

Content



- Introduction
- Literature Survey
- Phase 1: Named Entity Recognition (NER)
- Phase 2: Image Generation Pipeline
- Phase 3: Map Generation
- Phase 4: Adventure Book Generation

Project scope proposed initially

This is an extension to the initially proposed scope

Dungeons and Dragons



- Tabletop Role playing game
- A DM designs or follow a pre existing adventure
- Use props & images for immersion



A group of people playing D&D from the Critical Role web series.

Background	Reference	Character Hook
Acolyte	Player's Handbook	Temples are vexed by a widespread magical curse that is causing
100000		people to waste away while also preventing the dead from being
		raised. An expeditionary force is headed to Chult to find the origin
		of the curse, and you've decided to join it.
Anthropologist	Appendix A	A wizard-merchant named Syndra Silvane is wasting away, and she
- anim opologist	Appendix A	
		has hired you to join an expedition to Chult to find a cure. Syndra
		believes that your expertise will prove invaluable, and you've wanted
Archaeologist	Appendix A	to study the cultures there.
Aichaeologist	Appendix A	You have learned that a wizard-merchant is planning an important
		expedition to Chult-a place full of lost cities you're anxious to ex-
Charlatan		plore. Your request to join the expedition was approved.
Charlatan	Player's Handbook	After a few successful scams, you've gotten into some trouble with
		local authorities and criminal gangs. You were about to make a run
		for it when you heard a rumor that a merchant named Syndra Sil-
		vane is offering good pay for an assignment far from home.
Cloistered Scholar	Sword Coast Adventurer's Guide	A terrible curse is sweeping across Faerûn, and a dying merchant
		is gathering adventurers for a bold mission to destroy the source
		of the curse, which lies deep in the jungles of Chult. Your expertise
		could prove helpful, so you've asked to join the expedition.
Criminal	Player's Handbook	Local authorities grant you a full pardon for past crimes. In ex-
		change, you are to join an expedition to Chult and help put an end
		to the "death curse" affecting several wealthy citizens of the city.
Entertainer	Player's Handbook	Your sponsor is a retired adventurer and merchant named Syndra
		Filmer A trailing adventurer and merchant named Syndra
		Silvane. A terrible curse has befallen her, and she needs your help
		to end it. The quest will take you to Chult, a distant and exotic land
Faction Agent	Sword Coast Adventurer's Guide	where you can win fame and renown.
raction Agent	Sword Coast Adventurer's Guide	A widespread magical curse threatens to devour the bodies and
		souls of the living. The Harpers have learned that its source lies
Folk Hero	-	somewhere in Chult. You volunteered to join the expedition.
FOIK FIETO	Player's Handbook	You've garnered quite the reputation and caught the attention of a
		merchant and retired adventurer named Syndra Silvane. She wants
		you to explore Chult and discover the source of a curse.
Inheritor	Sword Coast Adventurer's Guide	A horrible curse threatens to claim the life of Syndra Silvane, a re-
		tired adventurer. She has blood ties to your family, and you have an
		obligation to help her. To do otherwise could Jeopardize your inheri-
		tance and your future.
Mercenary Veteran	Sword Coast Adventurer's Guide	A merchant and retired adventurer named Syndra Silvane needs
		trained warriors to join an expedition to Chult. You've always
		wanted to visit Port Nyanzaru and see the dinosaur races. They say
		the jungles are full of riches and danger—two things you crave.
Noble	Player's Handbook	Your family owes a favor to a merchant named Syndra Silvane. She's
	C. C	calling in the favor and asking for help. The family is counting on
		canning in the layor and asking for neip. The family is counting on
		you to fulfill its promise, and you've heard rumors that Syndra is
Outlander	Player's Handbook	planning an expedition.
o dianaci	Prayer's Hanapook	When you were young, you fled your homeland of Chult by stowing
		away aboard a ship. Now you're anxious to return home. A local
		merchant is mounting an expedition to Chult, and you've talked
Fa		your way into joining it.
Sage	Player's Handbook	A wizard named Syndra Silvane has fallen prey to a magical curse.
		She's mounting an expedition to Chult to end the curse, and she be-
		lieves your expertise will prove invaluable. You're inclined to agree.
Soldier	Player's Handbook	Your superiors in the military recommended you for an important
		expedition to Chult. The success of the mission is so vital that they
		promised to make you a captain if you help ensure the expedition's

INTRODUCTION

honey. It's more common and popular in Chult than been or ale. A mug of tej costs 4 cp in Port Nyanzaru or 6 cp in Fort Beluarian. A 1-gallon cask costs 2 sp in the city or 3 sp at the fort.

done well when paired with Small riders. Most competitors in the two-legged race are hadrosauruses and deinonychuses (again with Small riders). The unchained race sees racers on anything, including young allosauruses and very young tyrannosauruses.

BETTING	ON DIN	OSAUR KA	ACES		
Odds	Roll	Win on	Payout -	Racer Example	
7:1	d8	2-8	Wager + 7	Big Honker (crowd-pleasing young tyrannosaurus rex)	
5.1	dó	2-6	Wager + 5	Ubtao's Favorite (young triceratops on a streak)	
3:1	d4	2-4	Wager + 3	Banana Candy (well-trained hadrosaurus)	
2:1	d3	2-3	Wager + 2	Bonecruncher (young, ill-tempered allosaurus)	
1:1	Any	Even	Wager	Grung Stomper (young ankylosaurus)	
1:2	d3	1	Wager x 2	Scarback (young tyrannosaurus rex on its last race)	
1:3	d4	1	Wager x 3	Nasty Boy (allosaurus with famous trainer)	
1:5	d6	1	Wager x 5	Jungle Princess (deinonychus with frightened jackey)	
1:7	d8	1	Wager x 7	Mountain Thunder lazing dimetrodoni	

RACING DINOSAURS

Mount	Speed	Skill Check DC	AC	Hit Points	Attack	Damage	Constitution
Allosaurus, young	50/80	167	13	30	+6	8 (1d10 + 3)	15 (+2)
Deinonychus	40/60	12*	13	26	+4	6 (1d8 + 2)	14 (+2)
Dimetrodon	30/50	8	12	19	+3	8 (2d5 + 1)	15 (+2)
Hadrosaurus	40/50	10	11	19	+3	6 (1d10+1)	13 (+1)
Other	40/60	12	12	24	+3	6 (1d8 + 2)	16 (+3)
Triceratops, young	50/75	14	13	38	+5	6 (1d10 + 1)	15 (+2)
Tyrannosaurus, voung	50/700	18#	11	46	4.6	9 (1412 - 3)	17 (-3)

* If an Animal Handling check made with advantage due to lashing fails by 5 or more, the animal goes berserk and is out of the race.

WELCOME TO CHULT

With the exception of a few coastal settlements, Chult is untamed tropical wilderness: dense jungles and snaky rivers ringed by mountains, volcanoes, and sheer escarpments. Walls of mountains to the west, south, and east shield the interior from the sea and from the view of sailors. The rivers are so sluggish that it can be difficult determining which direction is upstream and which is down. The rivers pick up speed only where they thunder down through steep-sided gorges.

The safest entry points into this overgrown realm ("safest" is a relative term in this context) are on the north and east. The coast from the Bay of Chult to Refuge Bay offers beaches on which to embark into the uncharted jungle. Along the entire coast, the Bay of Chult is the only spot where travelers can find welcoming civ ilization. The rest of the peninsula is a breeding ground for bloodsucking, disease-bearing insects, monstrous reptiles, carnivorous birds and beasts of every variety. and murderous undead. The farther one moves from the coast, the more humid, hot, and inhospitable the

WEATHER

Chult is hot, humid, and rainy throughout the year. The temperature regularly climbs as high as 95 degrees F (35 degrees C) during the day and seldom falls below 70 degrees F (20 degrees C) even at night. A day without rain is rare, but rain varies from a steady mist to drenching downpours

Visibility in heavy rain is limited to 50 yards. Beyond that distance, only Huge or larger objects can be distinguished. Missile weapon ranges are halved during rain.

On days that receive heavy rain, there's a 25 percent chance of a full-blown tropical storm featuring sheets of rain, high wind, lightning, tall waves at sea, and immense surf along the coast. Guides who know the conditions in Chult recommend hunkering down and staying put on these days. Travel by river is impossible (canoes are swamped by waves within 15 minutes of launching). If characters insist on traveling by foot, every characte gains 1 level of exhaustion automatically and must make successful DC 10 Constitution check or gain another. Skill checks to avoid becoming lost are made with disadvantage on storm days.

RACES OF CHULT

Members of every race and nationality of Facrûn can be found in Port Nyanzaru, but few of them ever trave more than a hundred yards beyond the city walls. The native peoples must be divided into those that are com monly known and those whose presence is legendary or only suspected. The first group consists of Chultans (humans), aarakocra, Batiri (goblins), dwarves (including albino dwarves), pterafolk, and tabaxi. In the second category, weretigers are suspected to be present in small numbers; explorers and hunters sometimes report seeing grungs (frog folk) deep in the jungle; and yuan-ti are rumored to be making a resurgence in remote areas.

In addition to these known and suspected residents, a crew of frost giants has come to Chult seeking the Ring of Winter and an expedition of Red Wizards is hunting for the Soulmonger, but the presence of these bands is still unknown to the larger population.

Finally, the deep jungle is overrun by skeletons, zombies, ghouls, and other types of undead. Everyone in

warning horns, day or night, they scramble for the safety of the Market Ward and Merchants' Ward until the all clear sounds

OLD CITY

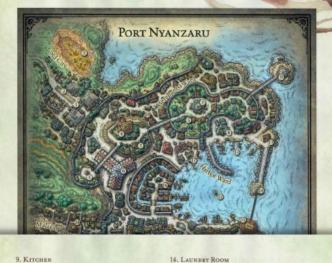
Three ancient, vine-covered ziggurats tower above this crumbling ward. The whole district is a juxtaposition of ancient and decaying (but still occupied) stone structures interspersed with flimsy new huts and longhouses

The Old City is run by "beggar princes" in a mocking parallel to the merchant princes of the city proper. They have no official authority, but each of them sits atop a web of debts, favors, incriminating information, and loyal muscle that allows them to get things done with an efficiency that the merchant princes sometimes every Unlike the merchant princes, whose positions are based on wealth that passes from generation to generation in the same family, the identities and even the number of the beggar princes changes continually

The Old City is run down but it isn't a slum or a haven for thieves. Most of the residents are lower-class laborers or struggling artisans who can't afford the higher rents of homes inside the city walls.

1. BEGGARS' PALACES

The Beggars' Palaces are the two largest ziggurats of the Old City. Shops and tenements of bamboo are ammed onto the terraces of the ziggurats in wild profusion. The upper levels are cleaner and less crowded than those near street level. A narrow wooden bridge connects the upper levels above the press and noise of



9. KITCHEN

The kitchen is large, high-peilinged, and well ventilated; otherwise, heat from cooking fires would make it unbearably hot. Stairs lead up to the laundry room

chant prince to whom the villa belongs.

10. LIBRARY

Chultans are fans of mythical romance, and the city has an active literary community with many book clubs. Port Nyanzaru's most successful povelists and playwrights concoct stirring tales about legendary heroes who pursue their lovers through unimaginable danger and save them from fates worse than death-which is more likely to involve a forced marriage to someone who doesn't love them than being eaten by an impossibly huge tyran nosaurus, although monstrous tyrannosauruses make regular appearances in these tales, too.

A character who spends a few hours sifting through the formulaic fiction and succeeds on a DC 15 Intelligence (Investigation) check finds useful books of Chultan lore mixed in with the romances. Use this as an opportunity to answer specific questions posed by the players. A library won't reveal the location of Omu, but a book of lore might mention the city and describe it as it was during its height.

The merchant prince's children, spouses, and other close family members live in these apartments. They are

12. MASTER SUITE

The merchant prince's suite is the most opulent room in tice-covered windows, causing the potted plants to sway carved in motifs suggestive of the merchant's business. combs, and so forth.

Treasure. The suite contains personal items equivalent to two rolls on the "Gems or Art Objects" column of the Treasure Hoard, Challenge 0-4 table in the Day geon Master's Guide. If the room contains jewelry, there is a 50 percent chance it is in a locked jewelry box reouiring thieves' tools and a successful DC 15 Dexterity check to open.

13. GUARD ROOMS

This guard room is intentionally set in the midst of the family rooms and the master suite. The guards are responsible for safeguarding the family during their most vulnerable time, when they're sleeping or relaxing in their own chambers. This station is always manned by 1d4 + 1 gladiators who are fanatically loyal to the mer-

The rooms on this level are decorated for comfort. The

II. FAMILY ROOMS

richly appointed with pillows and tiger-skin rugs.

the palace. A pleasant breeze blows through large, latgently. The headboard of the master bed is fantastically Connected rooms contain a small, private bath and a walk-in closet filled with expensive clothing, mirrors,

and righly rewarded for their devotion.

UPPER LEVEL

artwork and other decor reflects the taste of the mer-

The Emerald Enclave is a widespread group of wilderness survivalists who preserve the natural order by rooting out unnatural threats. They struggle to keep civilization and the wilderness from destroying one another.

LAWS AND PUNISHMENTS

EMERALD ENGLAVE

The merchant princes can create or alter laws by decree as they see fit, and judges are expected to interpret laws freely but wisely. Such a system would seem to invite abuse and bribery, but the Chultan courts are surprisingly free of Slavery isn't illegal in Port Nyanzaru, but it's frowned

Chultans are as scrupulous about their clothing as about

given how much everyone sweats in Port Nyanzaru's climate. Laundry is washed here every day, then hours are

spent fanning it to and fro on lines trying to encourage

The servants' quarters are isolated from the family sleeping areas for security. There is nothing remarkable

Every room in the house contains (at minimum) a small

water basin with running water coming either from a

rooftop basin or lifted from barrels or an underground

Port Nyanzaru is the lone civilized hub on a vast pen-

insula of great wealth. Many factions would like to tap

into that wealth, and they all have agents vying for influ-

ence in the city or scouting the jungle for opportunities. Some of these can aid the characters, and others will try

FACTIONS AND THEIR

REPRESENTATIVES

their bodies-a fastidiousness that's almost essent

the wet cloth to dry in the humid air.

15. SERVANTS' QUARTERS

16. RAIN TRAPS

to hinder them.

cistern by a water wheel.

upon, and the merchant princes won't do business with slavers or show them any legal favor.

Murder is illegal, but "sanctions"—legal permission to kill someone or to have someone killed—can be purhased from the merchant prince Jessamine. The price depends on the prominence of the intended victim, the The minimum price possible is 150 gp; the maximum is limitless. The availability of sanctions doesn't mean hey're common-it's rare for more than six to be sold in a year. That has less to do with price than with the fact that if word gets out who purchased a sanction, nothing can stop the friends of the victim from sanctioning the origina

Despite their high ambition, the merchant princes' power is entirely contained in the city. Elsewhere, Chult ruled by the red law of tooth and claw. An individual is either predator or prey, and the predator might be a mobof undead, a ravenous allosaurus, a grung hunting party. Jahakan pirates, or a Flaming Fist patro

LOWER LEVEL UPPER LEVEL

MERCHANT PRINCE'S VILLA

Map 1.2: Merchany Prince's Villa.

the DCs for different types of disosaurs are listed in the Racing Dinosaurs table. With each successful Animal Handling check, the first number listed as the dinosaur's speed is added to its "running tally." If the check fails, that dinosaur's tally doesn't increase that turn. When a racer's tally equals or exceeds 300, that dinosaur crosses the finish line. A racer can try to move at the higher listed speed by lashing the animal furiously;

ful DC 10 Constitution check at the end of this round or its speed is halved for the rest of the race. No initiative is involved. Riders can make their Animal Handling checks in any order, or all at the same time. If two or more dinosaura cross the finish line on the same round, the one with the highest tally wins. If the tallies are the same, the racers tied. In the event of a tie, brawls are likely to break out in trackside bet-

in this case, the Animal Handling check is made with

advantage, but the dinosaur must also make a success

each rider makes a Wisdom (Animal Handling) check;

In the unchained event, diposaurs can attack other dinosaurs if their tallies are within 50 of each other at the end of a round. Each dinosaur's pertinent combat statistics are listed on the Racing Dinosaurs table. If a dinosaur has half or fewer of its starting hit points, the rider makes Animal Handling checks with disadvantage

For simplicity, riders can't attack or be attacked, but DMs with a cruel streak might be tempted to relax that rule. Just be aware that if rider attacks are allowed, the

always on alert against tricksters and magicians trying to gain entrance to their master's home, so they have advantage on skill checks and saving throws against atsempts to distract, hamboozle, or charm them.

The ceiling of the entryway is two stories tall. In some homes, it's open to the sky. In those cases, the fountain is sunk into the floor and doubles as a drain for

2. SITTING AREA

This is both a waiting area and a meeting area for guests who aren't known to the villa's master, and therefore won't necessarily be given access to the inner chambers of the house. Decorative columns support the arched doorway from area 1. Rugs and comfortable pillows are arranged on the floor. Narrow windows look into the garden, but these double as arrow slits in case. visitors aren't as friendly as they claim to be.

3. GRAND HALL

A magnificently tiled floor is flanked by two sweeping staircases that ascend to the upper floor. Decorations are in keeping with the individual merchant prince's

Most of the right side of the house is devoted to accommodations for guests. Three private sleeping rooms with attached boths, plus two sitting rooms and access to the

Ferns, potted palms, orchids, and other exotic tropical flowers are cultivated in the garden, which is open to the sky. Chultan plants with useful properties or applications as poison might be found in the garden of a merchant prince with a horticultural bent, such as Jessamine. See appendix C for examples of rare plants.

6. DINING ROOM The dining table is very low, just a few inches above the floor. Diners relax on pillows as they eat.

In addition to daily baths, well-to-do Chultans love relaxing in saunas. Dry heat is a welcome change from Chult's high humidity. Most home saunas are heated with charcoal burners or hot stones, but in the homes of the merchant princes, a trapped fire elemental provides round-the-clock heat, and a servant or slave is always available to massage tired muscles.

The bath is the cornerstone of a Chultan mansion, even more so than the dining room or family rooms. The family relaxes here during private times, but there's nothing unusual about a merchant prince holding meetings in the bath with trusted advisors and close business associates. Being invited into the household bath is con-

Picture credit: Tomb of Annihilation (Wizards of the coast)

Research Problem



- Generating images consistent with the given narrative for a D&D adventure
 - Difficult to create art for a given narrative
 - Hard to maintain consistency of images across sessions
 - Images should be consistent with the lore
 - Pre existing or third party images needs to be incorporated



Literature Survey





Key Phrase extraction



Prompt generation



Image generation



Evaluation



Consistency

Literature: Key Phrase Extraction



Pre-processing & Summarization [1,2,3]

Key phrase extraction using machine learning [4,5]

Named Entity Recognition

[1] E. Papagiannopoulou and G. Tsoumakas, "A review of keyphrase extraction," Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery, vol. 10, no. 2, p. e1339, 2020.



Picture credit: Forgotten Realms Fandom Wiki

^[2] S. Vijayarani, M. J. Ilamathi, M. Nithya et al., "Preprocessing techniques for text mining-an overview," International Journal of Computer Science & Communication Networks, vol. 5, no. 1, pp. 7–16, 2015.

^[3] R. Srivastava, P. Singh, K. Rana, and V. Kumar, "A topic modeled unsupervised approach to single document extractive text summarization," Knowledge-Based Systems, vol. 246, p. 108636, 2022.

^[4] A. P. B. Veyseh, N. Meister, F. Dernoncourt, and T. H. Nguyen, "Improving keyphrase extraction with data augmentation and information filtering," arXiv preprint arXiv:2209.04951, 2022

^[5] S. Beliga, "Keyword extraction: a review of methods and approaches," University of Rijeka, Department of Informatics, Rijeka, vol. 1, no. 9, 2014.

Literature: Prompt generation



- Prompt Engineering and Design Principles [6]
- Control over Text Generation with Templates and Constraints [7]
- Generative Pretrained Transformer (GPT) Models
 - Task-specific finetuning [8]
 - Quantifying event boundaries in continuous narratives [9]

What's the difference between prompts?



Original Prompt	Extended Prompt	Original Prompt	Extended Prompt
A dream of a distant galaxy	A dream of a distant galaxy, concept art, matte painting, HQ, 4k	photo of a riverbank	photo of a riverbank, concept art, matte painting, HQ, 4k

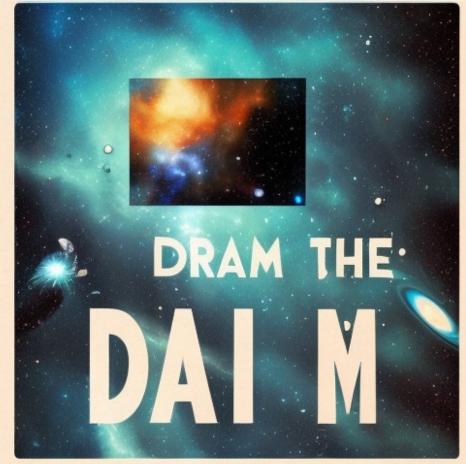
Negative prompts



"A dream of a distant galaxy"

Reduce artifacts

 out of frame, lowres, text, error, cropped, worst quality, low quality, jpeg artifacts, ugly, ...



THESE BY SERVICE ORANGE THE AND THE THE REAL PROPERTY OF THE P

Literature: Image Generation



- Image Synthesis with Generative Adversarial Networks [10]
- Variational Autoencoders (VAEs) for Image Generation [11]
- Diffusion models for Image Generation [12]
- Image Style Transfer [13]
- Image-to-Image Translation [14]

[10] Y. Lu, S. Wu, Y.-W. Tai and C.-K. Tang, "Image Generation from Sketch Constraint Using Contextual GAN," in Computer Vision – ECCV 2018, Lecture Notes in Computer Science, vol 11213. Springer, Cham, 2018

[14] J. Ho, C. Saharia, W. Chan, D. J. Fleet, M. Norouzi, and T. Salimans, "Cascaded diffusion models for high fidelity image generation." J. Mach. Learn. Res., vol. 23, no. 47, pp. 1–33, 2022.



^[11] AutoEncoders are Essential in Deep Neural Nets. Towards Data Science. Retrieved March 21, 2023, from https://towardsdatascience.com/autoencoders-are-essential-in-deep-neural-nets-f0365b2d1d7c

^[12] J. Ho, C. Saharia, W. Chan, D. J. Fleet, M. Norouzi, and T. Salimans, "Cascaded diffusion models for high fidelity image generation." J. Mach. Learn. Res., vol. 23, no. 47, pp. 1–33, 2022.

^[13] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2019, pp. 4401–4410.

Literature: Consistency



- Document-Level Sentiment Analysis [15]
- Document-Level Relation Extraction [16]
- Image Temporal Consistency [17,18]



Video 1: Example of using temporal consistent images to form a video[9].

Literature: Evaluation



- Evaluation Metrics for Text Generation [19]
- Image Evaluation [20]
- Consistency evaluation [21]







Elminster Picture credit: Forgotten Realms Fandom Wiki



Picture credit: Forgotten Realms Fandom Wiki



Phase 1:

Comparative Analysis of Named Entity Recognition in the Dungeons and Dragons Domain

Model Comparison



- **Pre-trained Models:** BERT[22], **ELECTRA**[23], and XLM-RoBERTa[24] are general pre-trained models that can be fine-tuned for NER.
- NER-specific Models: WikiNEuRal[25], RoBERTaNER[26], and BERT-CRF[27] are specifically designed or adapted for NER.
- Toolkits & Libraries: Spacy[28] and Trankit[29] are comprehensive NLP libraries/toolkits that provide functionalities beyond just NER.
- Domain-specific: StanfordAIMI[30] is a domain-specific, focusing mainly on medical tasks.
- Embedding Techniques: Flair[31] is unique in its use of contextual string embeddings.

^[22] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805.

^[23] Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. 2020. Electra: Pre-training text encoders as discriminators rather than generators. In International Conference on Learning Representations.

^[24] Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzman, Edouard Grave, Myle Ott, Luke Zettle- 'moyer, and Veselin Stoyanov. 2019. Unsupervised cross-lingual representation learning at scale. arXiv preprint arXiv:1911.02116.

^[25] Simone Tedeschi, Valentino Maiorca, Niccolo Campol-`ungo, Francesco Cecconi, and Roberto Navigli. 2021. WikiNEuRal: Combined neural and knowledgebased silver data creation for multilingual NER. In Findings of the Association for Computational Linguistics: EMNLP 2021, pages 2521–2533.

^[26] Jean Baptiste. 2022. roberta-large-ner-english: model fine-tuned from roberta-large for ner task. https://huggingface.co/Jean-Baptiste/ roberta-large-ner-english. (Accessed on 05/10/2023).

^[27] Fabio Souza, Rodrigo Nogueira, and Roberto Lotufo. '2019. Portuguese named entity recognition using bert-crf. arXiv preprint arXiv:1909.10649.

^[28] M. Honnibal and I. Montani, "spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing," unpublished.

^[29] M. Van Nguyen, V. D. Lai, A. P. B. Veyseh, and T. H. Nguyen, "Trankit: A light weight transformer-based toolkit for multilingual natural language processing," arXiv preprint arXiv:2101.03289, 2021. 32

^[31] A. Akbik, T. Bergmann, D. Blythe, K. Rasul, S. Schweter and R. Vollgraf, "FLAIR: An Easy-to-Use Framework for State-of-the-Art NLP," in Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: System Demonstrations, pages 54-59, Minneapolis, Minnesota, 20191.

D&D adventure books



Book	Counts		
	Words	Topics	
Lost mine of Phandelver (Baker and Perkins, 2014)	45947	29	
Hoard of the Dragon Queen (Baur et al., 2014a)	74243	45	
Rise of Tiamat (Baur et al., 2014b)	80065	48	
Curse of Strahd (Perkins et al., 2016)	154519	62	
Tomb of Annihilation (Perkins et al., 2017)	148605	35	
Candlekeep Mysteries (Perkins et al., 2021)	141104	106	
The Wild Beyond the Witchlight (Allan et al., 2021)	184135	60	



^[32] Wizards of the Coast, (2023, March 28). LOST MINE OF PHANDELVER. D&D Beyond. [Online]. Available: https://www.dndbeyond.com/sources/lmop

^[33] Wolfgang Baur, Steve Winter, and Alexander Winter. 2014a. Hoard of the Dragon Queen. Wizards of the Coast.

^[34] Wolfgang Baur, Steve Winter, and Alexander Winter. 2014b. Rise of Tiamat. Wizards of the Coast. [35] Christopher Perkins, Tracy Hickman, and Laura Hickman. 2016. Curse of Strahd. Wizards of the Coast.

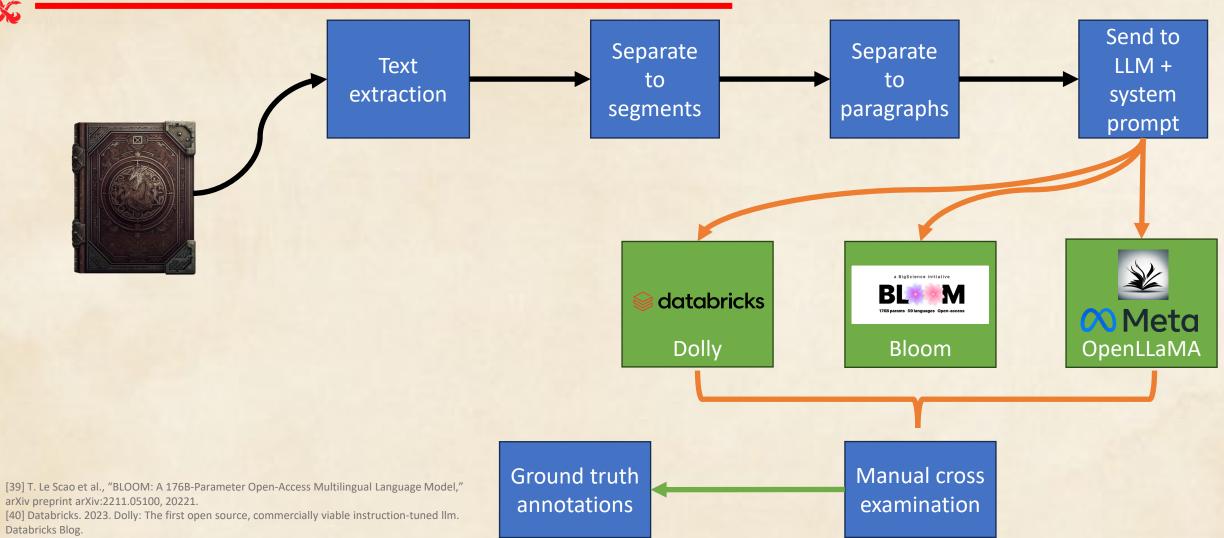
^[36] Christopher Perkins, Will Doyle, and Steve Winter. 2017. Tomb of Annihilation. Wizards of the Coast.

^[37] Christopher Perkins, Graeme Barber, Bill Benham, et al. 2021. Candlekeep Mysteries. Wizards of the Coast.

^[38] Stacey Allan, Will Doyle, Ari Levitch, and Christopher Perkins. 2021. The Wild Beyond the Witchlight. Wizards of the Coast.

Data Collection and Annotation





[41] Xinyang Geng and Hao Liu. 2023. Openllama: An open reproduction of llama.

Data Collection and Annotation



```
Input: Books;
Output: Named entities;
foreach book do
   segments \leftarrow divideIntoSegments(book);
   foreach segment in segments do
       paragraphs \leftarrow
        divideIntoParagraphs(segment);
       foreach paragraph in paragraphs
        do
           foreach LLM in LLMs do
               prompt \leftarrow
                createPrompt(paragraph);
               namedEntities \leftarrow
                LLM(prompt);
               processNamedEntities(namedEntities);
           end
       end
   end
end
removeDuplicates(namedEntities);
```

Annotation process

Chapter	Topic	Paragraph	Word Coun		
		The main			
		antagonists of this story are		131	
		131			
Introduction:	Adventure	three hags			
Into the	Summary	One of the			
Feywild	many novelties		43		
reywiid		of this adventure	43		
		is that			
		The characters			
		are drawn into			
		the adventure	31		
		by one of two	31		
		adventure hooks.			
		You choose			
		Chapter 1			
		describes the	40		
		Witchlight Carnival			
		The Monster			
		Manual contains			
	Running the Adventure	stat blocks	70		
		for most of the	72		
		creatures encountered			
		in this			
		Spells and			
		equipment mentioned			
		in the adventure			
		are described	31		
		in the Players			

Please identify and list all named entities					
in the following text using the BIO					
(beginning-inside-	outside) sche	eme:			
"The traveling ext	ravaganza kn	own as			
the Witchlight Car	rnival visits y	our			
world once every	eight years. Y	You			
have a dim memor	ry of sneakin	g into			
the carnival as a cl	hild without j	paying			
pair of elves named Mister Witch and					
Mister Lightwere decidedly unhelpful."					
B-Organization:	Witchlight	Carnival			

System prompt, Text to be annotated and result (Bloom, OpenLLaMA)

Mister

Mister

I-Person:

I-Person:

Using the BIO (beginning-inside-outside) tagging scheme, identify and label named entities in the provided text from D&D adventure books. The entities of interest are Persons (Person), Locations (Location), Organizations (Organization), and Miscellaneous. The text is: "The traveling extravaganza known as the Witchlight Carnival visits your world once every eight years. You have a dim memory of sneaking into the carnival as a child without paying... ...pair of elves named Mister Witch and Mister Light were decidedly unhelpful." Any ambiguities or errors in entity recognition should be removed.

System prompt, Text to be annotated (Dolly)

Witch

Light

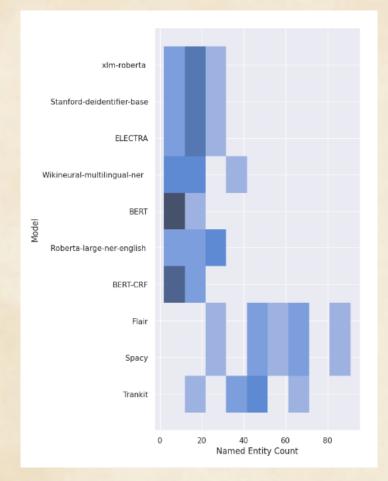
Result Comparison between LLMs



Book	Bloom		Dolly		Oper	Total Unique	
	Count	Recall	Count	Recall	Coun	t Recall	Entities
Lost Mine of Phandelver	21	0.47	32	0.73	40	0.91	44
Hoard of the Dragon Queen	58	0.89	62	0.95	60	0.92	65
Rise of Tiamat	54	0.88	57	0.93	53	0.87	61
Curse Of Strahd	92	0.90	96	0.94	10	0.99	102
Tomb of Annihilation	101	0.80	99	0.79	112	0.89	126
Candle keep Mysteries	60	0.87	61	0.88	64	0.93	69
The Wild Beyond Witch Light	66	0.84	67	0.85	7	0.89	79

Results





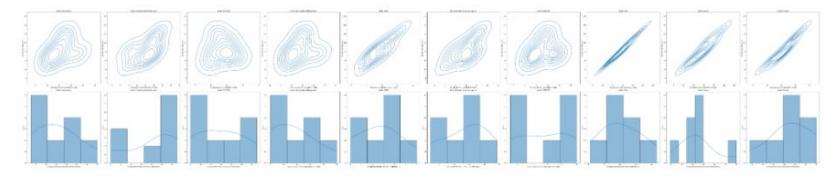
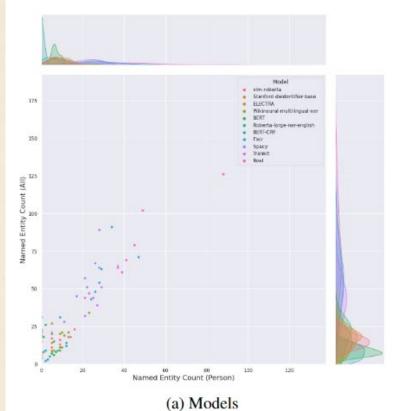
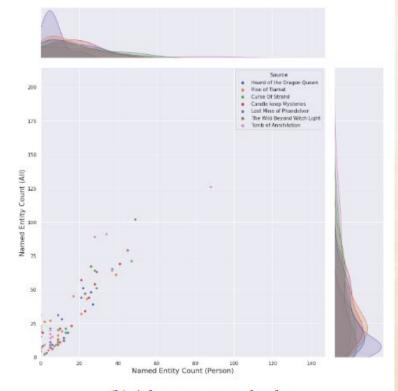


Figure 4: Distribution plot for each model

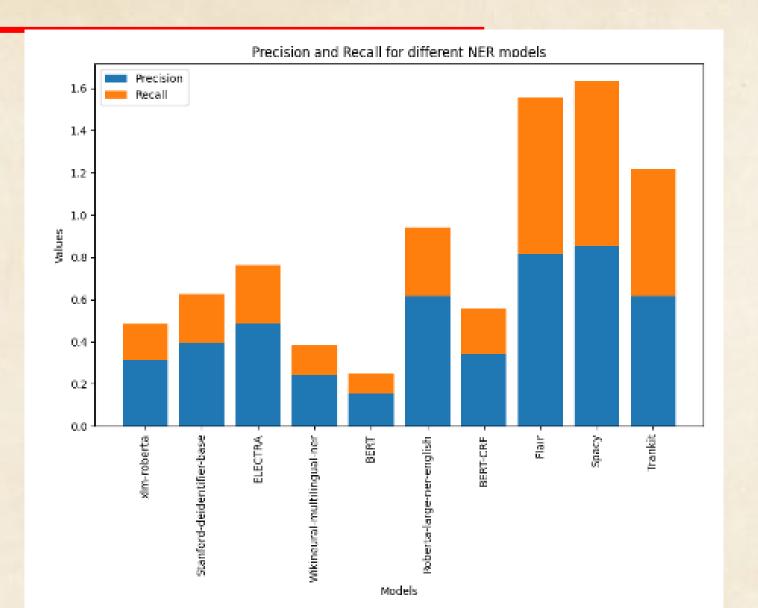




(b) Adventure sourcebooks.

Results





Recent Advances in Natural Language Processing (RANLP) 2023

Paragraph

One of the

is that...

of this adventure

The characters



Comparative Analysis of Named Entity Recognition in the Dungeons and Dragons Domain

Gayashan Weerasundara and Nisansa de Silva Department of Computer Science & Engineering University of Moratuwa, Sri Lanka

{gayashan.22, NisansaDdS}@cse.mrt.ac.lk

Abstract

Many NLP tasks, although well-resolved for general English, face challenges in specific domains like fantasy literature. This is evident in Named Entity Recognition (NER), which detects and categorizes entities in text. We analyzed 10 NER models on 7 Dungeons and Dragons (D&D) adventure books to assess domainspecific performance. Using open-source Large Language Models, we annotated named entities in these books and evaluated each model's precision. Our findings indicate that, without modifications, Flair, Trankit, and Spacy outperform others in identifying named entities in the D&D context

1 Introduction

Named Entity Recognition (NER) targets the identification and classification of textual entities, such as names and locations. In the diverse and intricate vocabulary of fantasy literature, like that of Dungeons and Dragons (D&D), NER becomes challenging (Zagal and Deterding, 2018). D&D, a prominent fantasy literature domain, spans content for its namesake tabletop game (Peiris and de Silva, 2022, 2023; Zhou et al., 2022). These narratives inhabit fictional realms like Forgotten Realms and Dragonlance, bursting with characters, locations, and objects (Gygax and Arneson, 1974).

NER's utility in fantasy literature is vast: from extracting information and summarizing text to character analysis and plot creation. However, conventional NER models, primarily trained on standard datasets like CoNLL-2003 (Tiong Vim Sang and Da Mauldar 2002) or OntoNota

2017), and the dynam dia (de Silva and Dou research emphasizing et al. (2021) points out, challenges, stressing th models specifically on

Fantasy NER has po vancements in image g cation might involve a leveraging NER tags to quently produce contex

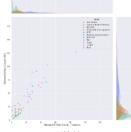
Our study contrasts 1 D&D books, each aven ual annotations of enti posed against model o assessments and named we glean insights into fantasy domain. Our ke

- · A pioneering, co evaluation on fant
- · An annotated D& studies.
- · A deep dive in strengths and pitfa
- · Discussions on N fantasy literature.

Following this, Sec NER and fantasy literat our data and annotatio unveils our methods ar 6 respectively discuss research, and Section



Figure 4: Distribution plot for each model



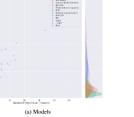




Figure 5: Frequency plots with respect to models and adventure sourcebooks

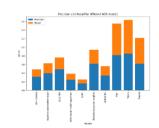


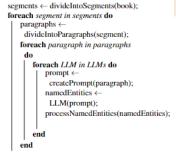
Figure 6: Precision graph for different NER models

models within the D&D context but also sets the stage for continued innovation at the confluence of fantasy literature and artificial intelligence.

7 Future Works

Based on our findings and limitations, we suggest some directions for future research. One direction is to fine-tune NER models on the D&D dataset and comparing their performance with off-the-shelf models. Additionally, other techniques such as transfer learning or domain adaptation could be explored to improve the performance of NER models in the D&D domain. Another direction is to use different data sources for NER in D&D, such as continued research and refinement to tailor these novels comics nodeasts or video games. A third

Input: Books: Output: Named entities; foreach book do



removeDuplicates(namedEntities); Algorithm 1: Named Entity Recognition using

Multiple LLMs

end

creature that

oes, villains,

r faction that

gion that has

sion, such as

) not belong

This contain

ls, Artifacts,

ich as guilds,

and each paragraph is iteratively fed into each of the LLMs with a prompt to identify named entities. The named entities identified by each LLM are then processed and saved. Finally, all named entities are checked for duplicates, and those duplicates are

After named entities were recognized, they were then mapped in to json objects for storage as shown in Figure 1. Nesting of objects is done according to the hierarchy as mentioned in table 2. Each of the named entities were nested in an array of entities as entity objects with corresponding attributes as mentioned bellow.

4 Experimental Setup and Results

The experiment was conducted to identify how ef-

Book	Bloom		Dolly		OpenL	Total	
		D		D II		D II	Unique
	Count	Recall	Count	Recall	Count	Recall	Entities
Lost Mine of Phandelver	21	0.47	32	0.73	40	0.91	44
Hoard of the Dragon Queen	58	0.89	62	0.95	60	0.92	65
Rise of Tiamat	54	0.88	57	0.93	53	0.87	61
Curse Of Strahd	92	0.90	96	0.94	101	0.99	102
Tomb of Annihilation	101	0.80	99	0.79	112	0.89	126
Candle keep Mysteries	60	0.87	61	0.88	64	0.93	69
The Wild Beyond Witch Light	66	0.84	67	0.85	71	0.89	79

Table 3: Result comparison between LLMs

JSON object

Please identify and list all named entities in the following text using the BIO (beginning-inside-outside) scheme:

"The traveling extravaganza known as the Witchlight Carnival visits your world once every eight years. You have a dim memory of sneaking into the carnival as a child without paying... pair of elves named Mister Witch and Mister Lightwere decidedly unhelpful." B-Organization: Witchlight Carnival I-Person: Mister Witch

Tab	lo 4	- D	rococc	of.	Annotatio

Mister

Light

I-Person:

Model	PER	LOC	ORG	MSC	All
XLM-RoBERTa (Conneau et al., 2019)	16	0	3	4	23
StanfordAIMI (Chambon et al., 2023)	0	0	1	18	19
ELECTRA (Clark et al., 2020)	10	0	1	10	21
WikiNEuRal (Tedeschi et al., 2021)	23	4	6	1	34
BERT (Devlin et al., 2018)	9	1	1	0	11
RoBERTaNER (Baptiste, 2022)	1	0	0	17	18
BERT-CRF (Souza et al., 2019)	12	0	0	0	12
Flair (Akbik et al., 2018)	28	14	6	4	54
Spacy (Honnibal and Montani, 2017)	21	11	7	18	57

Table 5: Statistics for the adventure book Candlekeep Mysteries. The NER tags are as follows, Person: PER, Location: LOC, Organization: ORG, and Miscellaneous: MSC

The resultant output is refined by filtering out corrupted values (e.g., "Strahd Von Zarovich" might

1
"book": "Candlekeep_Mysteries",
"chapter": 1,
"text": "The_Book_of_Inner_
Alchemy_is_one_of_Candlekeeps
"entities": [
{
• • • •
},
"entity": "B-Location",
"score": 0.9659823,
"index": 8,
"word": "Candlekeep",
"start": 42,
"end": 51
},
1 "
1
r

Figure 1: sample format of the JSON output

segmentation of terms, or misinterpretation of special characters. Such discrepancies can be mitigated using string manipulations and by crossreferencing outputs with a pre-curated list of named

Figure 2 displays entries that encountered corruption. These highlight instances where NER models incorrectly processed and extracted entities from the source material



Phase 2:

A Multi-Stage Approach to Image Consistency in Zero-Shot Character Art Generation for the D&D Domain

Literature: Generative Al & LLMs



- LLMs (Bloom, Dolly, OpenLLaMA etc.)
- Diffusion models (Stable diffusion [43], DALL-E etc [44])
- Multi-agent LLM Orchestration (Autogen [45], MetaGPT [46] etc)

^[46] Hong, S., Zheng, X., Chen, J., Cheng, Y., Wang, J., Zhang, C., ... & Wu, C. (2023). Metagpt: Meta programming for multi-agent collaborative framework. arXiv preprint arXiv:2308.00352.

Literature: Stable Diffusion







Experiment: Applying style to images



portrait (painting) of tabaxi, de Rivia closeup, suit, collar, formal attire, D&D, fantasy, intricate, elegant, highly detailed, artstation, concept art, matte, sharp focus, (brush strokes), (oil on canvas), hearthstone, art by Titian and Greg Rutkowski and Rembrandt van Rijn and Alphonse Mucha



No style transfer



With style transfer

Experiment: Iterative changes





Image of a dragon, concept art, matte painting, HQ, 4k (Stable diffusion)



make only the dragon red (InstructPix2Pix [36])



dragon is inside a cave



cave is a crystal cave

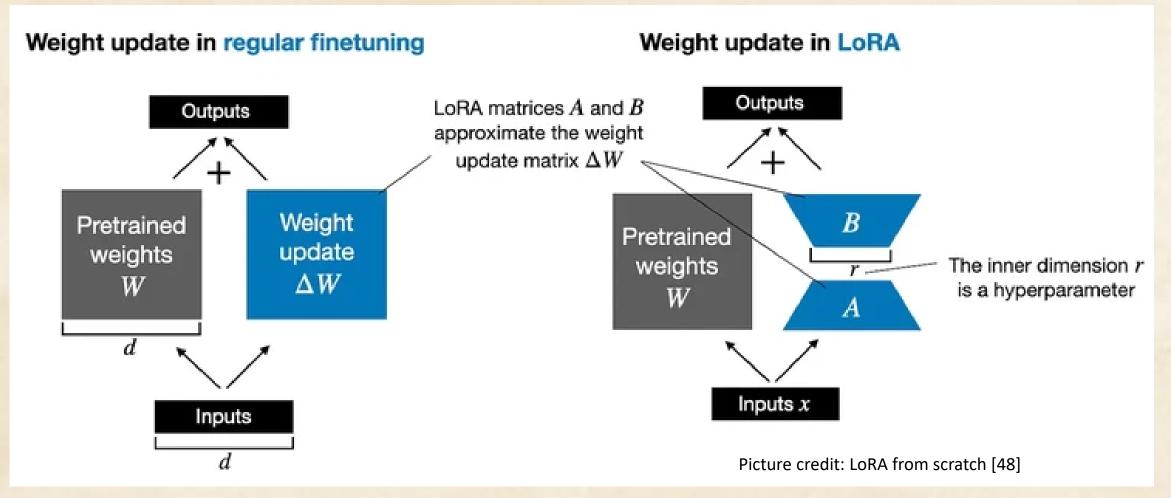


make only the dragon red

Image Generation Pipeline Text Pose Sheet Segmenting Image Encoder Processing with Yolo Input Prompt ControlNet Extracting Segmented Style **Images** Embedding Captioning Character With WD14 Sheet Lora Diffusion Model Remove Common Words and LoRA Consistent Insert a **Images Unique ID** 39

Literature: Low Rank Aduptation (LoRA)



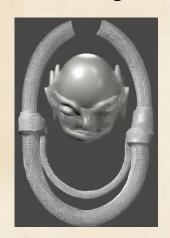


Experiment: Training a General Art Style





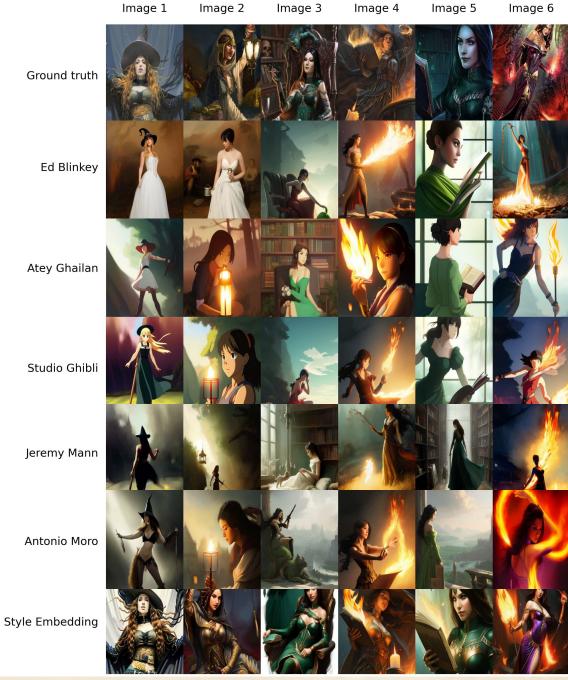
Elf in a general built model



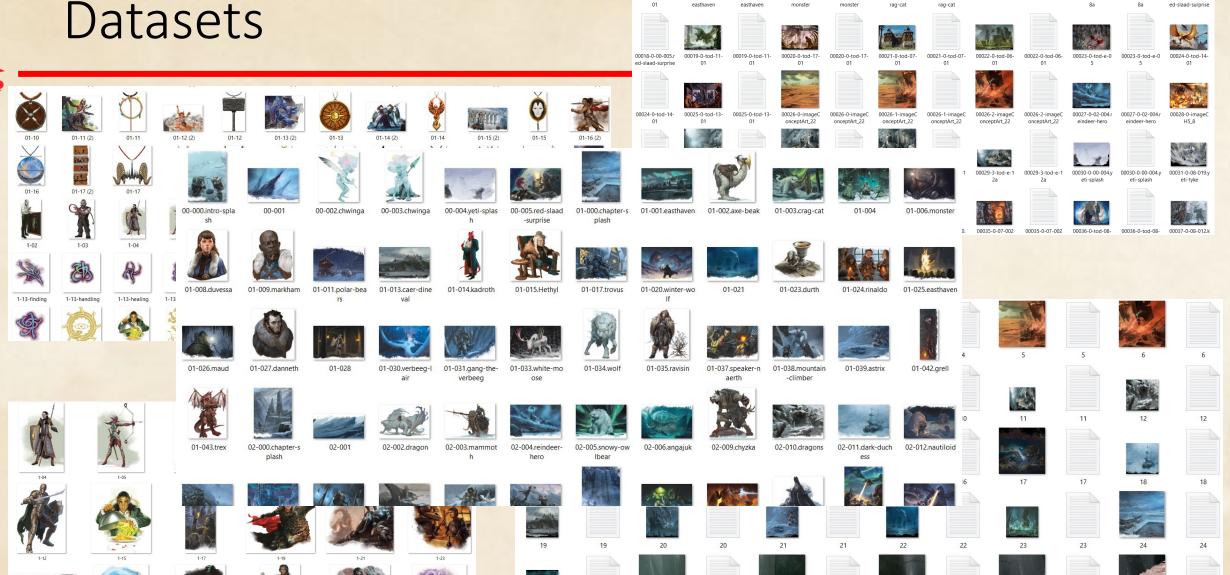




Elf in a style embedded model







25

26

27

27

28

00012-0-tod-03-

00013-0-01-001.

00014-0-01-006.

00014-0-01-006.

00015-0-01-003.c 00015-0-01-003.c

00016-0-00-001

00016-0-00-001

00017-0-tod-e-0 00017-0-tod-e-0

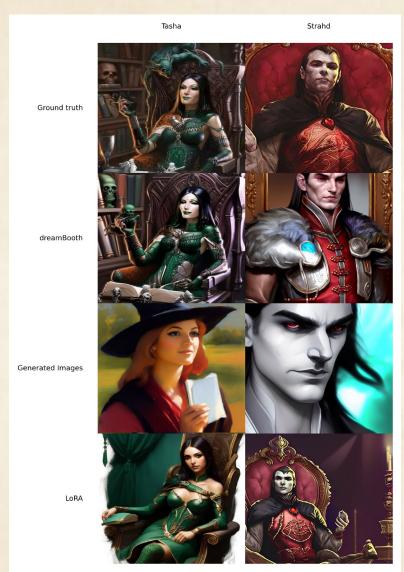
30

29

Experiment: Training a Specific Character Generation

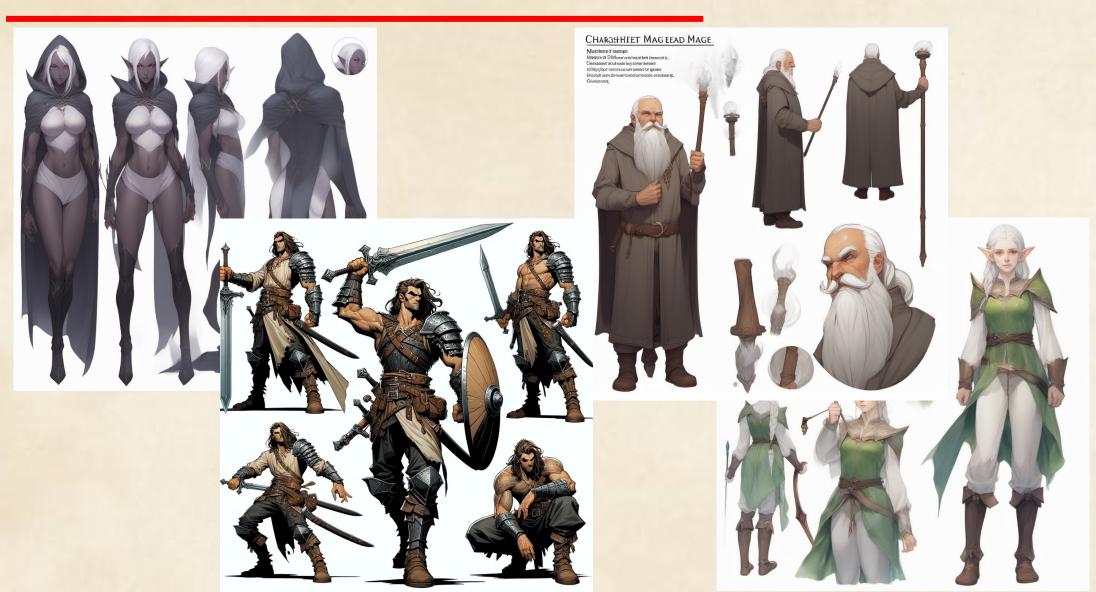


Dreambooth	LORA
Greater in Size	Minute in Size
More VRAM/ resources	Less VRAM/ resources
Need much less data but more dependent on dataset quality	Can supplement dataset quality with relevant text/ descriptions but need more data



Experiment: Character Sheet Generation











2fe7e1df2bb402 9069c1d8e5bd78 1461 (2)



2fe7e1df2bb402 9069c1d8e5bd78 1461 (3)



2fe7e1df2bb402 9069c1d8e5bd78 1461



3b6e4e3a41ba2a 844cddaa7a02f1 837f (2)

4e3e2c3d7a6e99

833854675dd644

7b35 (2)



3b6e4e3a41ba2a 844cddaa7a02f1 837f (3)

4e3e2c3d7a6e99

833854675dd644

7b35 (3)



3b6e4e3a41ba2a 844cddaa7a02f1 837f



4a29adbc7f4ba3 45ade3bfe8004b 8998 (2)





4a29adbc7f4ba3 45ade3bfe8004b 8998 (3)



55 45ade3bfe8004b

5abe1591d739d8

0cfc90f9a5d1ead

b72



55 af0fc392f3508a9 af0fc392f3508a9 3d8--character-s 3d8--character-s heet-art-d (2)







af0fc392f3508a9

5acbcf5a7dd402 5acbcf5a7dd402 cea130af060c40b cea130af060c40b 884 (3)



57



57



4db37ed3b0e9f4 35e951ea071ed6 d3dd (2)



5b9756952ce5b3 3e24fcf8309ad11 bbd (2)



4db37ed3b0e9f4

35e951ea071ed6

5b9756952ce5b3 3e24fcf8309ad11





4db37ed3b0e9f4

35e951ea071ed6

5b9756952ce5b3 3e24fcf8309ad11



5ce18605d49736 837f57e1b73c29 e5dd (2)



5ce18605d49736 837f57e1b73c29 e5dd (3)



4e3e2c3d7a6e99

833854675dd644

7b35

5ce18605d49736 837f57e1b73c29 e5dd



a37e (2)

5abe1591d739d8

0cfc90f9a5d1ead

b72 (2)

5d0b0941d5654a ab0aabd5e7b552



5abe1591d739d8

0cfc90f9a5d1ead

b72 (3)

5d0b0941d5654a ab0aabd5e7b552

a37e (3)



5d0b0941d5654a ab0aabd5e7b552 a37e



5acbcf5a7dd402

cea130af060c40b

884 (2)

5df1329776ec05 9ab91e8ce750cc 2290 (2)



5df1329776ec05

9ab91e8ce750cc 2290 (2)



5df1329776ec05



5f98388a6bb005 a609976f965741f 200



9ab91e8ce750cc 2290 (3)



9ab91e8ce750cc 9ab91e8ce750cc 220U (3)



5df1329776ec05 9ab91e8ce750cc ววดก



5eb4943fb1c493



5eb4943fb1c493 03e6b60a937469 16/12 (2) 16/12 (3)



03e6b60a937469



5eb4943fb1c493 03e6b60a937469 f185e882ea4376 16/12



3cc (2)









200 (2)

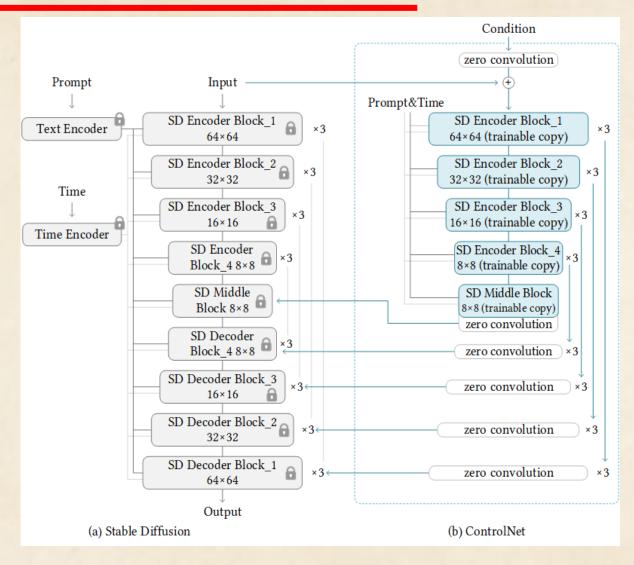


5f98388a6bb005 a609976f965741f 2UU (3)



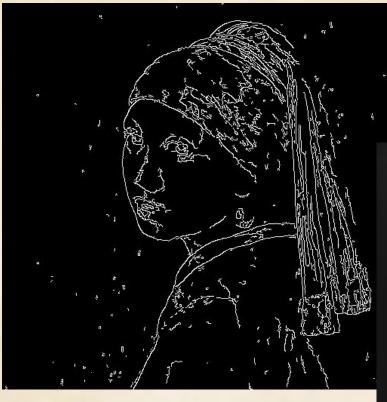
Literature: Controlnet



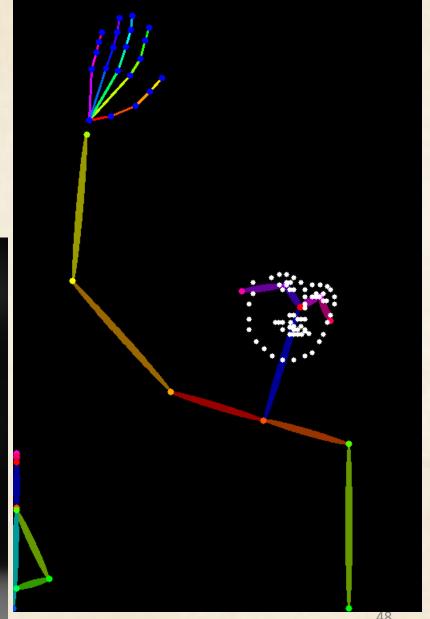


Literature: Controlnet





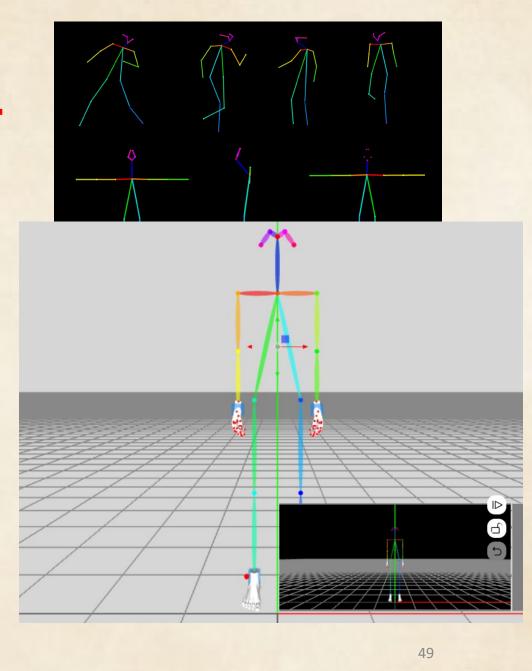




Experiment: Pose Sheets



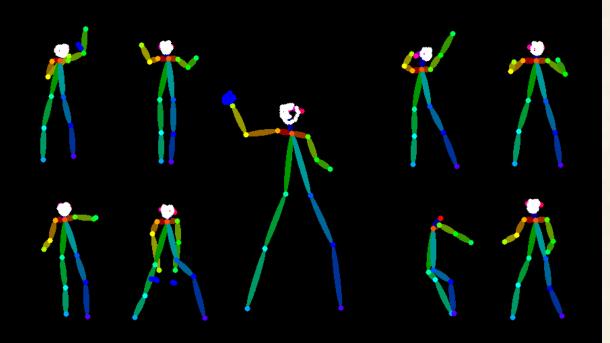




Experiment: Control Output

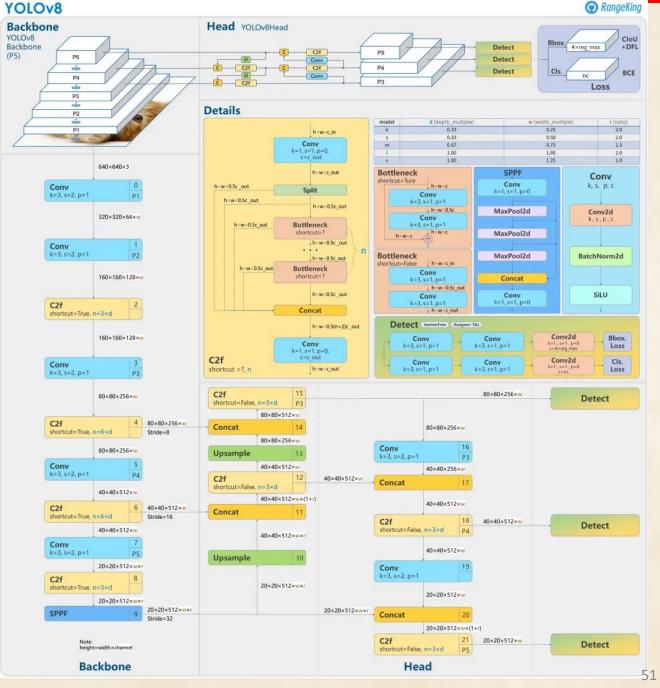






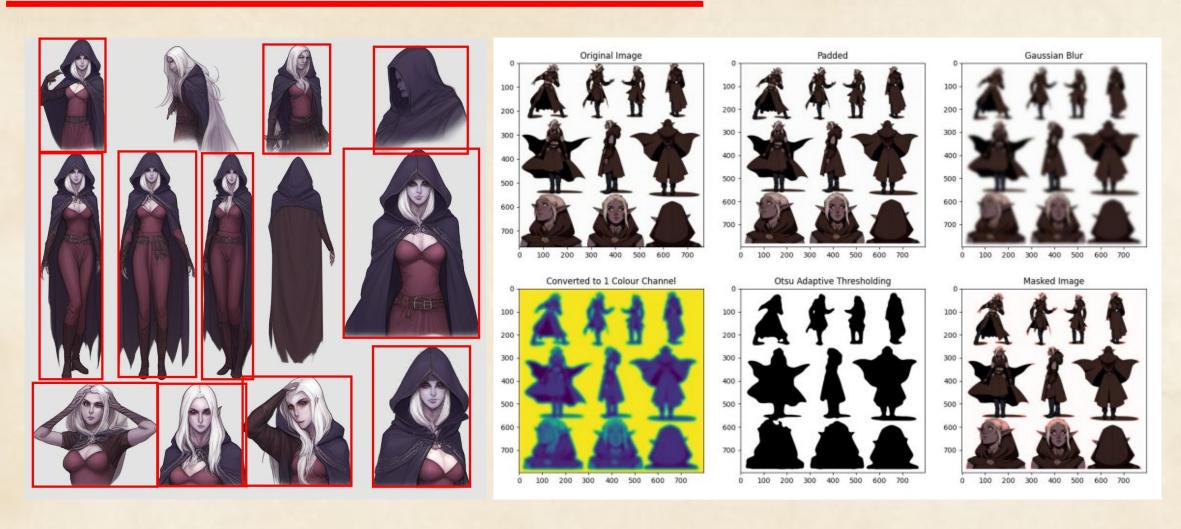
Literature: YOLO





Experiment: Separation and Training for Specific Character





Experiment: Clustering



- DINOv2 (feature extractor) [51]
- K-MEANS++ (clustering method) [52]



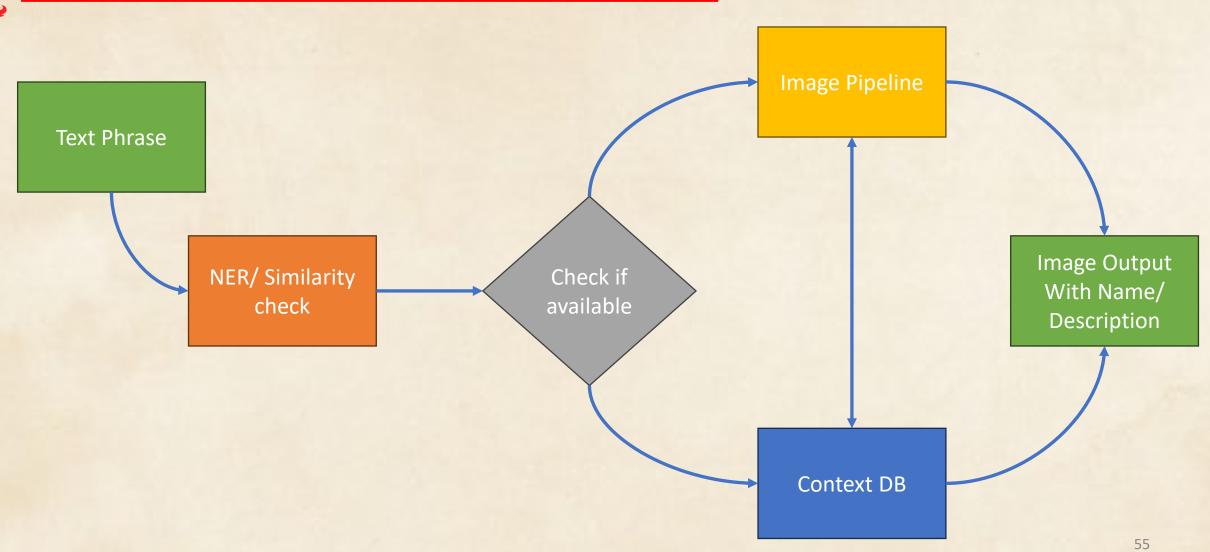






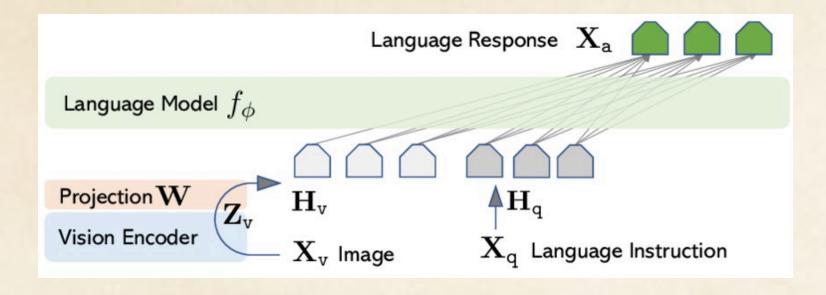
Workflow





Literature: LLaVA





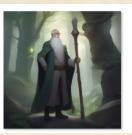
D:\Gayashan\Personal\MSc\ml\LLaVA
Image Path: ./ecom_arch (1).png

The diagram represents a cloud-based architecture on AWS for a live video streaming application with integrated payment processing. Clients interact with the system via an Authentication Layer that directs traffic through an AWS Application Load Balancer (ALB). AWS S3 and CloudFront distribute static content, while AWS ECS (Elastic Container Service) manages the application containers stored in AWS ECR (Elastic Container Registry). The application leverages AWS RDS with a MyS QL database for data management. Live video processing is handled separately, indicating a specialized workflow for real—time video handling. CloudWatch is used for monitoring the entire infrastructure, ensuring performance and availability. Additionally a Payment Gateway is integrated for processing transactions—indicating a monetization aspect of the live

Results









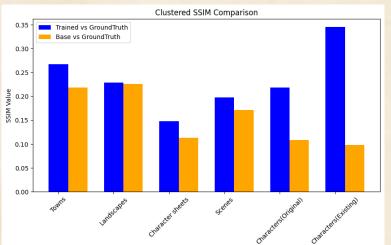
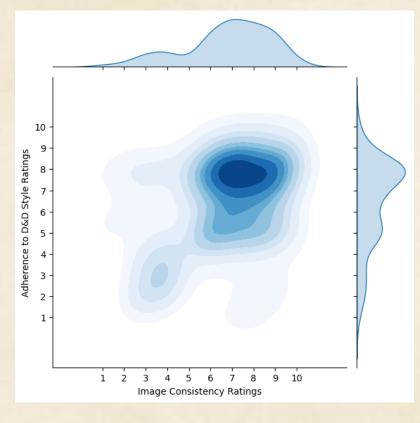


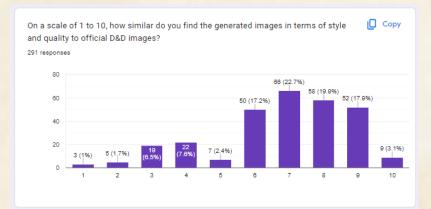


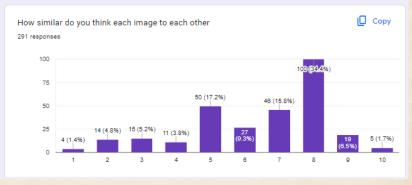




	Image to Image	InstructPix2Pix	Proposed Method
LLaVA (Liu et al., 2023)	6.4	7.5	8.3
FID (Heusel et al., 2017)	51.7	38.9	13.2







Specifications



Parameter	Pre-trained Character LoRA	Zero-Shot Character LoRA	D&D Style LoRA	Character-sheet LoRA
Number of Images	40 - 60	4 - 12	676	488
Repeats	10 - 15	20 - 40	2	4
Epochs	15 - 25	20 - 30	12	15
Precision	bf16	bf16	bf16	bf16
Learning Rate	0.0001	0.0001	0.0003	0.0002
Warmup	0	0	0 - 5%	0 - 5%
Resolution	1024x1024	1024x1024	1024x1024	1024x1024
Optimizer	Adafactor	Adafactor	Adafactor	Adafactor
Batch Size	2 - 4	1	4	4
Gradient Checkpointing	True	True	True	True
Xformers	False	False	False	False

Table 1: Training configurations for LoRA models trained for the proposed pipeline

Component	Specification
GPU	Nvidia RTX 4090
VRAM	24GB
RAM	61GB
CPU	16vCPU

International Conference on Agents and Artificial Intelligence (ICCART) 2023





Figure 8: Architecture of the pipeline used for the system

runpod.io³ using the Kohya ss GUI⁴. The parameters were fine-tuned, and multiple iterations were conducted to optimize the performance of the Stable Diffusion (SD) model and the LORA model.

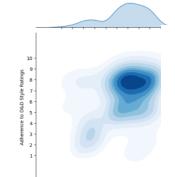
5 RESULTS

The results of our experiments were promising, showcasing the capability of our methodology in generating high-fidelity and consistent character images. The generated images exhibited remarkable consistency in style, pose accuracy, and visual quality, adhering to the predefined artistic style of the D&D universe.

To validate the effectiveness of our approach, we conducted a comparative study against established benchmarks, focusing on the structural similarity index measure (SSIM) (Zermani et al., 2021). SSIM is a widely recognized metric for comparing the similarity between two images, providing insights into the perceptual changes between the generated images and the ground truth.

Figure 9 presents the average SSIM comparison between the official images, baseline generations, and the images generated using the trained model per each category of images which indicates a greater consistency between official art with trained models compared to base models.

In our study, we evaluated the consistency and D&D style adherence of generated images through an



on a scale of 1 to 10.

Figure 10: Density Distribution of Feedb

1 2 3 4 5 6 7 8 9 1

Figure 10 density plot synthesize The x-axis shows character consiste the y-axis indicates style adherence. eas represent a higher consensus amo The results demonstrate that our me produces images that are both consis depiction and faithful to the D&D st by the concentration of responses in the

The plot reveals a notable cluster in the high-density regions, suggest agreement among participants. The a significant trend wherein our image proach has successfully maintained consistency in character portrayal at

In the evolving landscape of computer graphics, the pursuit of compelling challenge. This paper delineates a comprehensive r the Stable Diffusion (SD) model, an adept open-source image ge fidelity imagery. Beginning with the training of a generalized game Dungeons and Dragons, our research journeyed through st exhaustive character sheets, culminating in the generation of cor

1 INTRODUCTION

The computer graphics field has recently transformed with new image generation techniques, emphasizing realism and consistency, crucial in gaming, animation, and virtual reality. This is particularly relevant for the Dungeons & Dragons (D&D) universe, where consistent, high-fidelity character images are vital to maintaining its authentic high-fantasy (Peiris and de Silva, 2022) atmosphere. However, most recent studies (Peiris and de Silva, 2023; Weerasundara and de Silva, 2023) focused on natural language processing aspects of the game rather than image generation. Recent strides in diffusion models, especially the Stable Diffusion (SD) model (Rombach et al., 2022), noted for producing quality images consistently (Ho et al., 2020), have not been fully explored in D&D character generation.

Our paper explores using the SD model to create consistent, high-quality D&D character images. We discuss training a generalized art style, transitioning to character generation, and producing character sheets and consistent poses. Our approach combines recent research with the SD model's capabilities, aiming to advance computer graphics applications, particularly in the D&D realm. We detail our methodology, compare it with existing techniques, and analyze our results. Our goal is to present a new approach to computer graphics, focusing on the D&D universe.

In the evolv story visual notable con et al. (2022) transformers story contin as PororoSV et al., 2018) 2017), achie tion. Jeong et

creating stor Rahman et ally consiste et al. (2022) tent diffusio into the reas ers. The wo resolution is Liang et al. image synth

Other stu ing Maharai sistency in v ration of tex (2018)'s dev thesis mode

characters in various poses in a blank background. The first image shows a sample reference sheet.

foundation or our researcn, ınnuencing our approacn and objectives in text-to-image synthesis and story visualization.



Figure 4: Character sheets generated using the trained LORA

deavours. It juxtaposes the images generated using LoRA, DreamBooth, and the base model against the

official character art of Tasha and Strahd. The results were illuminating. The LoRA model, in particular, stood out for its exceptional capability.

3.3 Character Sheet Generation

Leveraging a generated character image, our objective was to create a comprehensive character sheet using a specialized LORA and a ControlNet (Zhang and Agrawala, 2023).

3.3.1 ControlNet

We selected controlnet-openpose-sdxl-1.01 due to its bias towards generating full-body images. First, we generated an image of a character and then keeping the seed value, prompt, scale and negative embedding consistent, we provided a reference pose sheet to the ControlNet. The system then extrapolated this image, generating a detailed character sheet inclusive of different poses and variations.

Figure 3 shows character sheets generated using ControlNet. Generated images consist of specific

3.3.2 Training a LORA

While ControlNet allows the generation of character sheets, it is not consistent. In a random generation, it is possible to get distorted images very easily. To make the generation more consistent, we need to embed the concept of a character sheet into the generated images. For that, we used a LORA model. To train a LORA, we need to create a labelled dataset of reference images. We used web scraping on Google images with keywords such as Character sheet and reference sheet and manually filtered undesired images. Resultant images were captioned with WD142 tagger to get preliminary labels. Next tags which are common to most of the images such as white background, Character sheet, reference sheet, turnover are removed while adding a common unique identifier as a trigger word. Finally, a LORA for that specific character can be trained on the labelled image dataset. Results for character sheet LORA can be seen in Figure 4

3.4 Separation and Training for Specific Character

Upon obtaining the character sheet, our focus shifted towards extracting undistorted and apt images that

³https://www.runpod.io/

⁴https://github.com/bmaltais/kohya_ss

a https://orcid.org/0000-0003-1419-8938

b https://orcid.org/0000-0002-5361-4810

https://bit.lv/3RKSPNa

²https://bit.lv/3RuKhc7

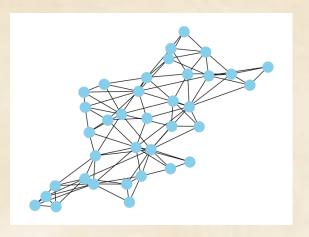


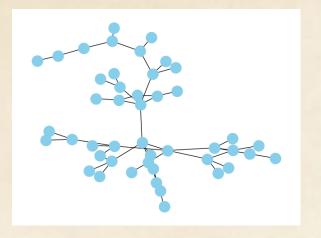
Phase 3:

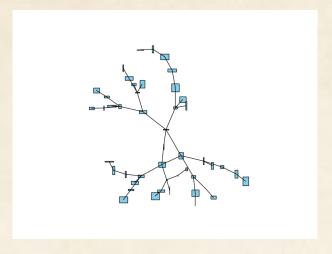
Map Generation

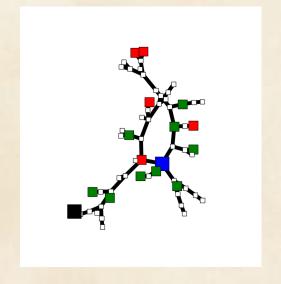
Experiment: Force Directed Graph Drawing

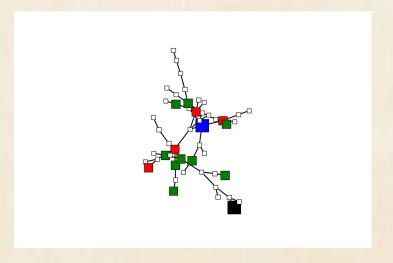


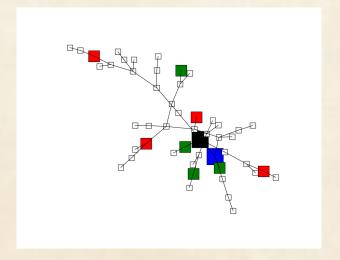






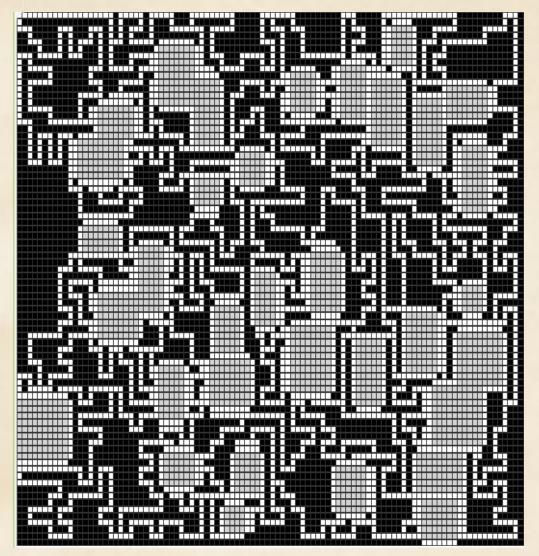


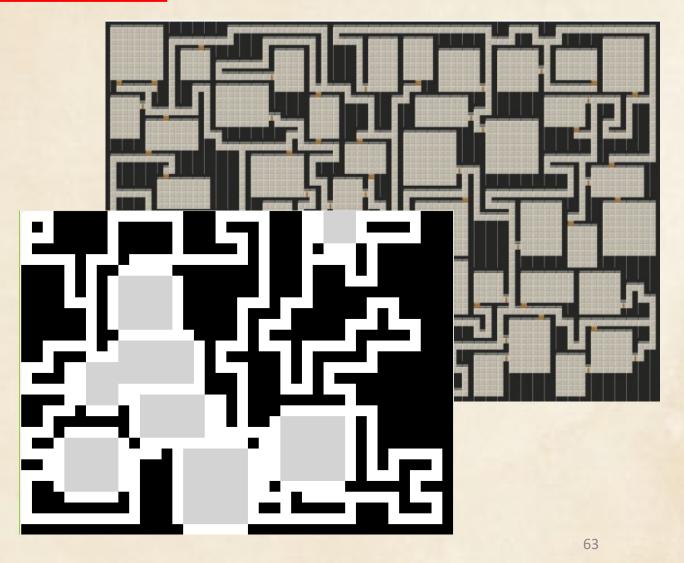




Experiment: Wave Function Colapse

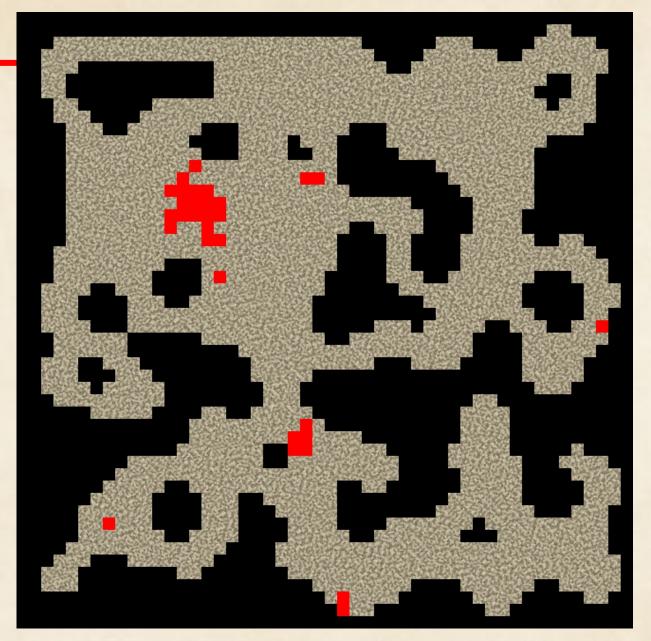






Experiment: Cellular Automata with Perlin Noise





Experiment: Consistent Dungeon Generation









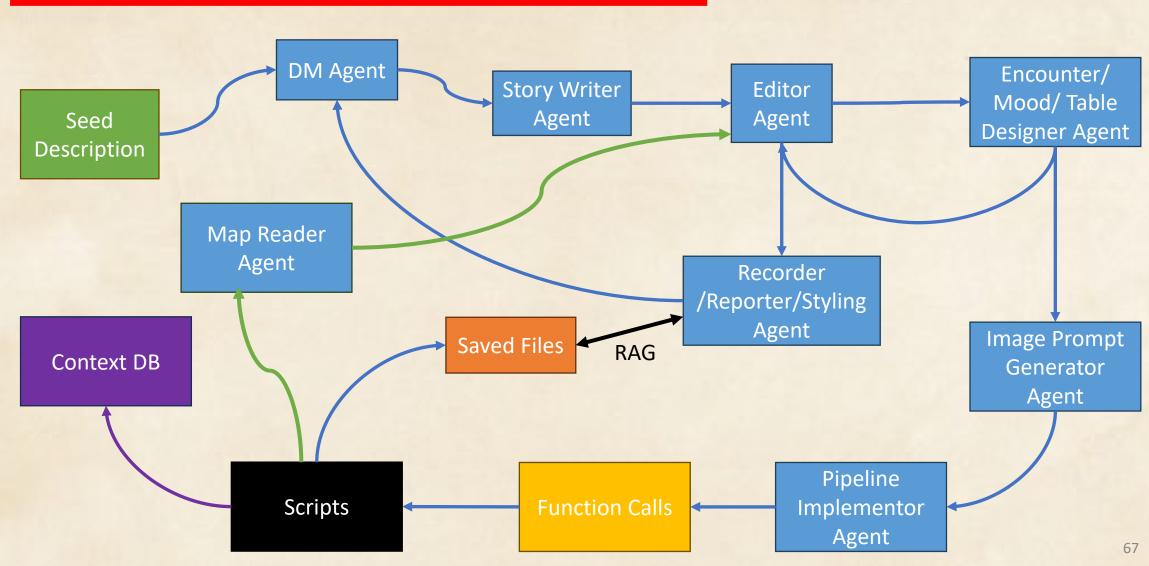
Phase 4:

Multi Agent Workflow for Adventure Generation

(This is beyond the initial scope)

Multi Agent Workflow





Agents

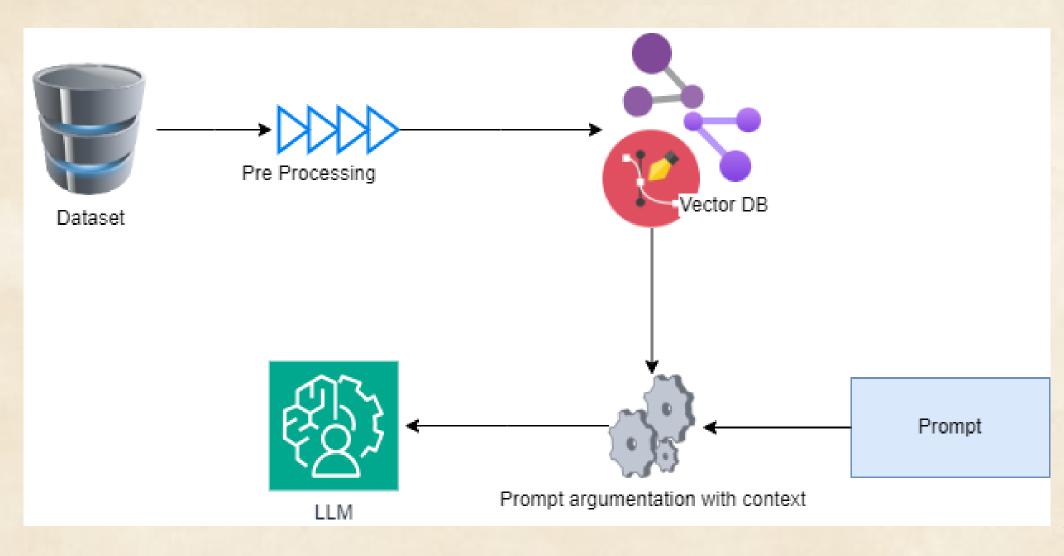


Role	Access
This agent is tasked with crafting immersive quests,	Chat history,
integrating storylines, NPCs, and encounters. It	Agent Control
serves as the central coordinator, aligning the contri-	
butions of other agents with the overarching narrative.	
Cordinate content creation agents with formatting	
agents such as Table Designer Agent, Stylize Agent	
and Mood Planner Agent	
This agent is tasked with writing the overall story.	Chat history
This agent gets feedback from the Dm agent, editor	
agent and reporter agent.	
Responsible for creating compelling NPCs, this	Chat history
agent builds characters that add depth and complex-	
ity to the campaign, ensuring they fit seamlessly into	
the storyline.	
Specializes in designing detailed combat scenarios,	Chat history
contributing to the dynamic and challenging aspects	
of the adventure.	
Identifies what type of table is needed for a given	Chat history
situation, for example an encounter table, a loot table,	
a stat block or any other, and create corresponding	
tags to generate a table based on markdown syntax	
	integrating storylines, NPCs, and encounters. It serves as the central coordinator, aligning the contributions of other agents with the overarching narrative. Cordinate content creation agents with formatting agents such as Table Designer Agent, Stylize Agent and Mood Planner Agent This agent is tasked with writing the overall story. This agent gets feedback from the Dm agent, editor agent and reporter agent. Responsible for creating compelling NPCs, this agent builds characters that add depth and complexity to the campaign, ensuring they fit seamlessly into the storyline. Specializes in designing detailed combat scenarios, contributing to the dynamic and challenging aspects of the adventure. Identifies what type of table is needed for a given situation, for example an encounter table, a loot table, a stat block or any other, and create corresponding

Mood Planner	This agent work on visual aids, settings, and atmo-	Chat history
Agent	sphere, enhancing the overall mood and tone of the	
	campaign. This agent identifies which areas need	
	visual aid and describe what type of visual aid is	
	needed. For example, if an image is needed in a cer-	
	tain place or a table is needed. It gives a description	
	of the mood and surrounding for important places	
	that require visual aids.	
Recorder Agent	Save content created by different agents as appropri-	Chat history,
	ate .txt, .csv files	Function
		Calls, Write
		access
Reporter Agent	Tasked with tracking and summarizing the cam-	Chat history,
	paign's progress, ensuring coherence and continuity	RAG Access,
	in the narrative and keeping track of special annota-	Function
	tions done by formatting agents.	Calls
Editor Agent	Ask agents to modify or improve their content, stop	Chat history
	the conversation or breakaway from a conversation	
	loop	
Stylize Agent	Add styling to different content created by agents	Chat history
Image prompt gen-	Generate tag based prompts for an image for a given	Chat history
erator Agent	situation	
Pipeline imple-	Use generated image prompts to generate images	Chat history,
menting Agent	through running required python script with given	Function
	arguments	Calls, File
		Execution

Literature: Retrieval augmented generation (RAG)





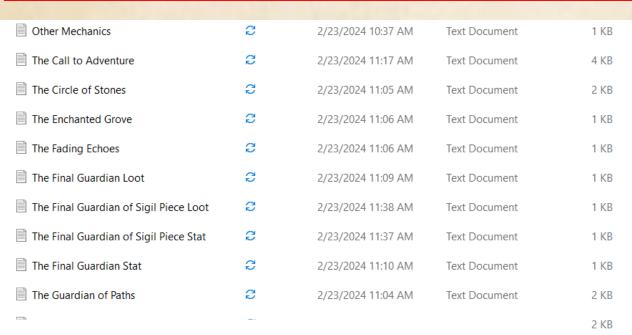
Tables



File	Content Type	Description
npc.csv	Non-Player Characters	Contains detailed information about
		NPCs including names, backstories,
		roles in the campaign, and other
		characteristics.
mood.csv	Campaign Mood Settings	Lists various mood settings and at-
		mospheres for different parts of the
		campaign, helping to set the tone
		and feel for each scene.
encounter.csv	Encounter Details	Provides specifics of encounters in-
		cluding type, difficulty, location,
		and involved NPCs or monsters.
lootTable.csv	Loot Tables	Details the loot available in vari-
		ous encounters or locations, includ-
		ing items, rarity, and any special at-
		tributes.
encounterTable.csv	Encounter Structuring Elements	Structuring elements for encounters,
		detailing the sequence and interac-
		tion of events within an encounter.
interactionTable.csv	Interaction Framework	Outlines the framework for NPC and
		player interactions, setting guide-
		lines for dialogue and engagement.

Styling





Rule	Annotation	Format	3 KB
Main title	_M	[content]_M	2 KB 1 KB
Heading	_H	[content]_H	0 KB
Table	_T	[content]_T	
Stat block	_S	[content]_S	
General text content	_G	[content]_G	
List	_	_[content]	
File	@	@[path]_[Heading]_[Content type]_[file type]	

72

The Summoning T

@./The Summoning.png I png

The adventure kicks off when the players are summoned by Archmage Eldrin, a respected and powerful wiz Archmage explains the dire situation—that the balance of the world is at risk due to the Sigil's destr provides the players with their first clue, pointing them towards the ancient ruins hidden deep within

Eldritch Forest T

The journey to the ruins is not straightforward. The Eldritch Forest is a living, breathing entity fil have to navigate through the forest's deceptive paths, facing challenges that test their survival skil mystical side of Eirandel, as they encounter spirits, ancient guardians, and possibly even the forest'

- @./The Whispering Glade.txt The Whispering Glade T txt
- @./The Guardian of Paths.txt The Guardian of Paths T txt
- @./The Mirror Lake.txt The Mirror Lake T txt
- @./The Circle of Stones.txt The Circle of Stones T txt
- @./The Enchanted Grove.txt The Enchanted Grove T txt
- @./The Fading Echoes.txt_The Fading Echoes_T_txt
- @./The Final Guardian.txt_The Final Guardian: The Entwined Serpent_G_txt
- @./The Final Guardian Stat.txt The Final Guardian: The Entwined Serpent Stat Block S txt
- @./The Final Guardian Loot.txt The Final Guardian Loot Table T txt

The Ancient Ruins T

Upon reaching the ruins, the players must breach its defenses, which include not only physical barrier filled with traps, puzzles, and enemies that guard the Sigil piece. This segment is designed to challe explore their characters' capabilities and how they can synergize as a group./n

The First Piece and Return T

The climax of Act 1 occurs when the players face the final guardian of the Sigil piece. This encounter learned so far about combat in Eirandel. Upon securing the Sigil piece, the players must make their wa a moment of realization that their journey is only just beginning. Archmage Eldrin, upon seeing the first piece recovered, referates the importance of their quest and hints at the dangers

and wonders that await them as they seek out the remaining pieces./n

- @./The Final Guardian of Sigil Piece Stat.txt The Final Guardian of Sigil Piece Stat Block S txt
- @./The Final Guardian of Sigil Piece Loot.txt The Final Guardian of Sigil Piece Loot Table T txt

Act 1 Objectives T

- Introduce players to the world of Eirandel and the main quest.
- Foster team cohesion and allow players to explore their characters' abilities.
- Challenge players with a mix of combat, exploration, and puzzle-solving.

Set the tone for the adventure, highlighting the significance of their mission and the mix of dark and light elements they will encounter.

THE SUMMONING



The adventure kicks off when the players are summoned by Archmage Eldrin, a respected and powerful wizard who has taken a keen interest in the prophesied reassembly of the Sigil. The Archmage explains the dire situation-that the balance of the world is at risk due to the Sigil's destruction and that its pieces must be recovered to prevent a cataclysmic event. Eldrin provides the players with their first clue,

reassembly of the Sigil. The ent a cataclysmic event. Eldrin ce of the Sigil./n

n and malevolent. Players will troduces the players to the er their progress./n

geon crawl segment of Act 1, bilities, encouraging them to

he players on everything they have eroes. However, the return is also

73

The Summoning_T

@./The Summoning.png_I_png

The adventure kicks off when the players are summoned by Archmage Eldrin, a respected and powerful wiz Archmage explains the dire situation—that the balance of the world is at risk due to the Sigil's destr provides the players with their first clue, pointing them towards the ancient ruins hidden deep within

Eldritch Forest T

The journey to the ruins is not straightforward. The Eldritch Forest is a living, breathing entity fil have to navigate through the forest's deceptive paths, facing challenges that test their survival skil mystical side of Eirandel, as they encounter spirits, ancient guardians, and possibly even the forest's

- @./The Whispering Glade.txt_The Whispering Glade_T_txt
- @./The Guardian of Paths.txt The Guardian of Paths T txt
- @./The Mirror Lake.txt The Mirror Lake T txt
- @./The Circle of Stones.txt The Circle of Stones T txt
- @./The Enchanted Grove.txt The Enchanted Grove T txt
- @./The Fading Echoes.txt_The Fading Echoes_T_txt
- @./The Final Guardian.txt_The Final Guardian: The Entwined Serpent_G_txt
- @./The Final Guardian Stat.txt_The Final Guardian: The Entwined Serpent Stat Block_S_txt
- @./The Final Guardian Loot.txt_The Final Guardian Loot Table_T_txt

The Ancient Ruins_T

Upon reaching the ruins, the players must breach its defenses, which include not only physical barrier filled with traps, puzzles, and enemies that guard the Sigil piece. This segment is designed to challe explore their characters' capabilities and how they can synergize as a group./n

The First Piece and Return_T

The climax of Act 1 occurs when the players face the final guardian of the Sigil piece. This encounter learned so far about combat in Eirandel. Upon securing the Sigil piece, the players must make their was a moment of realization that their journey is only just beginning. Archmage Eldrin, upon seeing the fi

a moment of realization that their journey is only just beginning. Archmage Eldrin, upon seeing the first piece recovered, refresaces the importance of their quest and hints at the dangers and wonders that await them as they seek out the remaining pieces./n

- @./The Final Guardian of Sigil Piece Stat.txt The Final Guardian of Sigil Piece Stat Block S txt
- @./The Final Guardian of Sigil Piece Loot.txt The Final Guardian of Sigil Piece Loot Table T txt

Act 1 Objectives T

- _Introduce players to the world of Eirandel and the main quest.
- _Foster team cohesion and allow players to explore their characters' abilities.
- _Challenge players with a mix of combat, exploration, and puzzle-solving.

Set the tone for the adventure, highlighting the significance of their mission and the mix of dark and light elements they will encounter.

THE SUMMONING



The adventure kicks off when the players are summoned by Archmage Eldrin, a respected and powerful wizard who has taken a keen interest in the prophesied reassembly of the Sigil. The Archmage explains the dire situation—that the balance of the world is at risk due to the Sigil's destruction and that its pieces must be recovered to prevent a cataclysmic event. Eldrin provides the players with their first clue,

reassembly of the Sigil. The ent a cataclysmic event. Eldrin ce of the Sigil./n

n and malevolent. Players will troduces the players to the er their progress./n

geon crawl segment of Act 1, bilities, encouraging them to

he players on everything they have eroes. However, the return is also

The Summoning T

@./The Summoning.png I png

The adventure kicks off when the players are summoned by Archmage Eldrin, a respec Archmage explains the dire situation—that the balance of the world is at risk due provides the players with their first clue, pointing them towards the ancient ruin

Eldritch Forest T

The journey to the ruins is not straightforward. The Eldritch Forest is a living, have to navigate through the forest's deceptive paths, facing challenges that test mystical side of Eirandel, as they encounter spirits, ancient guardians, and possi

- @./The Whispering Glade.txt The Whispering Glade T txt
- @./The Guardian of Paths.txt The Guardian of Paths T txt
- @./The Mirror Lake.txt The Mirror Lake T txt
- @./The Circle of Stones.txt The Circle of Stones T txt
- @./The Enchanted Grove.txt The Enchanted Grove T txt
- @./The Fading Echoes.txt_The Fading Echoes_T_txt
- @./The Final Guardian.txt The Final Guardian: The Entwined Serpent G txt
- @./The Final Guardian Stat.txt The Final Guardian: The Entwined Serpent Stat Block
- @./The Final Guardian Loot.txt The Final Guardian Loot Table T txt

The Ancient Ruins T

Upon reaching the ruins, the players must breach its defenses, which include not o filled with traps, puzzles, and enemies that guard the Sigil piece. This segment i explore their characters' capabilities and how they can synergize as a group./n

The First Piece and Return T

The climax of Act 1 occurs when the players face the final guardian of the Sigil p learned so far about combat in Eirandel. Upon securing the Sigil piece, the player a moment of realization that their journey is only just beginning. Archmage Eldrin and wonders that await them as they seek out the remaining pieces./n

- @./The Final Guardian of Sigil Piece Stat.txt The Final Guardian of Sigil Piece St @./The Final Guardian of Sigil Piece Loot.txt The Final Guardian of Sigil Piece Lo
- Act 1 Objectives T
- Introduce players to the world of Eirandel and the main quest.
- Foster team cohesion and allow players to explore their characters' abilities.
- Challenge players with a mix of combat, exploration, and puzzle-solving. Set the tone for the adventure, highlighting the significance of their mission and the mix of dark and light elements they will encounter.

THE WHISPERING GLADE Avoid minor natural Intelligence Basic guidance - A (Nature) DC 10 general direction to hazards like quicksand proceed that avoids or thorny thickets. natural hazards. Intelligence Detailed guidance -Discover a cache with a Identifies a safe path healing potion or herbs (Nature) DC 15 and hints at a nearby that can be used to hidden cache of craft a potion. supplies or a natural remedy. Wisdom Sense of unease -Advantage on the next (Perception) DC Detects a malevolent saving throw against a presence or a natural natural hazard 12 hazard nearby. encountered within the forest. "Gain a temporary +1 Wisdom "Clear warning -(Perception) DC Identifies the source bonus to AC for the of danger, allowing 18 next encounter within the party to prepare the forest, as the party or avoid it." is prepared for the danger." "The guide ensures Charisma Temporary alliance safe passage, reducing (Persuasion) DC A forest spirit 10 (to the forest appears and offers to travel time by half for a spirits) guide the party certain distance." through a small part of the forest. Insightful revelation - "Party gains a one-time Charisma (Persuasion) DC A forest spirit reveals use of a spell-like effect 15 (to the forest a significant shortcut (e.g., Pass without or the location of a Trace) without spirits) beneficial magical expending a spell slot."

THE GUARDIAN OF PATHS

Challenge Effect Result "Present a token Thalor allows safe Party members of nature's beauty passage and blesses receive a temporary (e.g., a rare flower the party with a hint of boon such as or a beautifully nature's favor. advantage on saving crafted wooden throws against natural hazards for item)." the next 24 hours. Plant a new tree "The party's actions The restored area heal a small part of the contains a natural or restore a forest, earning Thalor's remedy or herb that damaged area of the forest. respect and a safe can be used in passage." potion making. "The shared story or Thalor's wisdom Share a story or song that song pleases Thalor, grants the party an who offers wisdom insight into avoiding celebrates the about the forest or a or dealing with a spirit of nature. clue to the ruins." specific upcoming challenge. "Thalor gifts the "Engage in "Demonstrating strength and respect party a token of combat to prove nature's strength, strength and for nature convinces resolve (non-Thalor of the party's which can be used lethal, respect for worth, granting to summon a minor nature's passage." nature ally in a future encounter." balance)." Solve Thalor's Correctly solving the Discovering the riddle reveals a secret secret path or cache riddle about the path or a hidden cache shortens the journey forest and its inhabitants. of supplies. or equips the party with useful supplies.

THE ENCHANTED GROVE

Action	Result	Effect
Identify and	"The grove begins to	Party can harvest
neutralize the	heal immediately,	herbs for potent
source of the	rewarding the party	healing potions or
poison.	with the growth of rare herbs."	to sell for a high price.
Use magic to	"The magical	Each party member
cleanse the area of	cleansing purifies the	gains resistance to
its corruption.	grove, and the party is	a type of damage
	blessed with a	for 24 hours.
	protective aura."	
Convince local	"The animals help	Discover a shortcut
wildlife to assist in	spread the cure	that saves time or
the grove's	through the grove,	a treasure chest
restoration.	revealing a hidden path or treasure."	with magical items.
Offer a significant	"The spirits accept the	Receive a
personal sacrifice	sacrifice, fully	significant magical
to the spirits of the	restoring the grove	item or a

and granting the party permanent

ability score.

THE FADING ECHOES

Action	Result	Effect
Listen and	The echoes reveal a	Gain insight into
interpret the	secret about the	solving a puzzle or
echoes to learn	forest or the ruins	avoiding a dangeror
from the past.	that aids the party.	encounter.
Perform a ritual	"The ritual	Boon could be a
to honor the	strengthens the	temporary increase
guardians and	party's bond with the	hit points or advant
their sacrifices.	forest, granting a	on skill checks.



THE FINAL GUARDIAN: THE ENTWINED SERPENT STAT BLOCK

Attribute	Value
AC (Armor Class)	18
Hit Points	150
Speed	"30 ft., swim 60 ft."
STR (Strength)	18 (+4)
DEX (Dexterity)	15 (+2)
CON (Constitution)	17 (+3)
 ding Dungann Mac	tors with the personner

providing Dungeon Masters with the necessary information to bring these elements to life.

HIGH PRIESTESS VARA



THE SHATTERED SIGIL



THE SHATTERED SIGIL VENTURE: OVERVIEW

hattered Sigil" is an epic adventure set in the I lands of Eirandel, where the fabric of reality interwoven with the arcane and the dich in history and magic faces an

interwoven with the arcane and the cich in history and magic, faces an edented threat that could unravel its ve. The Sigil of Creation, an ancient as e power, once kept the forces of charbalance. However, it has been shatts cattered across the land, each hidder by forces both natural and supernat

growth and discovery. As the adventurers venture forth, they will find themselves tested not just by their enemies, but by the very land they seek to save. They will have to forge alliances with unlikely allies, from reclusive mages to ancient spirits of the forest, and confront their deepest fears and desires. The path to reassembling the Sigil is fraught with challenges that will require not only strength and courage but wisdom and sacrifice. The fate of Eirandel hangs in the balance, and the heroes stand at the center of this cosmic struggle, their actions echoing through the annals of time, shaping the future of this enchanted realm.

How to Run the Adventure



THE VOID CULT'S FINAL STAND TABLE

As the adventurers approach the climax of their quest to reform the Sigil of Creation, the Void Cult, led by the enigmatic High Priestess Vara, makes a desperate and formidable stand to thwart their efforts. This guide will help you, the Dungeon Master (DM), to create an epic and memorable confrontation that challenges your players both tactically and narratively.

The final stand occurs as the adventurers navigate the Dimensional Pathway to the Altar of Creation, a realm where the fabric of reality is thin, and magic is potent. This battleground should feel surreal and unstable, with the cult's dark magic causing the environment to shift unpredictably.

Cult Forces:

- High Priestess Vara: A powerful spellcaster and the cult leader. Her abilities should challenge the party's strengths and exploit their weaknesses.
- Void Cultists: A mix of melee fighters, spellcasters, and assassins. They are fanatically loyal and willing to sacrifice themselves for their cause.
- Dark Entities: Summoned creatures or constructs bound by the cult's magic, adding a wild card element to the battle./n

Encounter Design:



HIGH PRIE	STESS VARA ST	AT BLOCK
Attribute		Value
AC (Armor Class)	17 (mage armor)	
Hit Points	142	
Speed	30 ft.	
STR (Strength)	10 (+0)	
DEV	24 (-2)	

AC: 12 (15 with Mage Armor)
Hit Points: 98



THE MARKET SQUARE



The heart of Stonebridge, filled with various stalls selling food, clothing, tools, and trinkets. The square also serves as a place for villagers to gather and share news

THE VILLAGE SHRINE

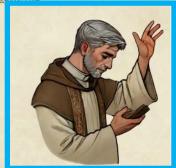
A small shrine dedicated to the gods of protection and harvest. It is tended by Father Eamon, a kind-hearted priest.

- Old Maud: An elderly woman who claims to have seen spirits near Blackstone Keep.
- Father Eamon: A gentle priest who provides healing and spiritual guidance to the villagers.





FATHER EAMON



VILLAGE RUMORS AND QUESTS

Upon arrival, the adventurers can gather information by interacting with villagers. Use the following table to determine what rumors they hear:

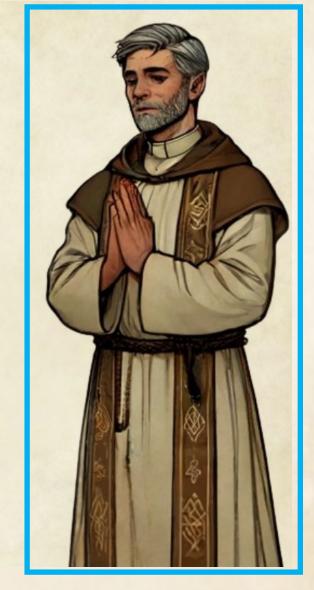
Rumor

- Basic guidance A general direction to proceed that avoids natural hazards.
- 2 Detailed guidance Identifies a safe path and hints at a nearby hidden cache of supplies or a natural remedy.
- 3 Sense of unease Detects a malevolent presence or a natural hazard nearby.
- 4 "Clear warning Identifies the source of danger, allowing the party to prepare or avoid it."
- 5 Temporary alliance A forest spirit appears and offers to guide the party through a small part of the forest.
- 6 Insightful revelation A forest spirit reveals a significant shortcut or the location of a beneficial magical anomaly.

INITIAL ENCOUNTERS

As the heroes explore Stonebridge, they may encounter the following scenarios:

Quest	Setup	Details
The	A distressed mother, Eliza, approaches	Tom disappeared while playing near the edge of the forest. The party can investigate the area, finding signs of a struggle and strange tracks leading into the vale.
The Desecrated Shrine	Father Eamon requests the adventurers' help in cleansing the shrine after it was desecrated.	The party must confront a minor undead threat (such as skeletons or zombies) and perform a ritual to restore the shrine's sanctity.









 Garrick, the Guard Captain: A stern but fair leader, concerned about the recent events.

THE MARKET SQUARE



The heart of Stonebridge, filled with various stalls selling food, clothing, tools, and trinkets. The square





OLD MAUD NATURALCRIT

APPENDIX

This appendix provides detailed descriptions for key characters featured in "The Shadows of Blackstone Keep" adventure. These can be used by dungeon masters to enhance the lore.

OLD MAUD



Old Maud is a frail yet resilient elderly woman who has lived in Stonebridge for as long as anyone can remember. With her long, gray hair tied back in a loose braid and her deeply wrinkled skin, she looks every bit the part of a village elder. Dressed in simple, tattered clothes, she moves slowly, often leaning on a sturdy wooden staff adorned with charms and trinkets. Despite her age and appearance, Old Maud possesses a sharp

reminder of what was lost to the curse of Blackstone Keep.



HILDA THORNBREW'S SHOP

Hilda Thornbrew, the friendly and knowledgeable innkeeper of The Black Boar Inn, also runs a small shop within the inn where adventurers can find various supplies, potions, and items to aid them in their quests. Known for her warm hospitality and fair prices, Hilda ensures that every adventurer is well-equipped before they venture into the dangers of Blackstone Vale.

HILDA THORNBREW'S SHOP:

As you enter the cozy common room of The Black Boar Inn, you notice a small corner dedicated to an array of items neatly displayed on wooden shelves. The warm glow of the fireplace illuminates the various potions, weapons, and adventuring gear available for purchase. Hilda Thornbrew, the stout and friendly innkeeper, stands behind the counter, ready to assist you with a welcoming smile.

Item	Description	Price
Healing Potion	Restores 2d4 + 2 hit points when consumed.	50 gp
Antidote	Cures poison and grants advantage on saving throws against poison for 1 hour.	25 gp
Torch	Burns for 1 hour, providing bright	1 cp



FATHER EAMON



Father Eamon is the gentle and kind-hearted priest of Stonebridge, devoted to providing spiritual guidance and solace to the villagers. He is a man of average height with a serene demeanor, short gray hair, and soft, compassionate eyes that seem to reflect the weight of his wisdom. Dressed in simple priestly robes and always carrying a holy symbol around his neck, Father Eamon is often seen tending to the small village shrine, performing blessings, and offering prayers for the wellbeing of the community. His voice is calm and reassuring, capable of soothing even the most troubled souls. Despite the darkness that has befallen the vale, Father Eamon remains a beacon of hope and faith, tirelessly working to protect and heal those in need. His deep connection to the divine grants him the ability to

people and will go to great lengths to protect the often coordinating with adventurers to safegua village from the encroaching darkness of Black Vale.



Despite his tough exterior, Garrick cares deep

SETTING THE SCENE

disappearances.



proposed by a coom and an amount and any order out

THE THRONE ROOM

The final confrontation site, a grand hall with a dilapidated throne at the far end.

KEY FEATURES:

- · The Betrayer: The corrupted former ally of Lord Blackstone, now a powerful antagonist.
- · Malevolent Spirit: The ultimate foe, whose defeat is necessary to lift the curse.

ENCOUNTERS AND CHALLENGES

The party will face various encounters and challenges as they navigate the vale. Use the following table to determine random encounters:

- 1 A group of animated suits of armor attacks the party.
- 2 Ghostly figures appear, offering cryptic warnings or attacking if provoked.
- 3 The party triggers a magical trap, releasing spectral guardians.
- 4 Cursed books fly off the shelves in the library, casting harmful spells.
- 5 A shadowy figure ambushes the party, attempting to sow fear and confusion.
- 6 The Betrayer himself appears, taunting the party and testing

NOTABLE QUESTS AND OBJECTIVES

_		
Quest	Objective	Details
Cleanse the	Disrupt the dark	The party must defeat dark
Forsaken	rituals and cleanse the	druids and purify the ritual
Chapel	chapel of its	altar to weaken the
	malevolent influence.	malevolent spirit's power.
Retrieve	Find an ancient tome	The party must navigate
the Ancient	in the Haunted Library	magical traps and cursed
Tome	that contains the ritual	books to retrieve the tome
	to banish the	with Eldrin the Sage's
	malevolent spirit.	assistance.
Confront	Defeat or redeem the	The party must face the
The	Betrayer, who has	Betrayer in the throne
Betrayer	become a powerful servant of the	room, making a choice to defeat him in combat or
	malevolent spirit.	find a way to redeem his corrupted soul.

CONCLUSION OF ACT III

By the end of Act III, the adventurers will have faced their greatest challenges and unraveled the deepest mysteries of Blackstone Keep. Their final victory over the malevolent spirit will lift the curse, restore peace to Blackstone Vale, and solidify their status as true heroes. The adventure concludes with the party emerging from the keep, triumphant and forever changed by their harrowing journey.

FINAL BOSS

The Betraver

Armor Class 18 (Plate Armor) Hit Points 150 (20d8+60) Speed 30ft.

	EX CO	N INT	WIS	CHA
18 (+4) 12 (+1) 16 (+	3) 14 (+:	2) 10 (+0	16 (+3)

Saving Throws Str +8, Con +7, Wis +4, Cha +7 Skills Athletics +8, Intimidation +7, Perception +4 Damage Resistances necrotic; bludgeoning, piercing, and slashing from nonmagical attacks Senses darkvision 60 ft., passive Perception 14

Languages Common, Infernal Challenge 10 (5,900 XP)

Legendary Resistance (3/Day). If the Betrayer fails a saving throw, he can choose to succeed instead

Dark Aura. A 15-foot radius of darkness surrounds the Betrayer. Creatures that start their turn in this aura take 7 (2d6) necrotic

Magic Resistance. The Betrayer has advantage on saving throws against spells and other magical effects.

ACTIONS

Multiattack. The Betrayer makes three attacks: two with his longsword and one with his necrotic blast.

Longsword. Melee Weapon Attack: +8 to hit, reach 5ft., one target. Hit 10 (1d8 + 4) slashing damage, or 11 (1d10 + 4) slashing damage if used with two hands, plus 7 (2d6) necrotic

Necrotic Blast. Ranged Spell Attack: +7 to hit, range 60 ft., one target. Hit 18 (4d8) necrotic damage.

Frightful Presence. Each creature of the Betrayer's choice within 60 feet of him and aware of him must succeed on a DC 15 Wisdom saving throw or be frightened for 1 minute. A frightened target can repeat the saving throw at the end of each of its turns, with disadvantage if the Betrayer is within line of sight, ending the effect on itself on a success.

LEGENDARY ACTIONS

The Betrayer can take 3 legendary actions, choosing from the options below. Only one legendary action option can be used at a time and only at the end of another creature's turn. The Betrayer regains spent legendary actions at the start of his turn.

Attack. The Betrayer makes a longsword attack.

Move. The Betrayer moves up to his speed without provoking opportunity attacks.

Dark Summon (Costs 3 Actions). The Betrayer summons 1d4 shadowy figures to unoccupied spaces he can see within 30 feet. The summoned shadowy figures act on their own initiative count and disappear after 1 minute.



Hilda Thornbrew: A stout and friendly woman with a wealth of local knowledge.



Pipeline Specifications



- Timing
- Zero shot character generation pipeline (in RTX 4090 | 24GB VRAM)
 - 10-20 minutes per pipeline (generation, processing, segmenting, captioning and clustering) * 6 => 1 to 2 hours
 - 15-30 minutes per LoRA
 - 1-2 seconds for generation
- ~1-4 hours for content generation
- 5-6 minutes for map generation

Conclusion



- Consistent image generation pipeline
 - Generating character sheets
 - Training a character specific LoRA
 - Saving to Context DB
 - Evaluation
- D&D based map generation
 - Map generation with cellular automata
 - Map to image
 - Using Maps to retrieve context
- End to end pipeline for d&d adventure generation
 - Multi agent content generation
 - RAG
 - Function calls with pre made scripts
- Published 2 papers (RANLP 2023, ICCART 2023, 3rd paper)
- https://github.com/gayashan2020/consistantChar.git

Reference



- [1] E. Papagiannopoulou and G. Tsoumakas, "A review of keyphrase extraction," Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery, vol. 10, no. 2, p. e1339, 2020.
- [2] S. Vijayarani, M. J. Ilamathi, M. Nithya et al., "Preprocessing techniques for text mining-an overview," International Journal of Computer Science & Communication Networks, vol. 5, no. 1, pp. 7–16, 2015.
- [3] R. Srivastava, P. Singh, K. Rana, and V. Kumar, "A topic modeled unsupervised approach to single document extractive text summarization," Knowledge-Based Systems, vol. 246, p. 108636, 2022.
- [4] A. P. B. Veyseh, N. Meister, F. Dernoncourt, and T. H. Nguyen, "Improving keyphrase extraction with data augmentation and information filtering," arXiv preprint arXiv:2209.04951, 2022
- [5] S. Beliga, "Keyword extraction: a review of methods and approaches," University of Rijeka, Department of Informatics, Rijeka, vol. 1, no. 9, 2014.
- [6] J. Oppenlaender, "Prompt engineering for text-based generative art," arXiv preprint arXiv:2204.13988, 2022.
- [7] S. W. McRoy, S. Channarukul, and S. S. Ali, "An augmented template-based approach to text realization," Natural Language Engineering, vol. 9, no. 4, pp.381–420, 2003.
- [8] A. Radford, K. Narasimhan, T. Salimans, I. Sutskever et al., "Improving language understanding by generative pre-training," 2018.
- [9] S. Michelmann, M. Kumar, K. A. Norman, and M. Toneva, "Large language models can segment narrative events similarly to humans," Jan 2023. [Online]. Available: https://arxiv.org/abs/2301.10297
- [10] Y. Lu, S. Wu, Y.-W. Tai and C.-K. Tang, "Image Generation from Sketch Constraint Using Contextual GAN," in Computer Vision ECCV 2018, Lecture Notes in Computer Science, vol 11213. Springer, Cham, 2018
- [11] AutoEncoders are Essential in Deep Neural Nets. Towards Data Science. Retrieved March 21, 2023, from https://towardsdatascience.com/autoencoders-are-essential-in-deep-neural-nets-f0365b2d1d7c
- [12] J. Ho, C. Saharia, W. Chan, D. J. Fleet, M. Norouzi, and T. Salimans, "Cascaded diffusion models for high fidelity image generation." J. Mach. Learn. Res., vol. 23, no. 47, pp. 1–33, 2022.
- [13] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2019, pp. 4401–4410.
- [14] J. Ho, C. Saharia, W. Chan, D. J. Fleet, M. Norouzi, and T. Salimans, "Cascaded diffusion models for high fidelity image generation." J. Mach. Learn. Res., vol. 23, no. 47, pp. 1–33, 2022.
- [15] S. Behdenna, F. Barigou, and G. Belalem, "Document level sentiment analysis: a survey," EAI Endorsed Transactions on Context-aware Systems and Applications, vol. 4, no. 13, pp. e2–e2, 2018.
- [16] X. Han and L. Wang, "A novel document-level relation extraction method based on bert and entity information," IEEE Access, vol. 8, pp. 96 912–96 919, 2020.
- [17] G. Kim, H. Shim, H. Kim, Y. Choi, J. Kim, and E. Yang, "Diffusion video autoencoders: Toward temporally consistent face video editing via disentangled video encoding," arXiv preprint arXiv:2212.02802, 2022.
- [18] B. Kim, I. Han, and J. C. Ye, "Diffusemorph: Unsupervised deformable image registration along continuous trajectory using diffusion models," arXiv preprint arXiv:2112.05149, 2021.36

Reference



- [19] "Corridor Crew. (2023, 03 07). VFX Reveal Before & After Anime Rock, Paper, Scissors [Video]. YouTube. https://www.youtube.com/watch?v=ljBSmQdL_Ow"
- [19] T. Zhang, V. Kishore, F. Wu, K. Q. Weinberger, and Y. Artzi, "Bertscore: Evaluating text generation with bert," arXiv preprint arXiv:1904.09675, 2019.
- [20] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," IEEE transactions on image processing, vol. 13, no. 4, pp. 600–612, 2004.
- [21] A. Cherian and A. Sullivan, "Sem-gan: Semantically-consistent image-to-image translation," in 2019 ieee winter conference on applications of computer vision (wacv). IEEE, 2019, pp. 1797–1806.
- [22] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805.
- [23] Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. 2020. Electra: Pre-training text encoders as discriminators rather than generators. In International Conference on Learning Representations.
- [24] Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzman, Edouard Grave, Myle Ott, Luke Zettle- ´ moyer, and Veselin Stoyanov. 2019. Unsupervised cross-lingual representation learning at scale. arXiv preprint arXiv:1911.02116.
- [25] Simone Tedeschi, Valentino Maiorca, Niccolo Campol-`ungo, Francesco Cecconi, and Roberto Navigli. 2021. WikiNEuRal: Combined neural and knowledgebased silver data creation for multilingual NER. In Findings of the Association for Computational Linguistics: EMNLP 2021, pages 2521–2533.
- [26] Jean Baptiste. 2022. roberta-large-ner-english: model fine-tuned from roberta-large for ner task. https://huggingface.co/Jean-Baptiste/ roberta-large-ner-english. (Accessed on 05/10/2023).
- [27] Fabio Souza, Rodrigo Nogueira, and Roberto Lotufo. 2019. Portuguese named entity recognition using bert-crf. arXiv preprint arXiv:1909.10649.
- [28] M. Honnibal and I. Montani, "spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing," unpublished.
- [29] M. Van Nguyen, V. D. Lai, A. P. B. Veyseh, and T. H. Nguyen, "Trankit: A light weight transformer-based toolkit for multilingual natural language processing," arXiv preprint arXiv:2101.03289, 2021.
- [31] A. Akbik, T. Bergmann, D. Blythe, K. Rasul, S. Schweter and R. Vollgraf, "FLAIR: An Easy-to-Use Framework for State-of-the-Art NLP," in Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: System Demonstrations, pages 54-59, Minneapolis, Minnesota, 20191.
- [32] Wizards of the Coast, (2023, March 28). LOST MINE OF PHANDELVER. D&D Beyond. [Online]. Available: https://www.dndbeyond.com/sources/lmop
- [33] Wolfgang Baur, Steve Winter, and Alexander Winter. 2014a. Hoard of the Dragon Queen. Wizards of the Coast.

Reference



- [34] Wolfgang Baur, Steve Winter, and Alexander Winter. 2014b. Rise of Tiamat. Wizards of the Coast.
- [35] Christopher Perkins, Tracy Hickman, and Laura Hickman. 2016. Curse of Strahd. Wizards of the Coast.
- [36] Christopher Perkins, Will Doyle, and Steve Winter. 2017. Tomb of Annihilation. Wizards of the Coast.
- [37] Christopher Perkins, Graeme Barber, Bill Benham, et al. 2021. Candlekeep Mysteries. Wizards of the Coast.
- [38] Stacey Allan, Will Doyle, Ari Levitch, and Christopher Perkins. 2021. The Wild Beyond the Witchlight. Wizards of the Coast.
- [39] T. Le Scao et al., "BLOOM: A 176B-Parameter Open-Access Multilingual Language Model," arXiv preprint arXiv:2211.05100, 20221.
- [40] Databricks. 2023. Dolly: The first open source, commercially viable instruction-tuned Ilm. Databricks Blog.
- [41] Xinyang Geng and Hao Liu. 2023. Openllama: An open reproduction of llama.
- [42] Weerasundara, G., & de Silva, N. (2023). Comparative analysis of named entity recognition in the dungeons and dragons domain. arXiv preprint arXiv:2309.17171...
- [43] Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Om- mer, B. (2022). High-resolution image synthesis with la- tent diffusion models. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 10684–10695
- [44] Betker, J., Goh, G., Jing, L., Brooks, T., Wang, J., Li, L., ... & Ramesh, A. (2023). Improving image generation with better captions. Computer Science. https://cdn. openai. com/papers/dall-e-3. pdf, 2(3), 8.
- [45] Q. Wu, G. Bansal, J. Zhang, Y. Wu, S. Zhang, E. Zhu, B. Li, L. Jiang, X. Zhang, and C. Wang, "Autogen: Enabling next-gen Ilm applications via multi-agent conversation framework," arXiv preprint arXiv:2308.08155, 2023.
- [46] Hong, S., Zheng, X., Chen, J., Cheng, Y., Wang, J., Zhang, C., ... & Wu, C. (2023). Metagpt: Meta programming for multi-agent collaborative framework. arXiv preprint arXiv:2308.00352.
- [47] T. Brooks, A. Holynski, and A. A. Efros, "Instructpix2pix: Learning to follow image editing instructions," arXiv preprint arXiv:2211.09800, 2022. 97
- [48] Aisuko (2024) Lora from scratch, Kaggle. Available at: https://www.kaggle.com/code/aisuko/lora-from-scratch (Accessed: 10 June 2024).
- [49] Weerasundara, G., & de Silva, N. (2024, January). A Multi-Stage Approach to Image Consistency in Zero-Shot Character Art Generation for the D&D Domain. In ICAART (3) (pp. 235-242).
- [50] Hussain, M. (2023). YOLO-v1 to YOLO-v8, the rise of YOLO and its complementary nature toward digital manufacturing and industrial defect detection. Machines, 11(7), 677.

Thank You

