20.net/> Accessed: 13-12-2022.
- [39] A. Thudt, U. Hinrichs, and S. Carpendale. The bohemian bookshelf: Supporting serendipitous book discoveries through information visualization. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI ’12, p. 1461–1470. Association for Computing Machinery, New York, NY, USA, 2012. doi: 10.1145/2207676.2208607
- [40] L. Van der Maaten and G. Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- [41] J. Whedon. Chosen. In *Buffy the Vampire Slayer*. United Paramount Network, May 2003. Season 7 - Episode 22.

- [42] R. Wieland. 2020 was the best year ever for dungeons & dragons, May 2021. <https://www.forbes.com/sites/robwieland/2021/05/19/2020-was-the-best-year-ever-for-dungeons--dragons/?sh=4d90efad4f37> Accessed: 30-11-2022.
- [43] W. Willett, J. Heer, and M. Agrawala. Scented widgets: Improving navigation cues with embedded visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1129–1136, nov 2007. doi: 10.1109/TVCG.2007.70589
- [44] F. Wong. *Dungeons & Daddies*. Acast, 2019.

APPENDIX

A ABILITY SCORES AND MODIFIERS TABLE

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

Fig. 6. Table of creature's ability scores and corresponding modifier values. Retrieved under fair use from the D&D Player's Handbook [36].

B INTERVIEWS

B.1 Questions:

Intro We want to know more about what DMs use during and before their campaigns. We are focusing on DnD specifically.

General

Do you do RPG on paper or do you use a web tool? What tools? (END) Can you share your notes?

Monsters

How do you choose a monster for an encounter? What factors do you take into consideration when choosing?

Character relationship

Do you keep track of the relationship between characters? What aspects of the relationship do you care about? How do you keep track of them? Do you quantify it in any way?

Party

How do you keep track of information about individual party members? Which aspects do you want to know about or refer to the most? Which sections of the character sheet do you use most frequently, as a DM? Do you track anything else not on the sheet? Now I want to ask you about how do you use this information When planning the next challenge (enemy encounter, puzzle, etc) for you players, do you consider their character information individually or as a group (e.g. strength, weaknesses, equipment)? Do you find yourself comparing characters? How do you do any of the above?

Dice Which numbers are the most important for you to know if a set of dice is lucky? Is it just 1 and 20, or do you take other numbers into consideration? When do you decide to not use a set of dice again? Is it low rolls, is it rolls of 1 etc.

B.2 Interviews

Interview 1 DMed only a couple of times on paper, did not know the tools well. After the pandemic they would DM online using roll20 or rRPG if they got the chance.

They consider ambiance, player's level and tastes. Uses the d&d books, internet in general, youtube campaigns with the same ambiance and the opinion of more experienced DMs as inspiration.

Yes, they use paper annotations and have a table to use during the game; if it is very important information they write it immediately. It is important to know that to make the game and role play more immersive

and interesting for both the players and the DM, the relationship information can be important in the future and can influence how an NPC is going to act. They note things like parent/children or trainer/trainee relations and quantify it in terms of "loves x neutral x hates".

For items and skills they leave it up to the player to keep track. When an information is very important they write it separately in private using paper, notes, tables and organize by character to not get lost. "I don't know". Skills, features, saving throws. Yes, for important stuff the player said or things that happened they ask the players to write it on the sheet and they write it down too to guarantee. They lean more towards a group focused decision, but go to an individual one if it is very important to the story. The game is played in a group so they want to highlight the group. If one person/character is having too much of the spotlight they try to find something to highlight the others. Yes, they compare strengths to balance fights.

They are not dice superstitious.

Interview 2 They use dice and maps online because they don't have physical resources such as miniatures, but use paper character sheets whenever it is possible. With the pandemic most games became online, but they run in person games when possible. Tools: roll20, rRPG firecast (not that good) and some websites to calculate encounters, recall spells and item prices.

Chooses according to the story, tries to make it more consistent with the scenario (no polar bears in the desert) and challenge level. They go by general knowledge, monsters they have seen in series, books, etc. Thinks of the monster first and then searches if there is a status for it in d&d.

Yes, at least a small note if they do not have time to detail it. They keep track specially for when an NPC reappears they remember that and keep verisimilitude. Example: "X had a good encounter with Y" or even writes more details to keep focus of the NPC's thoughts. The levels of the relationships are the most important: good or bad and how good and how bad. They put it as a scale, in the first encounter the characters are at 0 and after that on following interactions can increase positively or negatively. They do not classify the relationship, as a rivalry can be friendly or bad and it makes no difference whether the positivity comes from a friendship or a debt/promise.

Usually uses the DM notes in the character sheet in roll20 or paper. They keep track of the relationship with other characters, subjective aspects and impressions on actions. It can be the impression on a city level, also the reaction may change by culture. Annotates the story of the character, directions such as what the PC has done so far and where that is leading them; does not care much about technical aspects. Focuses on ideals, flaws, a little bit on the alignment and mainly the background. No. As a group because it is easier to adapt from the group to individually rather than the opposite. Uses encounter calculators for the challenge level of the group and thinks about the strengths and weakness of the characters classes, thinks about what can be explored from the PCs backgrounds, magical items. Yes, primarily when they think about changing the focus of the narrative. When they have been doing too much combat they try to change to political themes, for example. That is why they keep an eye on the background, to stimulate them to communicate and unite the characters.

They do not abandon dice when they are rolling bad, even because they do not have a lot of them. But when they roll too close to 10 (3 up or down) they get annoyed and think it is bad luck, most challenges are above that.

Interview 3 Pre-pandemic in paper and now online using roll20 and world anvil.

Either 1 - browses the Monsters Guide, finds a cool one that is going to be the thematic axis and selects one or two more to assist in the same thematic axis (example: kobolds accompanied by dragons or goblins with worgs) or 2 - creates new monsters.

Yes, recently the created a mechanic of social function of a PC with 20 alternatives that have to be applied to win group experience points that give abilities that favors the group. Example: a PC with the role of a traitor gains points when they advance a step in a betrayal narrative in relation to another player. Rarely keeps track of NPC relationships only when it is an outlier like a very positive interaction. The most

important aspects would be good neighborhood, frictions, prejudice, mistrust, good festive experience (getting hammered together). No.

DM shield notes, world anvil. Interactions with NPCs and ambitions (if the PC wants to be seen as a hero by the city they must keep track if the PC is taking steps on that direction so that they can recompense the actions and make it full circle).

They try to not frustrate the player or make their abilities useless. Skills, attributes, items and spells to see how they are going to design encounters so that the players can enjoy and express their characters' capacities. Yes, exhaustion levels. Stuff that is not on the sheet is easily forgotten and they think exhaustion levels are an interesting mechanic to explore. Privileges individual aspects. Players hardly architect an action flux coordinately, unless the situation is very atypical. They design according to the capacities that can be explored and if the players do not explore it at least they gave them the opportunity. To encourage collaboration and coordination, they think of a few encounters that demand more work, difficulty beyond the succession of individual initiatives. They complexity the scenario, such as putting the players and villain in a mobile pendular platform, adding various elements to pay attention. The very difficult encounter is made for them to think collectively and that is up to them. Yes, tanks can take more damage than a PC that has lower life and will zero if you hit. In order to not frustrate the player you have to prioritize the attack.

20 and 1. Three consecutives 1 (1/20³).

Interview 4 Online, even notes are all digital. They use roll20, discord (to take notes and schedule sessions), token generators, inkarnate for maps, reddit for images.

Challenge level using an encounter calculator, but mostly DMed pre-made campaigns to use. For adicional encounters they check if it fits with the original adventure. First they think about which monster would be cool then they check if it is in the desired level.

They prefer that the players maintain the relationships in their character sheets. A table of all can be too much stuff concentrated in one place, so this way they can go directly to the player. The most important factor is the type of relationship (good x neutral x bad) and whether it is "rivalry/nemesis/arch enemy" when bad or "super ally" when positive. It is interesting when the party itself embraces that. No, only once when the specific adventure required that.

Accesses the sheet directly. Mechanistically consults hp, attributes and abilities a lot (proficiencies and class abilities in general). For inventory they trust the players and do not pay much attention. No, the sheet is pretty thorough. They once quantified the PCs alignment using a table, but that was for a very specific campaign. For pre-made campaigns you do not have much of an option. For extra adventures for character arch they think of mechanics for each PC to have their time to shine. All characters are very unique. They compare players regarding roleplay, if one that is very immersed died the party would get loose itself and in terms of mechanics they notice things like if a character is more powerful.

1 and 20, rolling a lot of 11 is normal. If they roll three 1 they stay a good time not using that die and two 20 in a roll the die feels special.

Interview 5 Both, mostly online because of the pandemic. They use roll20 and discord, sometimes only discord with a shared spreadsheet and they tried notion for a homebrew but did not work out.

They use the Monsters Manual and browse by challenge level and check for the setting (they are not going to put a desert monster in the forest). Enjoys creating monsters using already existing stats.

No, usually the players themselves remember. When is something memorable they take into consideration but do not annotate, even because they DMed mostly on-shots.

When needed they use roll20 handouts and the players themselves take notes. Does not organize too much, prefers to improvise. Curses, promises, stuff that is relevant and can be used in favor or against the PC. Passive perception, main stats and whoever has the higher score for a group skill check. Yes, party inventory and money, pets because the players do not have access to each other's character sheet. By group and if they defeat too easily they increase the difficulty next time. No

They do not check only criticals. Too many fails, the die is cursed. A lot of successes, the die is good.

Interview 6 They start their ideas on paper and then when an idea can be expanded they use a doc. Uses google docs, dnd beyond, roll20, and dnd wikidot for rules.

Mainly coherence with the story and less the challenge level. If a monster fits the story they modify the strength. They research on dnd beyond, encounter generators, Monster's Manual and get inspired by other campaigns/games. A lot of times they just create it themselves and modify existing stats.

Tries to keep track. Lists by city all the NPCs and writes a brief summary of them and adds how they connect to other characters. If it is a family, they cluster them and put it on their bio. Top to bottom, macro to micro approach: city ζ group (work, where they go to) ζ individual. The potential to create stories, how the characters can evolve as you interact, the dynamics. If the relationship is friendly, why and how is it, how can that relate to other relationships, how can it evolve and alter the course of the campaign/story. PC-PC relations are not worth keeping track of, because each player already knows how they are going to act. PC-NPC also does not keep track, as it is more dynamic and they remember who does not like each other and the interaction to roleplay. The annotate more NPC-NPC to guide what their personalities are.

They create folders of the possible things they want to use, things the PCs mention about their past, how can they expand on that with the player without them even noticing. People from the PC's past, background, the way they see the world in order to try to make the world more dynamic/lived in. In game aspects they use passive perception and armor class a lot so that they do not need to ask the players all the time, and speed, saving throws, hp. It depends on what they want for the campaign, but usually in group. They try to make the monster not able to kill a character in one shot, but in general decide by group to not give too much time to one particular person or focus on one story only. Yes, in terms of comparing hp/tanks, to encourage someone to do something they have more ability on or not. In general tries to not stir it too much so that it is not forced.

Can not deny dice because they do not have a lot of them and are not superstitious about it. They use all the dice in the same frequency or use more the ones they are more use to or like the style more.

Interview 7 Online, using roll20 and dnd beyond.

Challenge level, quantity of PCs in the party and place of combat. They use <http://dnd5e.wikidot.com/>.

Yes, on a dedicated page. What interests them the most are conflicts of interest. It is fun to create conflict, for friendship and romance it is up to the players and does not matter much. Yes, they use a +/- relationship points system from the Strixhaven module.

They annotate on paper or a google doc (they have a master document with all the characters). It is useful for trauma, fears, backstory friends, enemies made throughout the game. Skills and saving throws. Yes, reactions they have to something, if the character gets scared they take note in case they need to scare them again or if a character found out something about another NPC. By group, if they are low level. They structure the dungeons, for example, with 3 challenges for one session and they want the players to pass it but not too high. If one character has high dexterity then they add a strength challenge, or if a lot of PCs are good with lockpicking they change the idea. To know how many PCs can pass an obstacle they think of who can surpass the dc (difficulty class) and if there are too many of them they redesign. Yes, in terms of possible interpretations. You expect an old and a young character to react differently to the problems brought to them and if they do not act accordingly to expectations they try to find out why.

Six consecutive results under 10 to change die and the opposite to consider it lucky. Easily superstitious with d&d.

Interview 8 Before the pandemic in paper, but now online using roll20 and google doc with the script.

They go by what fits in the plot and if there is an interesting map/scenario to go with it. Usually they use the d&d books, but to make sure it is compatible with the players' level they use an encounter calculator to balance medium-hard.

No, most of the time they keep it in memory, as most of their NPC are somewhat ludicrous they do not need to remember much details. Fun, they try to create characters that are very different from each other

to interact with the players.

Usually they do not take notes, but for specific stuff they used the DM notes in roll20. Which character/player has an item they give them in a previous session, if a relationship with an NPC was formed. Social matters like ideals, flaws, etc. to explore, because that is what the players use as a base for their stories. They do not care much about stats. Usually they propose challenges not focusing on individual characteristics, they leave it to secondary things. All players should have the same qualifications to complete a puzzle and battle. They try not to compare, but sometimes use the story from one PC to complement another.

They do not think about that, as they also do not own a lot of dice.

Interview 9 The pandemic turned it to mainly digital, with online rpg you can play with people outside your city. They are slowly going back to in person games like they used to play pre-covid and even then the online tools helped. They use roll20 and have not tried yet but want to test Owlbear Rodeo <https://www.owlbear.rodeo/> e Talespire <https://store.steampowered.com/app/720620/TaleSpire/>. Also name and token generators <https://rolladvantage.com/> <https://paper-tokens.firebaseio.com/>, base game book compilers (such as wikidot), inkarnate, <https://donjon.bin.sh/5e/dungeon/>, <https://tetra-cube.com/dnd/dnd-statblock.html>, <https://www.reddit.com/r/battlemaps/>.

They use the d&d books and <https://www.themonstersknow.com/>. First they like to have a general notion of what they want from that encounter to make sense. For an arctic region you expect polar bears, arctic wolves, monsters related to ice. Or you could add an unexpected factor to spark the players' interest. Why would a fire elemental be in an arctic region? They develop the why (not pure chaos) and focus on the narrative. They think it is interesting to break/subvert expectations (like adding a pacific orc) to engage or not (it is important to know your party). It is not worth investing much in political games if your party just wants to beat up monsters. Part of the entertainment of the DM is creating the world and describing things, they play too. To pick the monsters it varies, published adventures give you a limited option (but that does not impede you from adapting to your party). When creating their own adventure it's up to the proposal's discretion. To bring a mindflayer/demogorgon into a world with only humans and beasts (monstrous animals) breaks the rules of the world you created and the players' immersion. At the mechanical level, they observe the challenge level, the guide line system from the easy to mortal encounter. You can fill up an adventure with a lot of easy obstacles, because you have time. But for the deadliest there is a chance of a character dying. Sometimes the proposal is an absurd encounter to make players think outside the box or trigger an escape. If you achieve the purpose for that encounter without breaking immersion and not leaving anyone unsatisfied, then you did the work with mastery.

Usually they let the players free. They applied the Strixhaven mechanics to a homebrew once but they did not play long enough for that to come up. In their opinion it works well to let things flow naturally, 70% of the DM's work is improvising only 30% is actual planning. That includes planned NPCs. Sometimes you plan an awesome NPC you think the players are going to love but when you play that NPC is outshined by a random NPC you created at a pinch. They use paper or doc. They think it is interesting the collective storytelling that can be generated by the players feeling free to interact with the world. Another aspect is the organic relationships that can be formed by roleplaying above all, which is the special flavour rpg has; you can see players interacting mechanically elsewhere. They tried Strixhaven like mentioned, but generally do not quantify because that is organic and quantifying could break the immersion.

When playing virtually they check the character sheet on roll20 and take mental notes. Personality traits, ideals, flaws helps to know what to expect from the characters. That can also act as a summary of that character, but sometimes does not show all the layers involved. Yes, the PC's story and both the character's and player's behavior to make sure no one is being excluded. Also the PC's general vibe, the general idea the player has for the character (a fighter does not say much about a PC, but if you know battlemaster you can infer they are in a leadership

role), nuances that are only clarified with good communication. Ideals, passive perception, proficiencies, skills and the obvious ones for combat (hp, ac, speed, initiative). A sheet of paper can not transcribe all from the Player's Handbook and there is stuff you can not know by heart. No, the character sheet is pretty thorough. The particularities are in the player's administration of the sheet. Even though it is intuitive, it has a lot of information for a beginner. They say they have to work a lot on that point, they get stuck in the narrative and focus more on it making sense then on the players. Ideally it should be a bit of each. Example: 5 wizards have intelligence to spare, to make an encounter fun without knocking over players, you have to consider a strength test would fuck everyone up. Conversely, they could end up running over your combat and get too easy. For the character, planning and making an encounter fit well gets people to talking about it a long time later. You can use a character's background and make a monk's master reappear as evil and make the combat personal, touching the character and player. Mechanically they have to explore the player's abilities so that they do not abuse the encounter with an end-all spell for example. When you cut an hydra head another one grows so typically the character has to set fire to it, but the DM can subvert expectation and have the fire accelerate the process of growing another head and the PC has to use ice or lightning, improvise on the spot. They need also to pay attention to what they reward the players, careful not to shoot themselves in the foot. Yes, when paying attention to the party ecosystem. Ex if two PCs have 18 and 19 of strength and they keep stepping on each other's feet can generate internal conflict between players because they are fulfilling the same tank role/function. Without a secondary function you can resort to roleplay, but the DM must be aware of this. Having a party of 5 wizards has to be decided collectively so they don't run each other over and compete. It is interesting to compare how characters do by pulling two similar ones to another focus. Demand the players a little, one can be well immersed and cannot expect the same level of engagement from all but give a reminder to those who are not so much, to consider what to improve.

Dice is expensive, they don't have the luxury of switching cursed dice. They know ways to measure if a die is biased, using a saturated solution a standard die will float. They consider their superstition not for dice, but for person. A person is unlucky or lucky. Three 20, several 1. A bad roll is 1 to 7 more or less, 8 to 13 is acceptable/average, 14 upwards is a good roll.