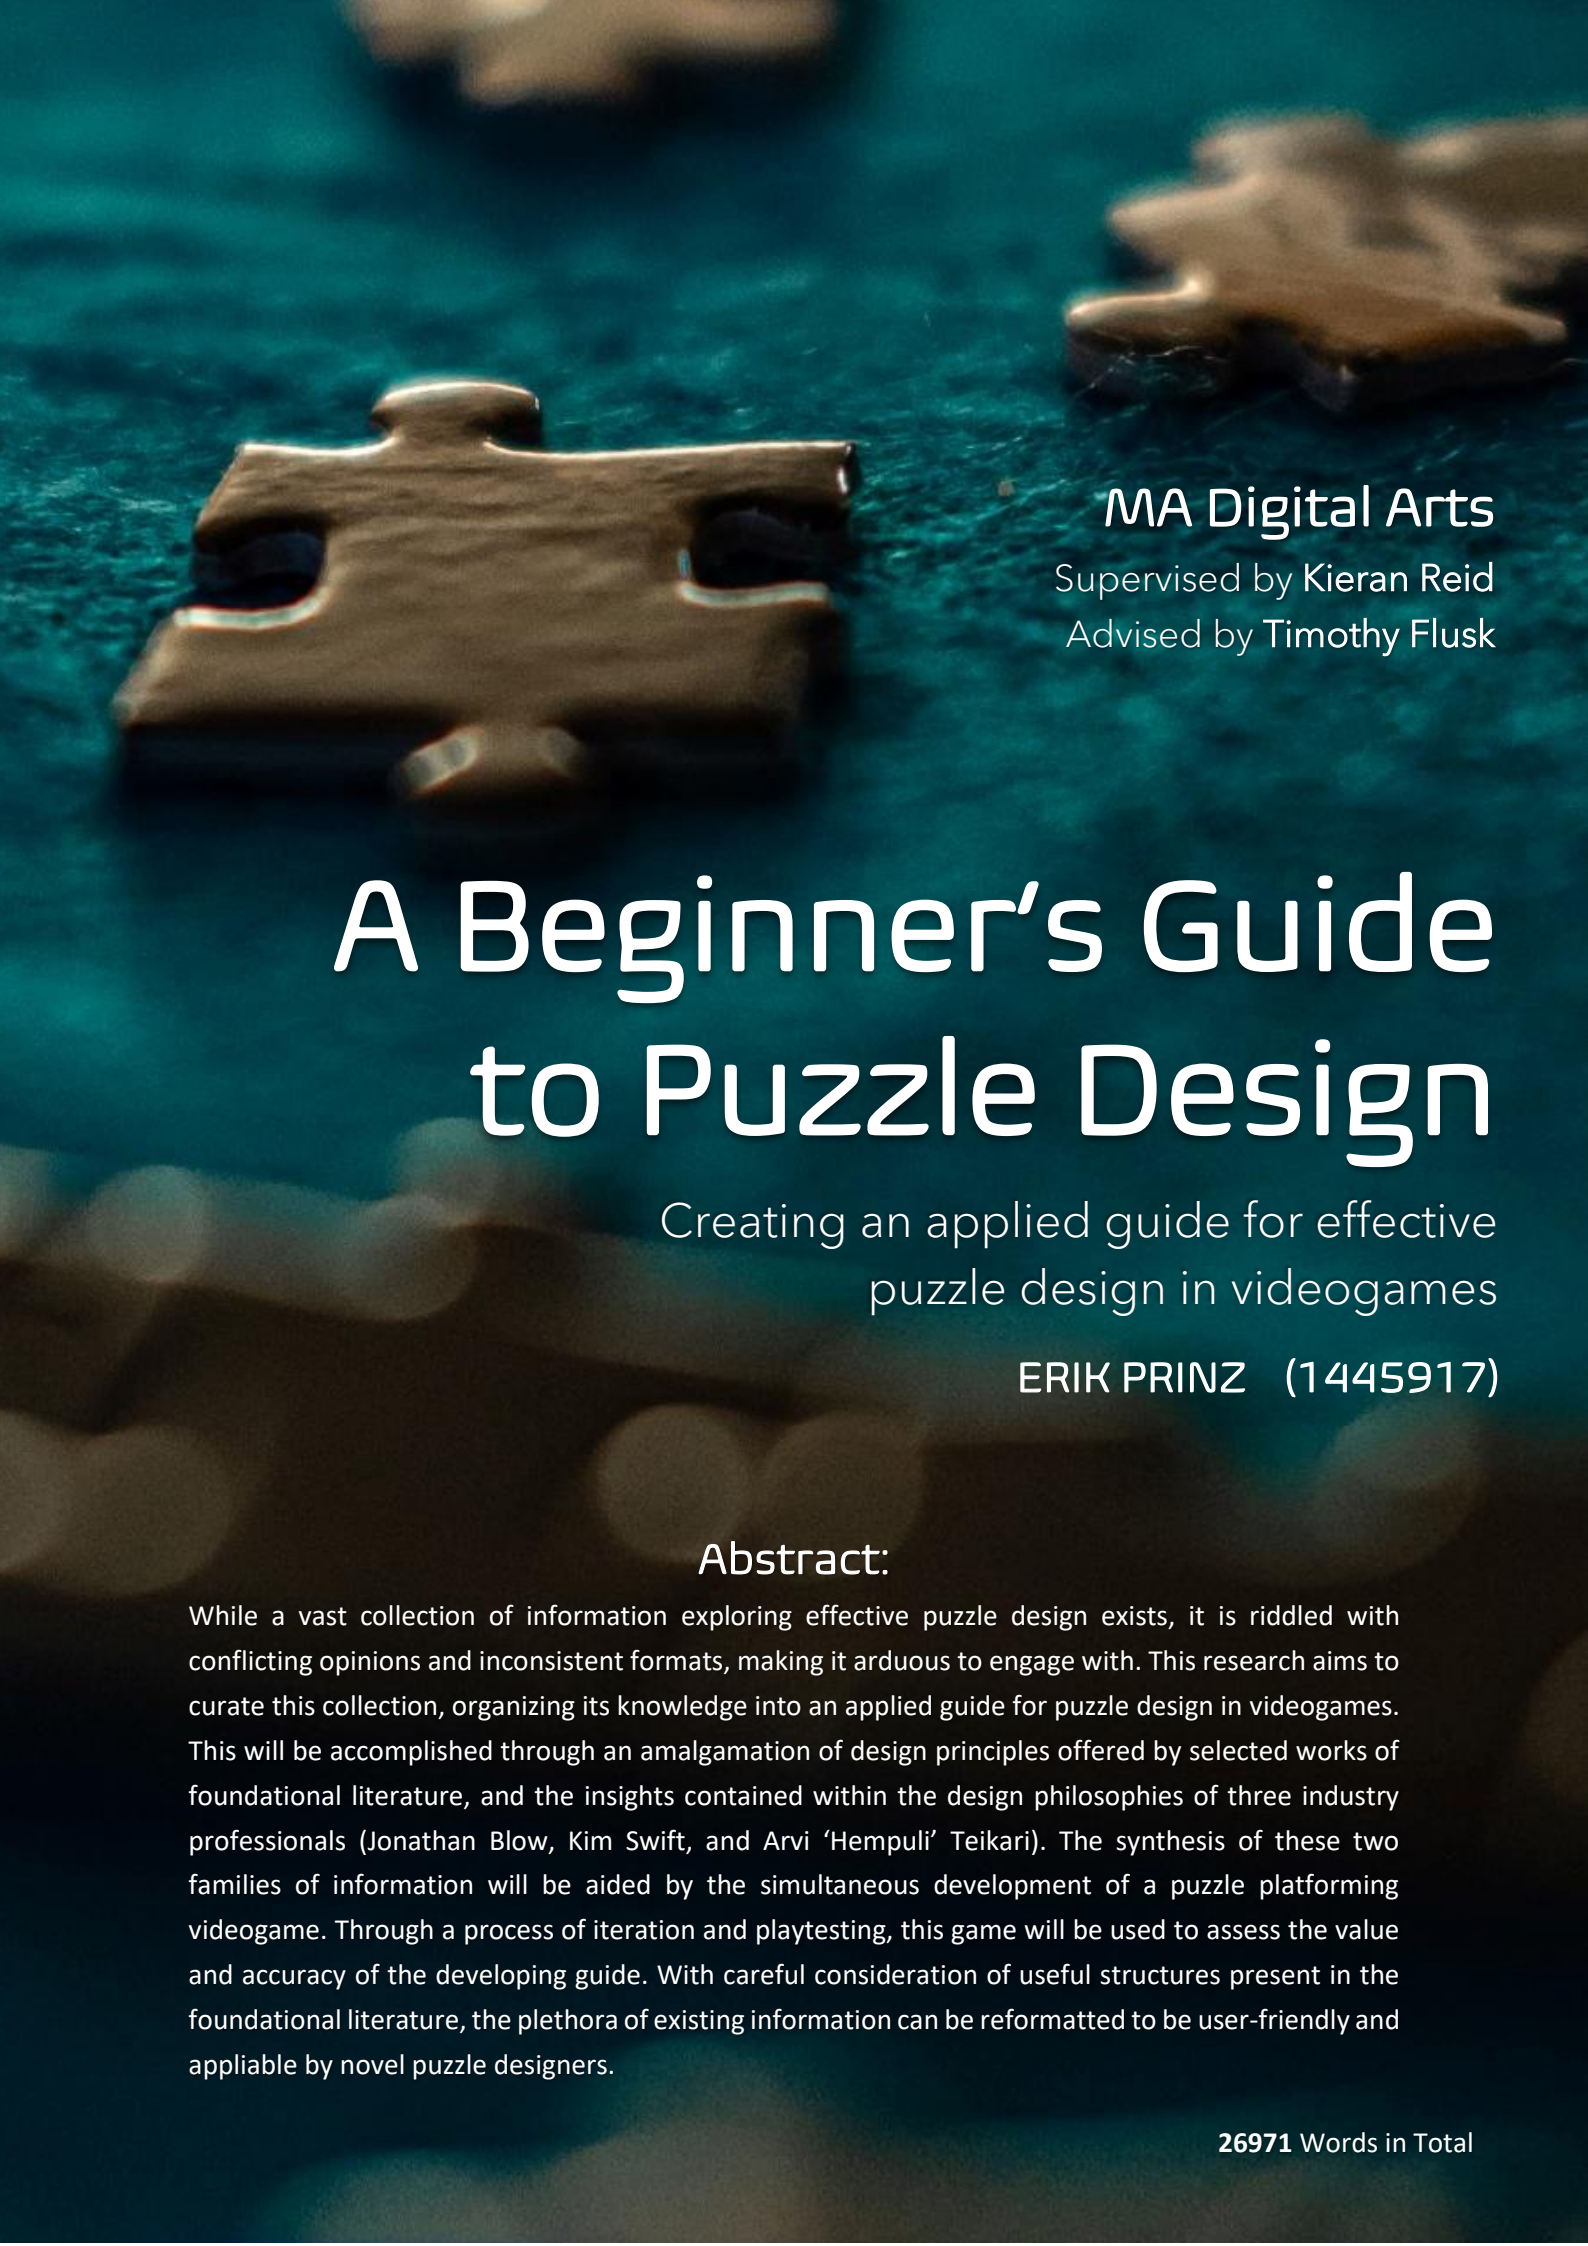

MA Digital Arts
Supervised by Kieran Reid
Advised by Timothy Flusk

A Beginner's Guide to Puzzle Design

Creating an applied guide for effective puzzle design in videogames

ERIK PRINZ (1445917)

Abstract:

While a vast collection of information exploring effective puzzle design exists, it is riddled with conflicting opinions and inconsistent formats, making it arduous to engage with. This research aims to curate this collection, organizing its knowledge into an applied guide for puzzle design in videogames. This will be accomplished through an amalgamation of design principles offered by selected works of foundational literature, and the insights contained within the design philosophies of three industry professionals (Jonathan Blow, Kim Swift, and Arvi 'Hempuli' Teikari). The synthesis of these two families of information will be aided by the simultaneous development of a puzzle platforming videogame. Through a process of iteration and playtesting, this game will be used to assess the value and accuracy of the developing guide. With careful consideration of useful structures present in the foundational literature, the plethora of existing information can be reformatted to be user-friendly and appliable by novel puzzle designers.

26971 Words in Total

Contents

Cover Image: Photo by Aaron Chavez on Unsplash

The Hors D'oeuvre

A Note to My Examiner

Thank you for taking the time to engage with and assess this submission.

Before we delve straight into creating a *digestible* way to engage with puzzle design, I felt it would be valuable to outline the exact contents of this submission, as well as the most efficient way to engage with each individual portion of it. This way, we can avoid any unnecessary repetition in the marking process, hopefully saving you as much time and effort as possible.

This submission is made up of the following individual products:

- This Dissertation (PDF)
- *In Search of Eureka – The Beginner's Guide to Puzzle Design in Videogames* (PDF)
- *Allegory* (Executable Game, in a Zip Folder)
- A Puzzle Design Playtest (Zip Folder)
  - Test A (Executable Experience)
  - Test B (Executable Experience)
  - Playtesting Questionnaire (PDF)
  - Consent Forms (Zip Folder)

With only the exception of the Puzzle Design Playtest, each of these individual components is imperative to the assessment of this project. Engaging with each of them in tandem with this Dissertation is by far the most effective way to process this project. The diagram below shows a timeline of when it would be most effective to engage with each of these components relative to the chapters of this Dissertation.

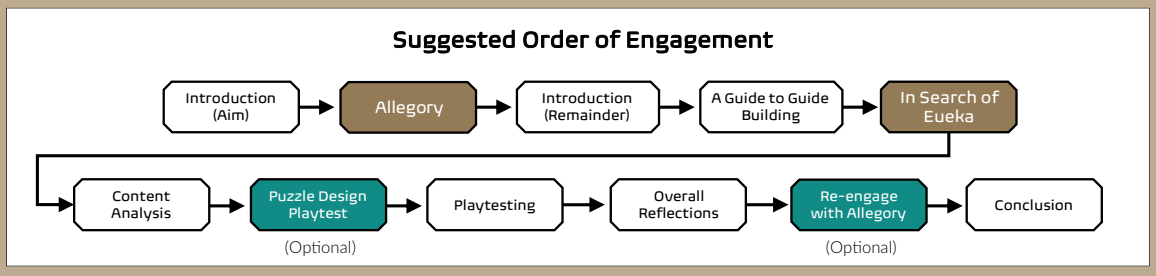

Figure 1: Suggested Order of Engagement for Various Submission Portions

This order of engagement will ensure that you always have the necessary context to fully understand the various chapters of this document, without needing to engage repeatedly with the supporting products.

Do not worry about remembering this order right now, though – I will continually remind you throughout this document when the best moments to switch focus are. Keep an eye out for text boxes like this one, containing the following symbol. These annotations are specifically designed to improve your marking experience.

Thank you again for your time and effort in marking this submission. I hope that you enjoy your time with it, and that it proves to be a stimulating and valuable experience.

The Appetizer

Introduction

The different media that give life to puzzles are vast in number, from the puzzle games played by masons in ancient Egypt to their modern counterparts seen in videogames today. Puzzles have captivated people for centuries, yet we still do not know why we are intrinsically drawn to them (Danesi, 2002). While we may not have an answer for the origins of this primal urge to seek out solutions, we have continually improved our understanding of what makes the solving process so engaging.

Captivating hooks used to catch the eye (Kim, as cited by Adams, 2015; Kim, as cited by Bates, 2004; Swift, 2007), toy-like interfaces that beg one to play with them (Rigney, 2012; Schell, 2015), tricky shifts in perception that catch one off guard (Schell, 2015; Selinker and Snyder, 2018), expertly designed difficulty curves to keep one on the edge of their seat (A Maze., 2014; Kim, as cited by Adams, 2015; Kim, as cited by Bates, 2004; Menzel, 2016; Parkin, 2008), all of which culminating in a fantastic "Eureka!" (Adams, 2015; Bates, 2015; Danesi, 2002; Schell, 2015; Selinker and Snyder, 2018). A veritable library of clever techniques to improve the solving experience already exists.

But this library is spread thin and far-between, hidden in the pages of lengthy tomes and the philosophies of the field's experts. For one inexperienced in puzzle design, it is no easy feat to engage with this information. As such, this research seeks to collect the most substantial advice from this library and format it such that designers of all skill levels can easily access and apply it.

Aim

This research aims to create a functional guide for creating puzzles specifically for videogames. The development of this guide will be iterative and done hand in hand with the creation of a puzzle videogame, to enable the testing of its usability and accuracy. The guide's steps will be formulated by curating information presented in foundational literature on puzzle design (Adams, 2015; Bates, 2004; Kim, 2006; Schell, 2015; Selinker and Snyder, 2018) along with the wisdom contained in the design philosophies of a group of acclaimed videogame puzzle designers (Blow, Teikari and Swift).

This guide will take the form of a clear process to follow when designing puzzles for videogames. A comparative study of the *goals* and *design philosophies* of the aforementioned puzzle designers will hopefully lead to useful conclusions about which portions of their philosophies are most vital to their success, and whether said portions should have a place in the curated guide.

The works and ideas of these designers will not, however, be considered in isolation. Their views will be contrasted with existing literature in the field of puzzle design. In this way, this research can highlight how these modern designers subvert or align with the more classic conventions laid out in existing literature. Understanding these relationships will not only allow for further curation of the insights gained from the selected puzzle designers, but will also highlight the classic conventions that still deserve a place in this contemporary guide.

While this theoretical analysis will form the foundation of the final guide, the finished product will be completed through a process of continual testing and iteration. A puzzle videogame will be simultaneously created, essentially existing as a testing ground for each iterative set of design steps. This provides a means of discerning the effectiveness of various philosophies and ideas as well as the usability of the guidelines being derived from them. Additionally, it offers a space for me to form my own observations and insights on the puzzle design process. At the end of this undertaking, the guide should be tangibly usable in practice rather than just another postulation about what makes a good puzzle.

Research Objectives

A. Curate information from several foundational works on puzzle design.
B. Curate the wisdom offered by three selected puzzle designers.
C. Synthesise the information gained from foundational works and dynamic informal sources into a single, concise collection of knowledge.
D. Develop a guide that is user-friendly and easily appliable by novice designers.
E. Effectively test the guide through the development of a puzzle videogame.
F. Improve my skills as a puzzle designer.

Research Questions

A. How feasible is it to define a universal methodology for puzzle design?

Before continuing, please play the entire ***Allegory*** experience.

Rationale

Now that we have outlined what the goals of this research are, I hope to illustrate why I believe they are obtainable, as well as why the research's outcomes are important. By first looking at how this research *can* be done, we can convey a deeper understanding of why it *should* be done later in this section.

An overwhelming sea of information about puzzle design already exists – the water is just somewhat choppy right now. While these waters may seem daunting to an inexperienced sailor, they offer a bountiful fishing ground for the seasoned academic angler. By this, I mean that there are a multitude of information sources carrying valuable knowledge about puzzle design that I can access.

These sources range from reference works aimed at teaching the reader to design puzzles[1], to academic works that address specific questions about puzzle design[2], to more informal sources which either engage with the opinions of industry professionals or summarise and curate existing knowledge about puzzle design[3]. To give our search more structure, I have sorted the sources available into two categories, namely *Foundational Literature* and *Dynamic Informal* work.

Foundational literature refers to existing academic and reference works about puzzle design. They generally go into sufficient depth when presenting theoretical content and offer valuable structures and formats that can be used as skeletons for the final guide. On the other hand, *Dynamic Informal* work is more subjective and less regulated. This means it is abundantly available and can immediately engage with new discoveries, making it the perfect sphere from which I can extract the philosophies of my selected puzzle designers, as it expresses the lessons learnt from their latest experiences. A more detailed discussion of these concepts can be found in *The Guest List* section below.

In addition to the myriad of information at my disposal, I have a unique means of testing this knowledge. A creative practice piece in the form of a puzzle platforming videogame will be created simultaneously with the guide to puzzle design. By using each iteration of the guide for actual puzzle development, I can investigate its usefulness in a practical sense rather than just a theoretical one – ensuring that it is applicable by new designers. Furthermore, feedback about these puzzles can then be collected from playtesters, allowing me to assess the value of the iteration used to create them[4].

With the abundance of relevant work available, and my unique, practical means of assessing said work, the research goals presented earlier should be obtainable. But just because something *can* be done does not mean it *should* be done, leaving the question, "Why do we *need* an applied guide to puzzle design for videogames?"

Well, put simply, puzzle design is daunting. When a new designer enters the field, they are bombarded with theories and opinions from *every* direction, leaving them no identifiable starting point or direction to move in. This is the result of three prominent issues with how the current knowledge in the sphere is framed.

Firstly, much of the foundational literature offers recommendations that are only relevant at very specific points in the development process, without directly addressing *when* one should consider them. This results in a lot of content being inapplicable to someone who has yet to start developing

[1] The most prominent examples being the works of Adams (2015), Bates (2004), Schell (2015) and Selinker and Snyder (2018).
[2] Examples include the works of Asbell-Clarke et al. (2016), Danesi (2002) and Linehan et al. (2014).
[3] Such as the works of Alexander (2014), Blow et al. (2013), Game Maker's Toolkit (2016, 2018), Kim (2006), Menzel (2019), Parkin (2008), Teikari (2019), The Game Overanalyser (2019), and Writing on Games (2016).
[4] See *Methodology* for more details.

their puzzles, as well as confusion when readers attempt to integrate advanced concepts into fledgling creations. The works of Bates (2004) and Schell (2015) both suffer from these issues to some extent.

That said, some works do a great job of integrating chronology into their content. But these often suffer from the second prominent framing issue: they are vague. These works will often include steps that are broadly described without specific details. While they do give the reader clearer direction, this improved digestion only lasts until the reader reaches one of these vague instructions. Kim's Eight Steps (as described by Adams (2015)) and *Puzzlecraft*'s chapter on crafting videogame puzzles (Selinker and Snyder, 2018, p. 202) are both examples of this phenomenon.

The final prominent issue is the amount of puzzle design knowledge that is hidden beneath a mountain of anecdotal irrelevancies – a problem most notable in the *Dynamic Informal* work. Extracting the valuable information that is actually worth implementing requires considerable amounts of sorting and cross-referencing. This is clearly not ideal for someone just entering the field.

With all these shortcomings of existing literature in mind, clearly there is room for a concise and complete guide to puzzle creation. By amalgamating the most interesting and useful information from the more comprehensive sources and presenting it in a format designed with chronology, priority and direction in mind, the final product should provide a far more accessible entry point into the ocean of puzzle design.

Methodology

The heart of this project lies in reformatting the plethora of existing puzzle design information to be *appliable* and *user-friendly*. This applied, descriptive research aims to produce a set of distinct instructions that take a designer from having nothing to their first completed puzzles. I specifically identify this research as descriptive, as it will only focus on identifying and organising the puzzle design ideologies that have a positive effect on the quality of puzzles produced through them - it will not go into great depth discussing *how* these philosophies achieve this goal.

As it is difficult to put forth a clear hypothesis or theory before the groundwork has been done, this research will take an inductive approach, collating all the most promising information available and using it to create something that can be effectively tested in future. Additionally, it will follow an interpretivist research philosophy, as I acknowledge that there will never be a single, correct way of creating a perfect puzzle. Each puzzle is different depending on the context in which it is created, and so even some of the valuable advice that makes its way into the final guide may prove to be irrelevant given the subjective needs of the user at the time.

Three main forms of qualitative data gathering and analysis will be used in the development of this guide, namely the *Amalgamation of Foundational Literature*, *Case Study Analyses* and a continual process of *Implementation, Testing and Iteration*.

Amalgamation of Foundational Literature

I will conduct a comparative study of the works of Adams (2015), Bates (2004), Schell (2015) and Selinker and Snyder (2018) using content and thematic analysis to amalgamate each of their sets of puzzle design guidelines into a single model. This model will form the very first iteration of the guide to puzzle design, and thus the *Implementation, Testing and Iteration* process can begin upon its completion. Content analysis is appropriate, as the summarizing of the various texts through the use of codes and identification of the frequency of ideas gives us a good grasp on the prominence of those concepts in the broader field, which in turn has implications about their priority and value.

Further iteration of this model in response to the testing results will be steered by additional secondary data found in supporting academic work such as that of Asbell-Clarke et al. (2016), Danesi (2002), Fernández-Vara and Osterweil (2010), Kim (2006), and Linehan et al. (2014). The insights these works provide will not only improve this initial amalgamation process, but also assist with the *Case Study Analyses* to follow.

Case Study Analyses

Once the initial iteration described above is complete, part of the continued refinement of this guide will take place through the examination of key case studies. Three renowned puzzle designers have been selected as case studies for this investigation, namely Jonathan Blow, Kim Swift and Arvi 'Hempuli' Teikari. A concise model[5] of each of these creators' design philosophies will be constructed through thematic and qualitative content analyses of the *Dynamic Informal works*[6] relating to each of them. If need be, the information in this model can be further supplemented through the descriptive analysis of specific levels from *The Witness* (Blow, 2015), *Portal* (Swift, 2007), and *Baba is You* (Teikari, 2019), in instances where that contained in the *Dynamic Informal work* is incomplete.

Once a model of each of their philosophies exists, they can be compared to one another using a table of common elements, akin to that used by Salen and Zimmerman (2004, p. 9) for comparing and amalgamating the multiple definitions of the word 'Game'. As Salen and Zimmerman (2004) point out, this may result in a lack of consideration of the context and nuance of each philosophy; however, the format allows for clearer pattern identification and comparative results. The conclusions of this exercise have implications about the uniqueness of each property outlined in these philosophies, thus streamlining the amalgamation of common ideas into the developing guide and directing the decision of which strategies need to be tested experimentally in the *Implementation, Testing and Iteration* process.

Implementation, Testing and Iteration

As stated previously, once the initial iteration of the guide for effective puzzle design is complete (through the content analysis of our *Foundational Literature*), the process of *Implementation, Iteration and Testing* can begin. By implementation, I refer to the fact that a puzzle platforming game (*Allegory*) will be created in tandem with this guide[7]. The relatively simple mechanic set necessary for this style of game should allow me to focus on the puzzle design within it.

All the groundwork up to this point would have been a process of inductive analysis – gathering and collating information into new, usable formats. However, this portion of our study will utilise a more deductive approach, giving us a means to verify the value and applicability of categorized information. This deductive analysis will take place over two phases. Firstly, any curated information available will be used in the creation of *Allegory*, allowing me to reflect on the value of the content and the effectiveness of its format.

Secondly, a targeted playtest will be designed after the guide has been completed. A set of puzzles will be specifically designed as a control group and then enhanced using the guide to create a second set. These two sets will be separately distributed to a group of playtesters, who will then be asked to complete an online survey about their experience with the puzzles. While the data collected will inherently be qualitative, the medium through which it is gathered will allow us to use some quantitative analysis techniques to identify trends and correlations in the data.

This semi-quantitative analysis, paired with my own qualitative reflection, will hopefully highlight valuable implications about the quality and usability of the Guide to Puzzle Design. These findings can then be used to refine the final product via an iterative process.

---

[5] See *The Dynamic Informal Design Model* for details on said model.

[6] The concept of *Dynamic Informal work* is defined in the Rational section.

[7] It should be noted that a team of three masters students worked together on *Allegory*. I was the lead Level Designer and Developer, and primarily used the project as a means to assess different approaches to puzzle design. Rodwin Malinga (the lead Composer, Tech Artist, and Project Manager) used the project to investigate the possible inclusion and effect of emergent music in videogames. Daniël Hartwell (the lead Sprite Animator, 3D Modeler, and Artist) used the project to explore the affective nature of colour theory.

The Guest List

The starting point for this research involved fishing for the most valuable existing guidelines, ideas, and theories about puzzle design. This section will briefly introduce you to the most notable references and case studies that were analysed throughout this process, to prepare you for the discussions to follow. In hopes of calming the choppy waters, I had sorted the information available into two categories, namely *foundational literature* and *dynamic informal* work.

Foundational literature refers to prior research and reference material within the field of puzzle design. These works most commonly aim to accomplish one of two core goals. They either attempt to teach the reader how to create effective puzzles in a general sense (Adams, 2015; Bates, 2004; Schell, 2015; and Selinker and Snyder, 2018)), or they will attempt to address a specific question regarding puzzle design in more depth (Asbell-Clarke et al., 2016; Danesi, 2002; and Linehan et al., 2014).

Dynamic Informal works refer to less regulated information sources, such as interviews, video essays, conferences, and documentaries. The design philosophies of Jonathan Blow, Kim Swift and Arvi 'Hempuli' Teikari will be extracted from these dynamic works and used as a set of *Core Case Studies*. Through their analysis, we will be able to exemplify, expand on and challenge the principles discussed in the *Foundational Literature*.

Foundational Literature

This section introduces the most noteworthy reference work I engaged with in the puzzle design sphere, focusing on a selection of four core works by Adams (2015), Bates (2004), Schell (2015) and Selinker and Snyder (2018). These works were primarily chosen for the similarities in structure and purpose that they share with this research. Let us briefly look at their formats and intentions to highlight the context in which their advice is presented.

Game Design (Bates, 2004)

Game Design is a reference book in which Bates attempts to teach the reader about the entire videogame development process. Because of this monumental scope, Bates is often forced to sacrifice detail and formatting. Additionally, many of his principles are written specifically for the design of *narrative* or *adventure* games, reducing their relevancy in the *generalized* guide we hope to create (see *Generalisability* section for details). That said, in exchange for navigating his somewhat unfocused train of thought, Bates rewards us with a myriad of valuable insights on puzzle design.

The Art of Game Design (Schell, 2015)

In *The Art of Game Design*, Schell (2015) focusses on teaching the reader all the *design* related intricacies that go into game development, disregarding much of the minutia of implementation. This scope aligns better with our research, and allows Schell to present his recommendations for effective puzzle design in a well-structured set of ten principles. These offer a trove of valuable knowledge, as well as a useful format for structuring our guide.

Fundamentals of Puzzle and Casual Game Design (Adams, 2015)

Adams (2015) takes this focus even further in his book *Fundamentals of Puzzle and Casual Game Design*, concentrating solely on the design of puzzle and casual games. This allows Adams to go into significantly more depth, especially when discussing the possibilities offered to puzzles by computers and Scott Kim's Eight Steps for puzzle design. These steps are arguably the most similar existing work to our proposed puzzle design guide, and so will form an integral part of our framework going forward.

Puzzlecraft (Selinker and Snyder, 2018)

In contrast to these three works about videogame design, *Puzzlecraft* (Selinker and Snyder, 2018) is primarily about puzzle design, and attempts to teach the reader how to create a vast array of specific puzzle types through concise, step-by-step guides. These applied guides are exceptionally useful as framing devices, as they are previous attempts at what we hope to create. That said, the surplus of puzzles Selinker and Snyder discuss prevents them from going into sufficient detail in their guide to videogame puzzles, leaving room for improvement.

Puzzle Designer Case Studies

The *Foundational Literature* offers us a structured and curated starting point, but its stagnant nature somewhat closes it off from new ideas. Thus, I will use a collection of *Dynamic Informal* sources to supplement, challenge or exemplify the *Foundational Literature*.

This will be accomplished through a comparative study of three *case studies*, namely the philosophies of the prominent puzzle designers Jonathan Blow, Arvi 'Hempuli' Teikari, and Kim Swift. Said philosophies will be extracted from the sea of *dynamic informal* sources available[8]. The following sections outline the most noteworthy beliefs of each designer, as well as prominent games designed by each of them, which can be examined to supplement our analysis later.

Jonathan Blow

I believe Blow views games as artforms, and so strives to reveal deeper meaning through his work (Alexander, 2014; Davies, 2015; Parkin, 2008; The Game Overanalyser, 2019; Writing on Games, 2016). He proposes that making puzzles is not about creating new, unique things, but rather discovering those that already exist (Game Maker's Toolkit, 2016). This forms the heart of his development ideology, which I refer to as *Interactive Emergence*.

Blow starts by establishing a system that he can play with (or creating a "construction set" (Kim, as cited in Adams, 2015)), and then does just that. As he plays, he notes any patterns or strategies that arise in his playstyle[9], insistent on showing those that are interesting or unexpected to his players (Game Maker's Toolkit, 2016; Parkin, 2008; The Game Overanalyser, 2019). Some of these will be forged into puzzle rules[10], while others will appear more subtly through solutions and player perspective shifts.

But rather than just identifying these properties, Blow wants players to *understand* their emergent nature (Game Maker's Toolkit, 2016; Parkin, 2008). Thus, Blow uses his puzzles to take the player on a curated version of his initial exploratory experience. This is evident in *The Witness*, in which learning the rules of each puzzle *is* a puzzle in and of itself (a fantastic variation of Schell's seventh principle regarding pyramid structure (2015)). To discern how each puzzle functions, players must hypothesise, experiment, and reflect on results (Writing on Games, 2016). This scientific examination that Blow demands creates a necessity for a deeper understanding of mechanics and their applications.

Blow's philosophies challenge the core literature, subverting the proposed concepts of embedding puzzles into the diegesis (Bates, 2004; Schell, 2015), avoiding trial and error (Bates, 2004), misdirection

[8] These include interviews with acclaimed puzzle designers (Alexander, 2014; Parkins, 2008), articles on websites dedicated to the digital arts and videogame spheres (Davies, 2015; Kim, 2006), video essays (The Game Overanalyser, 2019; Writing on Games, 2016), talks and conferences with professionals from the puzzle game industry (Blow et al., 2013; CS50, 2018; Gamelab Conference, 2019; Kim, 2008; Menzel, 2019; Swift, 2007), and documentaries (Game Maker's Toolkit, 2014; Game Maker's Toolkit, 2016; Millard, 2021; O'Dwyer and Jayne, 2017).

[9] Much like what Selinker and Snyder (2018, p. 180) refer to as "seeds" when describing Masyu puzzles.

[10] Such as those that demand segregation of spaces, creation of specific shapes and the use of symmetry in *The Witness*.

and trickiness (Bates, 2004; Selinker and Snyder, 2018), and several of Schell's principles[11] (Schell, 2015). That said, Blow masterfully balances these subversions with the application of other recommendations from the core works.

Puzzle design, according to Blow, is about fully exploring each mechanic, allowing it to naturally reveal unforeseen aspects about itself. Design need not happen with fixed goals or solutions in mind, as these will naturally surface – they are emergent in nature.

Arvi 'Hempuli' Teikari

Scanning through the long list of mechanics of *Baba is You*, Teikari selects unexplored mechanical combinations that he believes will surprise or amuse his players. These will each become the focus of a new puzzle, marking the starting point of Teikari's design process (Gamelab Conference, 2019; Couture, 2018; Game Maker's Toolkit, 2019; Talk & Dev, 2020; The Game Overanalyser, 2019).

Framing a new combination as the moment of insight, Teikari uses a process he calls "*Reverse Engineering*" to obfuscate this solution, thus creating an engaging level (Gamelab Conference, 2019; Game Maker's Toolkit, 2019; The Game Overanalyser, 2019). This process introduces multiple steps that lead up to the insight in each puzzle, balancing the inspiration with sufficient effort (as recommended by Selinker and Snyder (2018))[12]. Once the player has engaged with a specific combination, Teikari can use it in later puzzles to obscure a *new* goal (Game Maker's Toolkit, 2019), creating a naturally appropriate difficulty curve.

Teikari's design in *Baba is You* subverts many foundational puzzle design principles. The level elements' lack of intrinsic purpose makes it nearly impossible to embed them in a diegesis (as recommended by Bates (2004) and Schell (2015)), and the player-defined ruleset conflicts with Schell's principles regarding easily understood goals, giving a sense of solvability and perceptual shifts (Schell, 2015). These subversions are partly what make *Baba is You* so unique, but only function because of Teikari's proficiency with a myriad of other foundational design principles.

Kim Swift

Swift prefers to think of her creations as toys, as they encourage players of all ages to use their imaginations (Rigney, 2012). I think this sentiment really sums up her design philosophy, which prioritises *play* and the *player*. Developing with players in mind – the heart of Swift's puzzle design – is manifested in two key ways.

Firstly, Swifts emphasizes the need to *train* one's players. By using deliberate level design, one can teach their player how to engage with a mechanic on a deeper level (Valve, 2007; Swift, 2007). Linehan et al. (2014) discuss the prevalence of this concept in *Portal*'s difficulty curve at a macro scale[13], but it is evident in the micro design as well. Swift has a masterful understanding of her players' capabilities as they enter a level, as well as what she needs them to comprehend by its end – a concept given a lot of weight by Selinker and Snyder (2018). Thus, Swift defines very explicit goals for every level, creating a concise player experience.

Secondly, Swift heavily prioritises playtesting, adamantly stating that it should be a continual practice throughout development and should begin as early as possible. At the first signs of playtester confusion, the *Portal* team would either introduce additional training levels, or they would decisively

[11] Specifically, Schell's principles about making the goals of puzzles easily understood, using hints to extend interest, prioritising the moment of insight, and avoiding the use of perceptual shifts.

[12] This balance is what allows Teikari to exhaustively explore mechanical combinations without inciting the eureka sickness that Adam Millard – The Architect of Games describes (with reference to *Superliminal*) (2021, 6:00).

[13] Said curve applies the "sawtooth" progression described by Kim (Kim, as cited by Adams, 2015, p. 13).

remove or rework puzzles and mechanics (CS50, 2018; Valve, 2007; Swift, 2007). Player experience always took priority over interesting ideas.

Overall, *Portal* shows a near flawless execution of the discussed puzzle design principles, making it an exceptional means to exemplify the concepts in the foundational literature.

Amuse-Bouche

A Guide to Guide-Building

Firmly understanding the content of the Foundational Literature and Dynamic Informal work will not get us anywhere without an effective way to communicate this understanding. The heart of this project lies in reformatting the plethora of existing information to be *appliable* and *user-friendly*; thus we should draw inspiration from formats that have come before. By looking at *four key characteristics* of the formats apparent in the foundational literature, we can deduce what structure, order, and level of detail would be most valuable to readers of the guide.

A Note on Intention and Audience

Before we delve into the critical analysis ahead, I'd like to point out an important consideration. Throughout this chapter, I will be making extensive commentary on the design decisions made by authors and content creators regarding the way in which their work is presented. However, this is not meant to imply that any single structure is universally superior to another.

While each of the works analysed in this chapter *are* connected by the common theme of puzzle design, they were each created within a unique *context*. The organization of a piece's content is generally informed by its target audience, medium and the intended use of the piece[14]. In this section, I'll be highlighting the design decisions that align well to the goals set out in the *Aim* section, while discussing those that jeopardise these aims. Remember that our goal is to make something easy to engage with for puzzle designers just entering the field. This context will guide the choices we're about to make.

Key Characteristics

In examining the existing information available on puzzle design, I found that the ease of engaging with the portrayed concepts varied greatly from source to source. This was in no small part due to the format with which each of these works was presented. I have identified four key characteristics that have a notable effect on the usability of a source, namely *Consistency*, *Point Format*, *Order of Presentation*, and *Generalisability*. Gaining an intricate understanding of how these characteristics influence usability will ensure the most effective use of them in the final guide to puzzle design. Each of these characteristics is explained below, as is their effect on the reader's engagement.

Consistency

Consistency is a hallmark of a work that feels organized and easy to engage with. The concept can be applied to one's work in three ways. Essentially, if several connected ideas are being presented, they should all use the same *format*, have the same conceptual *framing*, and be explored with a similar complexity or *level of detail*.

Consistent Format

A consistent format is imperative as it allows the audience to directly compare the various points being made to one another. When one juxtaposes two points with the same structure, they are essentially comparing apples to apples, making it a simple matter to discern a point's relative value and importance. Additionally, this format will create a degree of systematic understanding in the audience. This means that even if a reader encounters a particularly difficult point, the fact that

[14] To illustrate this, consider a chapter in a reference book and a YouTube video designed for the layman. Both works may be about the same subject matter, but they simply cannot be presented in the same fashion.

they are familiar with its format can give them contextual clues, easing the process of breaking it down and analysing it.

Selinker and Snyder (2018) apply this concept particularly well, as each of their individual guides has an identical structure. They consist of a detailed written explanation accompanied by an example of the specific puzzle type being discussed. This explanatory text always opens with a concise description of that puzzle's format and rules, and then moves on to an ordered list of actionable, chronological steps for creating an instance of that puzzle. Even the steps themselves follow a consistent format, starting with a general instruction and ending with an example of how this instruction was applied in the accompanying puzzle.

Another strong application of this concept can be seen in Schell's set of ten principles for good puzzle design (2015). While Schell's format is not quite as compartmentalised as Selinker and Snyder's (2018), it is still commendably consistent. Here, Schell (2015) uses a concise thesis statement for the heading of each principle. This statement will either have a clear directive, or will immediately highlight the value that a principle offers[15]. This functions as a useful summary for the reader, framing their expectations on their first read and easing the process of re-engaging with the work, should they need to.

Finally, certain portions of Bates' sixth chapter, *Designing the Puzzle*, also showcase this concept when read in isolation. This can be seen most prominently in his taxonomy of puzzle types[16] (Bates, 2004, p. 120). In this subchapter, each puzzle type Bates presents follows a consistent structure, opening with a generalized explanation of the features of said puzzles, and ending off with an example showcasing how that puzzle may appear in a videogame. Even these examples remain consistent, as all of them are framed around the goal of opening a door. In this specific section, Bates' points bear a striking resemblance to Selinker and Snyder's (2018) puzzle guides.

That being said, Bates' work does not remain consistent across subchapters (Bates, 2004). While they each contain information surrounding the same general topic, the structure of each of Bates' sections is unique to that section, making it tough to compare points that are not directly grouped together[17]. When considering all these examples, it becomes clear that the final guide should commit to a single format, not only for each individual section, but across the entirety of the product.

A note on Visual Format Consistency

The concepts discussed above can be further accentuated through the use of consistent visual representation within a piece. One can immediately communicate the function or category of a section of text by committing to a distinct visual style for text of that type. Selinker and Snyder (2018) are particularly good at doing this. Throughout *Puzzlecraft* (Selinker and Snyder, 2018), example puzzles always appear on a grey background, the thesis statement of each guide step is bolded and italicized, and the applied examples given in these steps always appear in italics. Schell and Bates also apply these concepts to some extent, with Schell (2015) placing each of his lenses in specifically coloured and numbered text boxes, and Bates (2004) using a grey background to represent information provided by other designers.

[15] For instance, the principles "*Parallelism Lets the Player Rest*", or "*Pyramid Structure Extends Interest*" (Schell, 2015, p. 248) immediately communicate the positive consequence of applying that principle.

[16] In fact, this taxonomy is structured so well that it has been praised as a "must read" in the Informal Dynamic work (The Game Overanalyser, 2019, t. 12:45).

[17] This is discussed further in the *Consistent Framing* section below.

While these techniques may not seem integral, their value quickly becomes evident when engaging with a work. If done well, the reader will never have to search the text for context. The very appearance of the segment they are reading will inform them of its purpose, speeding up the process of analysis and understanding.

Consistent Framing

When presenting someone with information, it is often possible to convey the same core concept in different ways. For instance, you could list all the actions that a person *should* be taking to achieve some goal, or give an exhaustive list of actions they should *avoid*. The end result will be the same – your audience will only perform acceptable actions (provided that they listened to you, of course). I have taken to referring to this concept as the *framing* of information.

The choice of how one frames their work is generally governed by the conciseness and completeness of the final presentation. That said, when it comes to the creation of our guide, I believe that this choice should be a commitment adhered to across the entire work. Let's look at some examples in the Foundational Literature to illustrate why.

Schell (2015) has used a consistent framing for the majority of the information he presents about puzzle design. Each principle has been constructed to read as a goal or instruction which will have a *net positive effect* on the quality of the reader's puzzles[18]. Thus, there is little room for confusion about what should be applied and what should be avoided. The value of Schell's commitment to this framing quickly becomes evident when comparing his principles with Bates' sixth chapter, *Designing the Puzzle* (Bates, 2004).

In said chapter, Bates (2004) explores puzzle quality by discussing *bad* puzzle types, characteristics of *good* puzzles, methods to adjust puzzle *difficulty,* and finally concludes the chapter with a broad exposition about puzzle *design*. The changing goals and inconsistent framing in each subsection result in incomplete or redundant information and difficulties in discerning the relative value of each point. This can make it tough to actively engage with Bates' core arguments, a stark contrast to Schell's clearly presented principles (Schell, 2015).

These examples showcase the benefits of sticking to a consistent framing of information across the entire guide. It reduces the chance of reader confusion, allows the audience to easily compare the value of various points, helps prevent redundant information, assists in making the guide as complete as possible, and eases the process of giving each point a priority rating[19]. As such, our guide will commit to a single means of framing its content.

Consistent Level of Detail

If several connected ideas are presented, they should all be presented with similar complexity and level of detail. The chosen depth one goes into should be selected based on the target audience of the piece and the goals it is meant to accomplish.

This is important for a number of reasons. Firstly, it ensures that the reader can engage with the entire work at a uniform pace. This smooths out the process of working through a piece, which generally results in a better experience when applying its advice. Secondly, it increases the chances that the reader's needs are consistently being met in every section. Finally, it gives one's readers a strong indication of the relative complexity or weight of each section. If everything is written at the

[18] With the sole exception of Schell's 10th principle, "*Perceptual Shifts are a Double-Edged Sword*" (2015, p. 250), which functions as a warning for a plausibly bad practice in puzzle design.
[19] Explained later in the section regarding *Priority*.

same level of detail, the reader can easily surmise that the longer a section is, the more there is for them to do or to consider when working through the concepts discussed in that section.

There are various examples that showcase the importance of this principle, but none as valuable as Selinker and Snyder's *Puzzlecraft* (2018). For the most part, these two authors apply this concept excellently throughout the book. Each of their individual guides breaks up the process of creating a specific puzzle type into bite-sized steps of similar complexity. After working through a few of the guides, the reader gets a good idea of how long an average step in a guide will take and how difficult it will be to implement. This, in turn, means that a reader can use the number of steps in any individual guide as an immediate (and relatively accurate) indication of how challenging and time consuming it would be to create a puzzle of that type.

That being said, the pair fail spectacularly in this regard with their explanation of the creation of videogame puzzles (Selinker and Snyder, 2018, p. 202). Each of the steps they offer for this individual guide are, in reality, lengthy processes requiring a lot of iteration and testing. It's not as if Selinker and Snyder aren't aware of this (they even point this out to the reader explicitly), but it does mean that the instructions given here are not nearly as direct and succinct as they are in the pair's other guides. This isn't truly the fault of the authors' though; it was simply an overly ambitious goal to address something as nuanced and broad as a 'Videogame Puzzle' in the limited space they had designated for an individual guide.

Another great example to look at is *Kim's Eight Steps*, as presented by Adams (2015, p. 12). For the most part, these steps remain relatively consistent throughout the chapter. While steps One (*Find inspiration*), Two (*Simplify*) and Seven (*Devise a sequence*) do go into slightly more detail than the rest of the guide, the difference is small enough to be overlooked. All in all, the steps are written with a commendable amount of depth, each of them giving clear and somewhat actionable instructions. That is until one considers step Five (*Construct the puzzles*). Here, Adams records some brief details about what makes a good puzzle (according to Kim), but fails to give a clear process one can follow to actually create a puzzle[20].

The level of detail used in *Crafting a Videogame Puzzle* (Selinker and Snyder, 2018, p. 202) and *Kim's Fifth Step* (Adams, 2015, p. 12) is not the inherent problem – it is only in contrast to the subchapters around them that these sections become problematic. The average level of detail throughout a work will attract a specific audience, and each section one presents creates and reinforces certain expectations in said audience. When a single section is not what the attracted audience is looking for, or when a section does not live up to the developed expectations, it will inevitably feel out of place. In the worst-case scenario, it will not meet the needs of a large portion of your readers. This can frustrate one's audience, shatter immersion or even cause readers to give up without engaging with the rest of the work.

As such, we should strive for a consistent level of detail throughout our guide, such as that seen in the majority of the guides developed by Selinker and Snyder (2018), seven of Kim's Eight Steps (as cited by Adams, 2015), and in the persistently elegant Principles for Good Puzzle Design discussed by Schell (2015).

[20] A process such as this only surfaces later in Adams' *Sokuban* analysis, in which he recommends starting with a solved level and working one's way backwards to a starting point (Adams, 2015, p. 15). But even so, this suggestion feels very specific to that case study and is completely separate from the initial presentation of the Eight Steps.

Guide Criteria - Consistency

Various types of Consistency have proven to be a virtue when used in our foundational literature, generally making each of the works somewhat easier to engage with. Using a consistent *format* and *framing* of information improves the reader's ability to compare the relative value of different points, as well as their ability to process and understand more challenging points. It also reduces the chances of redundant information and incomplete sections. Due to these benefits, the final guide will commit to a single, consistent format. Additionally, the content of the guide will be consistently framed as directives that result in a positive effect on the reader's work.

To further reinforce these benefits, the guide should use a consistent *visual representation* system for the various categories of information it contains. This supplies the reader with context at a glance, making any portion of text easier to process. Finally, the entire guide should adhere to a consistent level of detail. The specific depth the guide will go into will be informed by the sections to follow (specifically *Point Format* and *Generalisability*).

Guide Criteria: Consistency
• A *single* format should be committed to and applied across the entire guide. • Every point in the guide should be *framed* in the same way. • The guide should represent each of the key types or categories of information it contains with a distinct visual style. • The entire guide should have a consistent level of detail that is appropriate for the intended audience.

Point Format

Due to the complexity and nuance of this field, we will inevitably be offering a ton of information to our readers. How we compartmentalize that information into bite-sized chunks will govern how easy the work is to digest. Works from our Foundational Literature that use some form of point-based structure generally stand out as the easiest to engage with and the most effective to apply in practice.

By point-based structure, I refer to the procedure of clearly segregating one's work into elementary ideas or instructions, then concisely and completely dealing with each of these in turn. The works of Kim (as cited by Adams, 2015), Schell (2015) and Selinker and Snyder (2018) all apply this philosophy to some degree. In these works, there are two apparent ways of approaching this process – either framing each individual point as an actionable *Instruction*, or as an overarching *Goal*. Each are uniquely valuable, as discussed below.

Instructions

Works that frame their elementary ideas as instructions will present their readers with a list of things to do, as well as a time frame or order in which they should be done. These works are great at getting the ball rolling, allowing the designer to immediately start putting theoretical concepts into practice. They guide the reader through the creation process, allowing the designer to focus their energy and decision making on the product itself, rather than the methodology needed to create it.

We can see this philosophy applied relatively effectively in the numerous guides of Slinker and Snyder's *Puzzlecraft* (2018). Here, the authors break up the process of creating a specific type of puzzle into clear and distinct phases, each of which is then presented to the reader as an instruction for them to apply. Each instruction is made up of a concise thesis statement followed by a more detailed description of the phase. The thesis statement is actionable and gives the reader a strong idea of what they will accomplish by the end of that step, while the description that follows it breaks the statement down into smaller, more direct instructions. Kim (as cited by Adams, 2015) uses a very similar format for his *Eight Steps for Puzzle Design* as well.

This format is particularly useful for beginners, as it gives them something they can immediately engage with and reduces the amount of critical analysis necessary from them. Additionally, the fact that it yields visible results quickly is a good confidence boost for someone new in the field. That being said, this format does not come without its shortcomings. A particularly notable one is that this format may allow readers to follow it blindly. By taking away the need for critical analysis and reflection, this format may reduce the amount of knowledge retained by the reader *after* completing the design process.

In the worst case, the reader may simply become reliant on having a specific list of instructions to follow. This not only means that their growth as a designer is limited, it also implies that the reader would be unable to work around any missing steps in the process offered by the guide, should it be incomplete. This may sound negligible, provided the author puts sufficient effort into ensuring that their guide covers all the necessary steps. However, due to the second shortcoming of this format (discussed below), this occurrence isn't all that inconceivable.

This format is excellent at conveying how to implement simple and relatively short processes. But the more nuanced and lengthy a process becomes, the harder it is to explain in this way. We have already seen the evidence for this in our foundational work. In their individual guides, Selinker and Snyder (2018) do an exceptional job of teaching the reader how to create a wide array of more conventional puzzles, such as various forms of crosswords (ibid., p. 86) and number logic puzzles (ibid., p. 131). It is only when they come to more intricate puzzle formats, such as *videogame*

puzzles, that they drop the ball (ibid., p. 202). The same can be seen in Scott Kim's Eight Steps (as cited by Adams, 2015), where the most complex steps are presented somewhat ambiguously, making them far less serviceable[21].

This is potentially because elaborate tasks often don't have a single, objectively correct implementation method. There are some general phases in the process that *must* come before others due to obvious dependencies, but aside from these, the order and necessity of small-scale steps really comes down to the personal preference of the designer[22]. Puzzle design is, after all, an artform in its own right (Alexander, 2014; Davies, 2015; Indie Cade, 2011), a fact which necessitates some degree of designer agency and creative freedom. This is very likely the reason that the Foundational Literature which attempts to break down videogame puzzle design into specific recipes tends to focus on broader strokes rather than the nitty gritty (Kim, as cited by Adams, 2015; Selinker and Snyder, 2018).

Goals

Works that frame their content as a set of goals will present their readers with a list of things to consider over a larger portion of the design process. Unlike instructions, these goals usually don't come with a specific time frame or defined order. Instead, they are there for puzzle designers to *always* keep in mind when making a design decision, guiding them to the choice that will yield the best final product.

These works are valuable for a myriad of reasons, the clearest being the versatility and longevity of the advice that they offer. The goals normally describe some form of end point the designer should aim to reach, and so are not heavily affected by the minutia of the problem at hand. In most cases, the information can be meaningfully applied regardless of the type of puzzle one is creating. Thus, even if new puzzle types arise, the work will remain relevant.

Consequently, this also means that these works inherently create a degree of independence in their readers. This format better lends itself to *Generalisability*[23], and so is more likely to allow designers to freely explore any puzzle type they have in mind. But perhaps even more importantly, users of these works are still expected to make the specific design decisions themselves. This gives them valuable practice at tackling problem solving with critical thinking, a skill which can be broadly used to face future challenges, without the need to directly reference *any* guide.

Schell's ten principles are a stellar example of this format, showcasing just how versatile it can be (Schell, 2015). After claiming that these principles can be useful in any game genre, Schell convincingly backs up his opening statement by showing the reader a myriad of examples in which the principles are either successfully applied, or could have been used to improve an experience. What makes these examples so compelling is how different they all are. Schell frames this as a discussion about videogames, but references puzzles present in everything from toys to word games to mystery novels. This ridiculous amount of variance does a fine job of showcasing just how broadly applicable Schell's principles are.

A further powerful point about these examples is that they are just that – *examples*. Schell never explicitly tells the reader that the principles have to be used in this fashion, he merely offers these use cases to strengthen the reader's understanding of the concepts he has portrayed. The burden of actually employing these concepts within one's specific context still lies upon the reader. While

[21] Both Selinker and Snyder's and Kim's mishaps are discussed in more detail in the *Consistent Level of Detail* section above.
[22] This matter is further discussed in the *Chronology* section below.
[23] This concept is explored in more detail in the *Generalisability* section.

this can be seen as a positive, as it gives the reader room to practice their critical thinking, it is also the most notable downfall of the goal orientated point format.

For a designer to be able to effectively evaluate their project and accurately steer it in the direction of these goals, they require an intricate understanding of the goals themselves. This is why many works that take this approach (such as Schell's ten principles (2015), or the Introduction of Selinker and Snyder's *Puzzlecraft* (2018, p. 6)), spend a lot of time on elaboration and offer a multitude of examples. Unfortunately, this requirement for comprehension can be daunting to new designers, a feeling that is potentially exacerbated by the overload of information and elaboration that this format somewhat demands.

These difficulties become even more prominent in works with structures that have a degree of ambiguity, such as Bates' sixth chapter, *Designing the Puzzle* (2004). In this chapter, Bates explains several useful design goals to the reader, but they are masked under an inconsistent categorization structure[24]. Selinker and Snyder's introductory chapter does a slightly better job at presenting design goals to the reader, but at times the fluid paragraphing they use to convey the information fails to adequately emphasise exactly which notes are the imperative goals the reader should focus on (Selinker and Snyder, 2018).

It takes a deft hand to present a collection of goals with enough detail for the reader to fully understand them, but not so much detail that they become overwhelming or lack focus. While Schell's ten principles (2015) do suffer from some of the above-mentioned issues, they are still conceivably the strongest example of this style of point formatting within our collection of foundational works. Thus, they will certainly be referenced as a starting point for any goal orientated points made in the guide.

Guide Criteria - Point Format

Overall, it appears that the use of a Point Format is advantageous for the type of work we hope to create. It helps keep things concise, clear, and segregated, allowing the reader to deal with a single concept at a time – all of which are invaluable qualities for a guide. The structure of the points we present will be similar to those seen in the works of Kim (as cited by Adams, 2015), Schell (2015) and Selinker and Snyder (2018). By this, I mean that each point will be made up of a succinct and clearly emphasised thesis statement, followed by a more detailed description that unpacks any nuance that is necessary for the reader.

Regarding the choice between Instruction and Goal orientated points, I feel that both formats have a place in the final product. As the shortcomings of each align very well with the strengths of their counterparts, they can be used in conjunction to support one another. Instructions will be used to give the reader a starting point, direction and an initial boost in confidence, while Goals will be used to strengthen their skills as designers, and ensure the longevity and versatility of the guide. In general, I will avoid making information redundantly available in both formats. If the point being made can be presented in either format, I will prioritise the inclusion of Goals over Instructions. Finally, I will be sure to clearly communicate whether a point is a Goal or Instruction to the reader, as well as how each should be used.

[24] As discussed in depth in the *Consistency* section.

Guide Criteria: Point Format

- The information in the guide should be presented with a variation of Point Formatting.
- Each point in the guide must be made up of a concise and emphasized thesis statement, followed by a more detailed description.
- Points framed as Instructions should be used to give the reader direction, making it easier to start the design process.
- Points framed as Goals should be used to train the reader's critical thinking and ensure the longevity and versatility of the guide.
- Information should not be repeated in both formats.
- Goals should take priority where possible.
- Which format each point takes (as well as how to use each of them) must be made clear.

Order of Presentation

Many successful guides or instructional works are presented as a collection of bite-sized points that the reader can digest individually[25]. As we plan to commit to a similar format, we need to give careful consideration to the order in which we present these points. This order is generally informed by two core principles, namely the *Chronology* and *Priority* of each point. As we discuss these topics, it is valuable to keep in mind that this guide is targeted at *beginner* puzzle designers, and so we hope to make it as clear and accessible as possible.

Chronology

By Chronology, I refer to the ideal of considering the point in the videogame's development process at which a specific instruction is most relevant. To as much a degree as possible, the guide should aim to present information to the reader not only at a point where they are capable of engaging with it, but also at the point where they will gain the most value from it. Finding this sweet spot will require an intricate understanding of the dependencies existing between the instructions within the guide.

The clearest examples of works that engage with Chronology are those that are presented as an ordered list of instructions, namely Selinker and Snyder's subchapter on *Crafting a Videogame Puzzle* (2018, p. 202) and *Scott Kim's Eight Steps*, as presented by Adams (2015, p. 12). In both these examples, the reader's ability to implement any individual step is completely reliant on them having engaged with all the steps that precede it. The reader is expected to work through the instructions in the same order in which they encounter them. This is useful, as it gives the designer a clear direction – there is no room to be confused about when they should be implementing something.

But this design choice comes with its fair share of trade-offs, the most prominent of which can be clearly seen in the two above-mentioned examples. Works that give an explicit order such as these are often written with less detail, discussing concepts in broader strokes rather than with specificity. Sometimes, this can even lead to instructions that are too vague to effectively apply, such as Kim's fifth step, *Construct the Puzzles*[26] (Adams, 2015, p. 13).

This is most likely because the act of designing a videogame puzzle is as much an artform as it is an engineering challenge. While the over-arching process may remain relatively consistent, the order in which one tackles the minute steps which *make up* these broader processes is somewhat of a personal choice. Writing with these broader strokes allows the reader to conserve some of their artistic agency, but this can lead to a lack of applicability if the reader is new to the field. As such, our guide needs to find an effective balance between specificity and afforded freedom, in order to remain clear to new designers while not being completely prescriptive.

Priority

Priority is connected to the relative value of each individual point offered by the guide. When curating advice to include in the final guide, we will obviously select that which will likely improve the quality of the puzzles created with it. But even if all the information presented has the potential to improve the final product, certain philosophies and suggestions will have a predictably greater effect on a finished puzzle than others. As the reader may not have the scope to implement every

[25] Such as the works of Selinker and Snyder (2018), Schell (2015) and Kim, as presented by Adams (2015).
[26] This concept is further discussed in the *Consistent Level of Detail* and *Instructions* sections above.

single point from the guide, it is important that they are made aware of which points have the greatest potential so that they may prioritise the implementation of these.

The value of this extra effort is made clear by looking at the works of Schell (2015) and Bates (2004). While he never explicitly states this, it appears that Schell has ordered his ten puzzle design principles based on their importance, starting with the most valuable principle (Schell, 2015, p. 243). This is useful, as if a reader begins applying these principles in order, they are likely to get the most utility possible from them, even if they end up stopping short of the tenth principle due to time, scope or resource constraints.

In addition to this clear advantage, Schell's principles also feel natural to engage with, as there is an inherent implication of what he wants the reader to focus on (Schell, 2015). In comparison, Bates' work feels somewhat disorganised (Bates, 2004). Within his sixth chapter, *Designing the Puzzle*, there is admittedly a vague semblance that the points made earlier in each subchapter are more important than those that appear later in said subchapters. But this structure is not emphasised, and it is made even harder to see by the fact that this priority-based order resets each time a new subchapter begins.

This lack of prominence leads to a feeling that Bates wants the reader to focus equally on every point that he makes (Bates, 2004). When one considers the volume of information contained in this chapter, this expectation can come across as daunting to new designers or leave them fatigued well before they have reached the chapter's end.

Even though this issue is more prominent in Bates' work, both authors could have benefitted from further emphasising this concept. Reflecting on this, it becomes clear that our guide should explicitly denote the priority or potential value of each of the points it includes. To effectively implement this design, a clear and consistent means of identifying the value of a given point must be defined.

Guide Criteria – Order of Presentation

This guide should take the responsibility of ordering design processes away from the reader to as large a degree as possible. This is because the intended audience are novel designers, who may lack the experience necessary to decide on an effective order themselves. But unfortunately, due to the need to respect the designer's artistic agency and allow them space to experiment and grow, there is little value in completely setting this order in stone.

As such, I feel the most effective presentation is to use an order based *both* on chronology and priority. This order can be paired with clear hints to help the designer make an informed decision about what advice to use and when to use it. To accomplish this, we can split the guide into several sections that can objectively be organised chronologically, based on the steps laid out by Selinker and Snyder's chapter on *Crafting a Videogame Puzzle* (Selinker and Snyder, 2018, p. 202), and Kim's *Eight Steps* (as cited by Adams, 2015, p.12).

The specific instructions in each section can then be ordered by their priority, as well as be given a clear priority rating for the reader to consider. This value will be based on the prominence of the point across our four foundational works as well as its appearance in our three designer philosophy case studies. The reader can then be told to consider each point in the order in which they appear, but to choose the advice that best suits them in each section. This takes most of the burden off of the novel designer, but still gives them some agency, along with the information needed to make reasonable decisions.

Guide Criteria: Useful Order of Presentation
• The broader steps of the design process should be sorted chronologically within the guide. • Smaller scale instructions should be completable in an order of the reader's choice. • The guide should provide a suggested order for smaller-scale instructions based on their potential value. • The guide should have a clear and consistent means of determining a point's priority.

Generalisability

Generalisability refers to how applicable an instruction or information is to a broad range of topics, problems, or puzzle types. Usually, more generalisable works are more valuable, as they do not inherently limit the reader to a certain style or type of content, and their versatility means that they can remain relevant for far longer than more specific works. The potential issues caused by specificity are illustrated well by the works of Bates (2004) and Selinker and Snyder (2018).

In Bates' sixth chapter, *Designing the Puzzle*, he often relies so heavily on the puzzles existing in some broader narrative that some of his proposed principles are ineffectual in more abstract puzzle games (Bates, 2004). Additionally, Bates often resorts to explaining concepts solely through examples that illustrate them. This is good at creating an immediate, working understanding for the reader, but may not create the general understanding they need to apply the concept in different use cases. Both these issues show how specificity can reduce the range of topics one's work can be applied to effectively.

In a slightly different vein, the individual guides in Selinker and Snyder's *Puzzlecraft* (2018) showcase the potential obsolescence that goes hand in hand with specificity. While these guides excellently describe a long list of puzzle types, that list is finite. Should new puzzle types arise, the majority of the content in *Puzzlecraft* will not prepare the reader to create in this new medium[27]. This issue may necessitate the release of revised editions that address newer fields and topics. This solution has already been employed by Selinker and Snyder, whose original version of *Puzzlecraft* (Selinker and Snyder, 2013) was not quite as extensive as the edition we have been referencing thus far (Selinker and Snyder, 2018). This approach is merely a stopgap solution though, as the inherent problem persists even in updated works.

While generalization is a great means of tackling these problems, too much of a good thing is harmful. Scott Kim's Eight Steps, as presented by Adams (2015), serve to illustrate this warning. In these Eight Steps, Kim has cleverly taken the direct, instructional guide format and presented it in such a way that it is more broadly applicable. These steps serve as a sort of scientific method for puzzle design, and so can be used indefinitely for novel ideas and puzzle formats. But this large scope, paired with the steps' concise format, results in some ambiguous directives, like that of the fifth step (*Construct the Puzzles*) which is too broad to practically apply[28]. This shows that an over-commitment to generalisation can leave one's work vague and unserviceable.

These examples highlight that there is a balance between generalisation and specificity to aim for – a point at which one's work will remain useful into the future in a myriad of contexts, while still being practically appliable with clear goals or instructions. Of our Foundational Literature, the *Ten Principles of Puzzle Design* presented by Schell (2015) are arguably the best attempt at finding this balance. Each piece of goal-orientated advice Schell offers is specific and efficient, but still described such that it can be applied to any puzzle type. As such, we will aim to produce a similar level of generalisation in our guide to puzzle design.

Guide Criteria - Generalisability

Our guide needs to be generalised enough such that it can assist with the creation of a wide array of puzzle types for any genre of videogame. Additionally, its directives should be framed in a way that their use is not limited to existing puzzle structures, allowing the guide to remain valuable into

[27] This isn't entirely true, as the book does include some more general advice in the opening chapters that remains applicable. Additionally, by practicing creating the puzzles described by the book, the reader will probably hone a skillset and way of thinking that will still leave them significantly more prepared for new puzzle types than a complete layman.

[28] This shortcoming is discussed in more detail in the *Consistent Level of Detail* section.

the foreseeable future. That being said, the target audience of this guide are beginner designers, who particularly benefit from clear and concise instruction. As such, the process of generalising the work should not be taken so far as to leave any of the guide ambiguous or vague.

Guide Criteria: Generalisability
• The Guide should be generalized enough to be applicable to a wide variety of puzzle types, including those that have yet to be created. • The Guide must remain clear and actionable – the scope should not be so generalized that it becomes ambiguous.

Guide Design

It's finally time to discuss the actual appearance and structure of the guide. Armed with the insight of the previous section, we are now able to make effective decisions regarding the delivery of the puzzle design information hidden in our library of Foundational Literature and Informal Dynamic Sources. This section will showcase the final guide design, as well as explain how it meets all of the criteria laid out in the previous chapter.

Criteria Breakdown

This breakdown is simply a summary of all the criteria we have already encountered. I have collated each point into a single, labelled collection, so that it can be easily referenced in the explanations that follow.

Guide Criteria Breakdown	
Consistency	
CRIT A:	A *single* format should be committed to and applied across the entire guide.
CRIT B:	The guide should represent each of the key types or categories of information it contains with a distinct visual style.
CRIT C:	The entire guide should have a consistent level of detail that is appropriate for the intended audience.
Point Format	
CRIT D:	The information in the guide should be presented with a variation of Point Formatting.
CRIT E:	Each point in the guide must be made up of a concise and emphasized thesis statement, followed by a more detailed description.
CRIT F:	Points framed as Instructions should be used to give the reader direction, making it easier to start the design process.
CRIT G:	Points framed as Goals should be used to train the reader's critical thinking and ensure the longevity and versatility of the guide.
CRIT H:	Information should not be repeated in both formats.
CRIT I:	Goals should take priority where possible.
CRIT J:	Which format each point takes (as well as how to use each of them) must be made clear.
Useful Order of Presentation	
CRIT K:	The broader steps of the design process should be sorted chronologically within the guide.
CRIT L:	Smaller scale instructions should be completable in an order of the reader's choice.
CRIT M:	The guide should provide a suggested order for smaller-scale instructions based on their potential value.
CRIT N:	The guide should have a clear and consistent means of determining a point's priority.

Generalisability	
CRIT O:	The Guide should be generalized enough to be applicable to a wide variety of puzzle types, including those that have yet to be created.
CRIT P:	The Guide must remain clear and actionable – the scope should not be so generalized that it becomes ambiguous.

Information Goals

Naturally, there is a massive collection of information that we can potentially share with our readers. But because this guide is for beginner puzzle designers, not *all* of the information available in our library of sources is useful to them. The following table contains a summary of all the content that I have considered including in the guide. This content has been organized based on its importance, either falling into the *Imperative* or *Nice to Have* categories.

Information Goals	
Imperative	
INFO A:	Heading or Thesis Statement.
INFO B:	Priority of the point.
INFO C:	Detailed description of the concept.
Nice to Have	
INFO D:	Number of authors or designers that Agree with the point.
INFO E:	Number of authors or designers that Disagree with the point.
INFO F:	Smaller-Scale Steps that make up larger goals or instructions.
INFO G:	References proving designer and author beliefs.
INFO H:	List of specific authors or designers that Agree with the point.
INFO I:	Whether a point is a Goal or Instruction.
INFO J:	Whether or not I personally agree with the point.
INFO K:	List of specific authors or designers that Disagree with the point.

Finalised Format

Finally, we see an exert from the guide itself, displaying the format and appearance of a single one of its points. Various sections of this design have been labelled from 1 to 7, so that these areas can be referenced in the discussion to follow. By the end of this chapter, the design presented should be justified, with reference to both the criteria we had set out to meet, as well as to the information we had hoped to include.

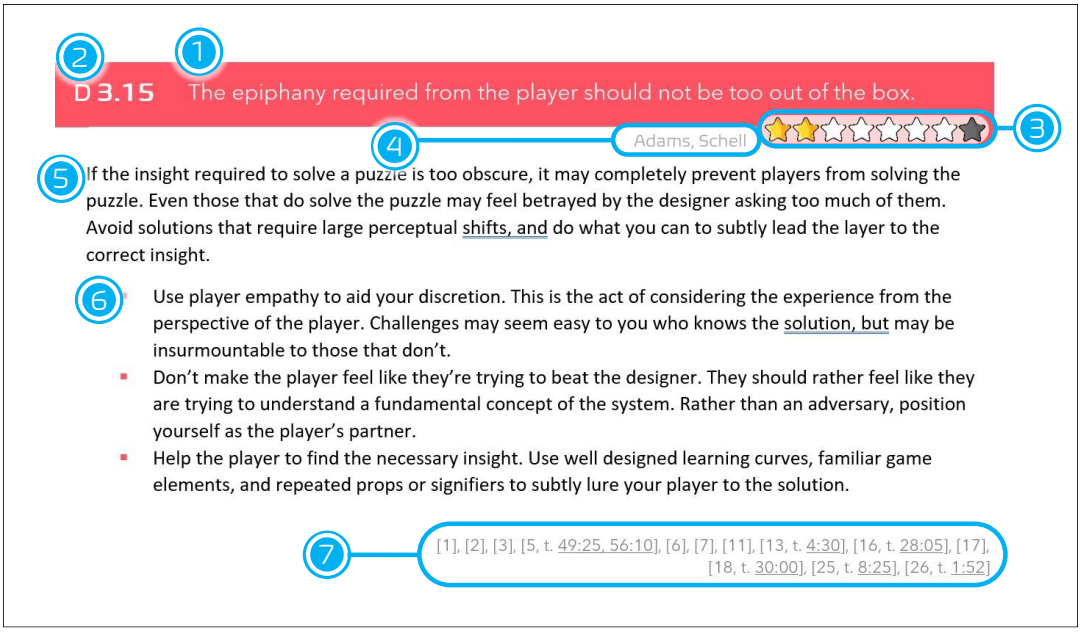

Figure 2: Format of a point from the Final Guide

The format seen in *Figure 2* will remain consistent throughout our guide (CRIT A). It portrays a wide variety of information succinctly, using a distinct visual representation and position for each unique category of information (CRIT B). Let's go through the labelled areas of this format to discuss both the purpose of each as well as how they address the guide criteria we have set out.

The guide will use a variation of point format most closely resembling that used by Schell (2015) in his ten principles (CRIT D). Area ① shows us the heart of this format – a concise and clear thesis statement that summarises the principle at the core of each point (CRIT E, INFO A). This heading will either take the form of an Instruction/Directive (CRIT F) or that of a more general Goal (CRIT G).[29]

As shown in Area ②, each of these points will have a unique alpha-numeric value associated with them – a *Point ID*, if you will. This identifier helps keep the guide organized, but also provides the reader with some important information at a glance. Firstly, the letter at the start of the ID tells the reader whether the point is a Goal (G) or a Directive/Instruction (D). In addition to this letter, the headings of points will also be colour-coded based on their types. The headings of Goals will be green, while those of Directives/Instructions will be pink (CRIT J, INFO I).

The numerical portion of the ID shows the reader which section the point falls under. For instance, the example in *Figure 2* showcases the fifteenth point in Chapter 3 of the guide. All of the guide's points will be organized into Chapters based on the different phases of the puzzle game design process. These phases will be determined by the distinct moments in the design process outlined in Scott Kim's Eight Steps (as cited by Adams, 2015), and Selinker and Snyder's guide to Crafting a

[29] As previously discussed in the *Point Format* section, the elementary information that makes up a point will not be repeated as both a Goal and an Instruction – if both representations were plausible, framing the point as a Goal would take priority (CRIT H, CRIT I).

Videogame Puzzle (2018). The order in which these chapters are presented will be the chronological order of the broader design phases that they represent (CRIT K).

While the reader would be expected to work through each of these chapters consecutively (based on how far along they are in the design process), the order in which they tackle the individual points *within* their current chapter will be based on their discretion (CRIT L). To assist with this decision, each point will have a visual representation of its priority rating, as shown in Area ③ (CRIT M, INFO B). This priority score is based on the number of authors (from our *Foundational Literature*) and puzzle designers (from our Case Studies) that agree with the point.

Each gold star in Area ③ represents a single author or designer that supports the point, while each *black* star represents one that explicitly contests the point (INFO D, INFO E). I have also reserved a single star for myself, which I will use to slightly raise the priority of points that I found particularly useful while designing the puzzle game (CRIT N). For readers that may be interested, Area ④ will provide a specific list of authors and designers that support the point (INFO H). If I have chosen to raise a noteworthy point's priority rating by one star, I will also add my name to this list for the sake of clarity (INFO J). As most beginner designers will probably benefit more from the priority rating itself than from the specifics behind it, Area ④ has been left visually subtle to prevent it from distracting readers.

Area ⑤ contains a more detailed explanation of the thesis statement in Area ① (CRIT E, INFO C). This space offers an opportunity to clarify any ambiguities in the thesis statement, explore any interesting subpoints relating to it, or discuss the reasoning behind the point. Additionally, when a thesis statement can be broken down into more elementary instructions or goals, these smaller, more actionable directives will be presented to the reader in Area ⑥ (CRIT P, INFO F). Note that Areas ⑤ and ⑥ will be written with a consistent level of detail across all of the guide's points. This level of detail will be targeted at those with only a basic understanding of the fundamentals of game design (CRIT C), and will remain generalized enough for the information to be applied to a variety of puzzle types (CRIT O).

Finally, we come to Area ⑦, which will be a full list of the sources from which the opinions of authors and designers were obtained (INFO G). I have opted to use the IEEE referencing style for this, as it is concise and unobtrusive, which prevents the reader from being distracted by Area ⑦. This decision better meets the needs of our beginner designers, who are less likely to engage with these sources – but the information will still be made available for our more advanced readers, should they need it.

As justified throughout this chapter, the format shown in *Figure 2* is not only viable, but particularly well-suited to guiding fledgling puzzle designers. It meets all of the criteria proposed in the *Criteria Breakdown* section, and contains the majority of the information presented in the *Information Goals* section. For the sake of clarity, INFO K was ultimately omitted from the design (INFO K being a list of specific authors or designers that *Disagree* with a given point). This information is still available to the readers through the sources in Area ⑦, but it was not given its own area to avoid overwhelming the less experienced readers that we hope to cater for.

At this point, please head over to ***In Search of Eureka*** – *The Beginner's Guide to Puzzle Design in Videogames*.

You needn't study every directive in detail, but I do ask that you briefly engage with the guide as a whole. That being said, if you'd like to engage more deeply with each directive, please be my guest!

You may notice that the format of this document differs slightly from that described in this chapter. These slight changes are discussed in detail in the *Main Course* section below.

Please note that the following two chapters reference this guide on several occasions. I've tried to include some concise context for these references within this document, but you will gain a significantly better understanding of the upcoming sections if you work through them in tandem with the guide.

Main Course

The Content Analysis

In the previous chapter, we engaged with the format of our four Foundational Works extensively. This process gave us a well thought out mould to prepare our guide within – but a fancy dish is worth nought with nothing to put in it. As such, the time has come to shift our focus from *how* information is presented in these works to the information itself. In this chapter, we will discuss the process used to populate the guide with content and reflect on how effective our planned structure was.

The Meat and Potatoes

Amalgamation Process

As pointed out in the *Rationale* above, there is an abundance of information about puzzle creation available to us. Right now, however, this information is scattered across several different sources and formats. To make it more digestible, we've undertaken the endeavour of sorting a large portion of it into a more user-friendly format. To facilitate this sorting process, I've decided to split the plethora of sources into two distinct groups, namely our *Foundational Literature* and our *Dynamic Informal Sources*[30].

As briefly described in the *Methodology* section, I began my organizational journey by performing two separate comparative studies – one on each of these information groups. While the experience was slightly different for each of them, the studies were both conducted through the use of a table of common elements. These tables were based on those that Salen and Zimmerman used to compare and amalgamate the multiple definitions of the word "Game" (Salen and Zimmerman, 2004, p. 9). While the format of our tables had to be adjusted slightly to suit our specific context, this sorting process was conceptually very similar to theirs.

As pointed out by Salen and Zimmerman (2004), this method of analysis is particularly good for pattern identification and finding comparative results. As such, it neatly led to the formation of two distinct models for puzzle design in videogames, which I shall refer to as the *Foundational Design Model* and the *Informal Dynamic Design Model* from hereon out. The specific experience of creating each of these models is described in more detail in their dedicated sections below.

A Tale of two Models

There were three notable reasons that I chose to construct two separate models first, rather than just doing a single, larger-scale amalgamation. Firstly, timing had to be considered. As the guide and the videogame had to be developed simultaneously (due to time constraints), I urgently needed a foundation to base my practical work on. Running two separate sorting processes allowed me to complete the *Foundational Design Model* relatively early. While it was still rough around the edges, it could be used as a framework for my practical implementation.

Secondly, the two sets of information differed significantly in format. Tackling them separately allowed me to refine my approaches based on the specific properties of each data set. This resulted in a more efficient process that bore more usable results. Finally, the third reason, while simple, was perhaps the most prominent one. Splitting up the process helped break a *monumental* task into two

[30] For a quick reminder, *Foundational Literature* refers to more structured, academic works, while our *Dynamic Informal Sources* are the more subjective and less regulated information sources.

slightly more achievable tasks. This helped bolster my motivation and confidence, and improved my ability to create milestones and schedules.

A Note on Examples A, B, and C

I have selected three points from the final guide to act as examples for this chapter, namely Point 4.14 (**Example A**)[31], Point 2.8 (**Example B**)[32], and Point 4.1 (**Example C**)[33]. As we discuss the process of forming the guide in the subchapters to follow, these three examples will help visualize the format that information took at various points in the collation and organization process.

The Foundational Design Model

This section will briefly showcase exactly how information was collated from our *Foundational Literature* into a usable model for puzzle design. Due to the more structured nature of our foundational works, this was the more straightforward of the two models to complete

Step 1: Data Collection

I began by combing through the works of Adams (2015), Bates (2004), Schell (2015) and Selinker and Snyder (2018) for raw data. In a single document, I recorded every distinct directive, goal, and piece of advice regarding puzzle design that these authors offered. I was searching for elementary concepts – small nuggets of information that couldn't be broken into smaller ideas. These were not yet organized in any particular order, but still retained a reference to where they were taken from.

Step 2: Thematic Selection Sorting

Once this long list was complete, I performed a thematic analysis of my own design on it. I will refer to this form of analysis as *Thematic Selection Sorting*, as the procedure I followed resembles that of a Selection Sort Algorithm, which is often used for ordering data sets when programming (Gaurav, 2022). Put simply, I would select a single concept from our list, then scan through the rest of the list for any other concepts that share a similar core idea to it[34]. All of these related points would then be removed from our list and stored together in a shorter list. The advantage of this algorithm was that each subsequent 'scan' would be faster than the previous one, as our unsorted list would continually shrink in size.

Eventually, the entire data set was sorted into smaller collections of information, each grouped by some distinct idea. Below, you can see both *Example A* and *Example B* in this phase of the process.

[31] Point 4.14 – "Give struggling players the option to learn the solution."
[32] Point 2.8 – "Define the rules and mechanics of your puzzles, ensuring that they are clear, concise, understandable, and concrete."
[33] Point 4.1 – "Teach the player with gameplay."
[34] By 'sharing a similar core idea', I refer not only to other concepts that have the exact same sentiment. All points that dealt with a particular idea in some way or another would be grouped together, regardless of whether their approaches or philosophies aligned or not.

Group of points sorted with *Thematic Selection Sorting* (Example A)
Giving players the answer to puzzles: • Give struggling players the option to freely learn the answer from within the game (Schell, 2015, p. 250). • Give the player the opportunity to simply learn a puzzle's solution, especially if they cannot progress further in the game without solving said puzzle (Bates, 2004, p. 85). • A puzzle designer should not freely offer the solution to a struggling player. Rather, they should give the solver feedback on their current progress or hints to extend their interest (Selinker and Snyder, 2018, p. 8).

Group of points sorted with *Thematic Selection Sorting* (Example B)
Points pertaining to the rules of puzzles: • Puzzles should have short problem statements that explicitly lay out their rules and goal (Kim, as cited by Bates, 2004, p. 86). • It is imperative that the rules of a puzzle are clear (Bates, 2004, p.85). • If the rules or limitations of the system prevent a player from taking a specific action, this needs to be communicated to the player with as little ambiguity as possible (Bates, 2004, p.134). • Defining the rules is a critical part of puzzle design. Be explicit about the board, pieces, available moves and goal (Kim as cited by Adams, 2015, p. 12). • Select an appropriate form of motion for your videogame, then catalogue the puzzle system's internal logic (Selinker and Snyder, 2018, p. 202). • Use the advantages that computers offer to ensure that the player cannot break the rules of the system (Adams, 2015, p. 16).

Step 3: Finding a Throughline

These smaller groups each marked the starting point for a single entry into the planned Table of Common Elements. The first step to converting a group into the more useful tabular format was to find the throughline between all the points. The common idea between the points would act as a foundation for this, but our throughline needed to be more concrete.

To make it so, I'd first reframe the core idea such that it would facilitate the direct comparison of all the elementary design principles related to it. For instance, I could have reframed the core idea of *Example A* to read, "*Players should never directly be given the answer to a puzzle*." I would then consider how each of these points positioned themselves relative to this refined core idea, keeping track of how common each approach was. In our previous scenario, only one of the three elementary concepts aligns with the reframed core idea, while the other two are in direct conflict with it.

Once everything was tallied up, the core idea would need to be reworked for a final time. This version would be a concise thesis statement that aligned with the sentiment of the majority of the elementary concepts related to it. Additionally, it would be framed as either a goal or a directive that a puzzle designer could act upon. If we were to adjust our current working idea for *Example A* in this fashion, we would settle on something along the lines of "*Give struggling players the option to learn the solution.*" This new throughline now aligns with two of the three elementary concepts within the *Example A* group.

Step 4: Applying the Table Format

Once our throughline was cemented in place, we could use it to formulate a small Table of Common elements. This could pan out in one of two ways. Sometimes, when the core idea and group of elementary concepts were relatively straightforward, the sentiment of *all* the elementary concepts could be encapsulated by a single thesis statement. In these cases, our Table of Common Elements would only have a single row. This row would contain our throughline, as well as a representation of the alignment of our *Foundational Works* to said throughline. Situations like these closely resembled a single row in Salen and Zimmerman's table (2004, p. 9). The table format of *Example A* showcases this result.

Small Scale Table of Common Elements (Example A)[35]

	Schell	Bates	Adams	Selinker & Snyder
Give struggling players the option to learn the solution.	✔	✔		x

More commonly, the throughline I settled on would not capture all the nuance of the elementary points related to it. As this information was still going to go through several steps of summarization, amalgamation and synthesis, I decided against omitting this nuance for now. Instead, I framed our throughline as the heading of a sub-table, which could exist as a single entry in Salen and Zimmerman's format while still retaining the information within the elementary concepts that fell under this heading (Salen and Zimmerman, 2004, p. 9). The table format of *Example B* is a clear example of this approach.

Small Scale Table of Common Elements (Example B)

	Schell	Bates	Adams	Selinker & Snyder
Make the rules of the puzzle clear, concise, understandable and concrete.		✔	✔	✔
Puzzles should have short problem statements that explicitly lay out the rules and goal.		✔		
The rules of the puzzle must be clear.		✔	✔	
Define the rules. Be explicit about the board, pieces, moves and goal.			✔	✔
Ensure the player cannot break the puzzle's rules.			✔	

As shown in the table above, this format gives the reader an overview of the general concept, while still affording them the option to read the specifics that make up said overview. This example clearly illustrates how the alignment of a *Foundational Work* to various elementary ideas directly correlates to its alignment with the overarching throughline. Note that even though the elementary concepts are still represented in this format, they too have undergone a process of summarization. Each sub-

[35] ✔ : Agreement of Foundational Work.
x : Disagreement of Foundational Work.
Blank Space: No Opinion in Foundational Work.

point in this table is distinct – any redundancy in our initial list of elementary concepts has been reduced to a simple checkmark representing an author's alignment with an idea.

Step 5: Ordering the Table

Finally, once every group of elementary concepts had been reworked in this tabular format, all that was left was to join the tables into a single, ordered Table of Common Elements. I organized each small-scale table based on the number of *Foundational Works* that agreed with its throughline, starting with those that all four foundational works agreed upon, and ending with those that were uniquely suggested by a single work.

At the very end of the table, I placed the principles that housed conflicts with no clear resolution. There were times when an equal number of works held conflicting opinions. In these cases, I simply aligned the throughline to the more positive side of the debate and continued as normal. Chances were that these conflicts would naturally be resolved in the amalgamation and synthesis steps that would soon follow.

A Finished Foundational Model

At the end of this process, I had created a workable model for puzzle design in videogames. The model was messy and filled with all sorts of loose ends, but it was sufficient for me to begin the development of the puzzles in *Allegory*. [36]

[36] For those wondering, the referencing for the elements in the table was not lost in this process, it was simply stored in comments on the working document. This sufficed, as I was the only one who would be using the model in this state.

The Dynamic Informal Design Model

Here we'll briefly discuss how information was extracted from our plethora of *Dynamic Informal Sources* and organized into a workable model. The aim of this second model was to represent the design philosophies of our three puzzle designer case studies – Jonathan Blow, Kim Swift, and Arvi 'Hempuli' Teikari. Due to the sheer volume of information that needed to be processed, the procedure discussed in the previous section had to be augmented to improve its efficiency. These changes and the reasons behind them should be made clear within the following steps.

Step 1: Data Collection

In a similar fashion to how we started developing the *Foundational Design Model*, this model also began with a search for elementary concepts. The key difference lies in *where* I was searching for this information. I sifted through a wide variety of sources, including articles, talks, workshops, video essays, documentaries and interviews[37].

Understanding this context is important for two reasons. Firstly, there was significantly more content to work through. The list of elementary concepts collected in this step was considerably longer than that collected for the *Foundational Design Model*. This meant that any sorting processes I planned to use would be less efficient.

Secondly, the storage and presentation of information is notably different from what we've previously engaged with. In our *Foundational Literature*, each author collates their views into a single, well-organized work. This means there is relatively little redundancy, and a proportionately low chance that the author will contradict any of their own opinions. However, this is not the case for our *Dynamic Informal Sources*.

These sources contained a significant amount of repetition, and had a far larger chance of containing self-contradiction. This is only natural, as a single designer would present their opinion at several discrete points in time to numerous different audiences. This not only meant that they would have to repeat the fundamentals of their beliefs for the new listeners, it also meant that their beliefs had the opportunity to change or develop over time. In addition to this, the format in which they presented information (usually verbally, and often in an impromptu fashion) made it nigh impossible for designers to check their responses. Thus, small nuances, inconsistencies and even little mistakes were to be expected.

Regardless of all this, at the end of this initial step we once again had a long list of elementary concepts that needed to be organized into a more usable format. Any inconsistencies or redundancies would be kneaded out over the organization, amalgamation and synthesis that was still to come.

Step 2: Thematic Selection Sorting

Once again, I went on to perform the thematic analysis of my own design on our new data set. This worked in largely the same manner as it did for the *Foundational Design Model*, so I'll refrain from discussing it in depth again. The only difference worth noting is that this process took exponentially longer this time, due to the nature of the Selection Sort Algorithm (Gaurav, 2022)[38]. Using our *Thematic Selection Sorting* on this larger collection of elementary concepts clearly showed that the procedure is only truly viable on relatively small data sets – we would need to investigate more efficient algorithms to deal with particularly large collections of data. Our list of elementary concepts

[37] A complete reference list for the *Dynamic Informal Sources* can be accessed in the Bibliography.

[38] This is specifically in reference to the *time complexity* of the selection sort algorithm, which is $O(N^2)$ (Gaurav, 2022). This property of the algorithm implies that any increase in the number of distinct data points being sorted would result in an exponential increase in the number of discrete steps required to sort the data.

for the *Dynamic Informal Design Model* is nearing the size limit for what I'd consider a viable data set to apply this process to[39].

Step 3: Skipping Steps

At this point, the collection of information had all been sorted into groups, within which the elementary concepts all shared a specific and distinct idea. In the first iteration of this organizational process, we converted this general idea into a more concrete throughline. This was ultimately to finalise a list of discrete headings for the final Table of Common Elements, as well as to remove redundancies within each group. In this second iteration, I felt that this step was less necessary as these goals were already naturally being addressed by existing work and the procedures still to come.

By 'existing work', I'm specifically referring to our efforts in the third step of developing the *Foundational Design Model*, through which we had already generated a list of discrete headings that could be used in our final comparative table. During the amalgamation step that was soon to follow, the groups of elementary concepts connected to the *Dynamic Informal Design Model* could naturally be sorted under these existing headings. This would likely be more efficient than going through the effort of generating a second list of headings, which may well have ended up conflicting with the existing list during the amalgamation process.

In terms of the 'procedures still to come', I'm referring to the synthesis that the information was still to undergo. Once the information in both models is grouped under distinct headings (during the amalgamation procedure), the final groups would still need to be moulded into concise, cohesive products. This synthesis process would *have* to involve the elimination of redundancies, as some of the ideas presented by the *Foundational Literature* would naturally crop up in the *Dynamic Informal Work* as well. Since we would be collapsing repeated concepts into distinct directives at this point anyway, we might as well collate the repeated ideas in the *Dynamic Informal Design Model* then, rather than doing this process twice at two different points in time. Doing this collation process in a single step also comes with the added advantage that less nuance will be lost to summarisation in the information's final format.

Due to these potential increases in efficiency, as well as the looming time constraints of the project, the data that had been sorted with *Thematic Selection Sorting* was left in its slightly unrefined state as we moved on to the next step in the broader amalgamation process.

Step 4: Forcing the Table Format

Even though we had made a well-considered decision to skip the identification of concrete throughlines in our data set, it simply wasn't feasible to leave the data in its original point form. As the next two phases of the broader amalgamation process would involve directly comparing the content of the two separate models, the format of each data set needed to at least resemble that of the other. As such, the elementary concepts of the *Informal Dynamic Design Model* were reorganized into a simplified version of the *Foundational Design Model's* tabular format.

Reorganising a single group of elementary concepts into this rudimentary table format simply comprised of assigning the table a heading (based on the core idea that connected the points in that group), and then giving each point its own row within the table. Then, all that was left was to record

[39] For those wondering where the upper limit is, there were approximately *208 distinct elementary concepts* in the *Dynamic Informal Design Model's* list, which took a considerable amount of time to sort. I wouldn't recommend using this method on data sets larger than this.

each of the chosen puzzle designers' opinions on these elementary concepts – using the same check mark representation present in Salen and Zimmerman's Table of Common Elements (2004, p. 9)[40].

Below, one of the groups of elementary concepts that made up *Example C* is shown in two different stages of processing. Here we can see what it looked like just after being sorted using *Thematic Selection Sorting*, as well as how it was augmented to fit into the basic Table of Common Elements format.

Group of points sorted with *Thematic Selection Sorting* (Example C)

Teaching through gameplay:

- Break your gameplay into baby steps. You must teach the player almost exclusively using gameplay (Game Maker's Toolkit, 2019; Swift, 2007; Writing on Games, 2016).
- Find a way to make training the player fun (Swift, 2007).
- By over-tutorialising or telling players what to do too frequently, you stunt their ability to think freely and reduce their puzzle-solving ability in that session (Blow, ten Bosch, and Martin, 2013).
- Designing a learning puzzle experience involves identifying the elements that can be used as moments of insight and finding a way of introducing them that's neither too overwhelming nor so guided that it robs the player of the revelation (Davies, 2015; Teikari, 2019).
- Teach the player through implicit learning rather than exposition (Thekla Inc, 2016; The Game Overanalyser, 2019).
- Remember that each attempt at a puzzle reveals information about how the world works. This can be used to lead players to the truth, or illustrate other interesting things about why certain ideas don't work (Indie Cade, 2011).

Simplified Table of Common Elements (Example C)

	Blow	Swift	Teikari
Teaching through gameplay			
Break your gameplay into baby steps. You must teach the player almost exclusively using gameplay.	✔	✔	✔
Find a way to make training the player fun.		✔	
By over-tutorialising or telling players what to do too frequently, you stunt their ability to think freely and reduce their puzzle-solving ability in that session.	✔		
Designing a learning puzzle experience involves identifying the elements that can be used as moments of insight and finding a way of introducing them that's neither too overwhelming nor so guided that it robs the player of the revelation.		✔	
Teach the player through implicit learning rather than exposition.	✔	✔	✔
Remember that each attempt at a puzzle reveals information about how the world works. This can be used to lead players to the truth, or illustrate other interesting things about why certain ideas don't work.	✔		

[40] Once again, the actual references that supported the conclusions about the designer opinions would be stored in comments on the document for later use.

A Finished Dynamic Informal Design Model

Once all the groups of elementary concepts had been worked into this simplified table format, they were all stuck together into a single table, forming the final version of the *Dynamic Informal Design Model*. This name was particularly befitting, as the final product was rather rough around the edges. Due to the lack of concrete throughlines, it wasn't possible to order this table based on the number of designers that agreed upon each point. That being said, just the fact that all the points had been sorted into clear groups made the process of Amalgamating the Distinct Models significantly more efficient[41]. While various compromises had to be made for the sake of time, the final product was perfectly suited to my needs, and so allowed me to enter the final phases of content analysis.

[41] Exactly why this was the case will be explored fully in the section to follow.

Amalgamating Distinct Models

With our two separate models for puzzle design done, we had reached the final phases of the sorting process. All that was left was to join them into a single product, then rework that product into something cohesive and appliable. This section focusses on how I went about amalgamating the *Foundational Design Model* and *Informal Dynamic Design Model* into the final model that would ultimately be refined into our guide to puzzle design in videogames.

To join our two tables of common elements into a single table, I used a simplified version of the *Thematic Selection Sorting* that we initially used to create them. The process was naturally simplified by two key factors of the work we had already done. Firstly, the throughlines we identified when creating the *Foundational Design Model* were perfectly usable as distinct themes to search for in the new data set. As such, the core of this final sorting step was isolating single entries from the *Foundational Design Model*, and searching for all the information relevant to these themes in the *Dynamic Informal Design Model* (which was essentially treated as the unsorted data set from our Selection Sort algorithm).

Secondly, this process was further simplified by the fact that the *Dynamic Informal Design Model* had already been partially sorted. When searching through the elementary concepts, I could focus my search on the headings that had some connection to the theme I was searching for. This significantly reduced the time each scan through the data set took.

Below you can see an augmented version of our original Table of common Elements for *Example B*[42]. This new table has concepts from the *Dynamic Informal Design Model* integrated into it.

Amalgamated Table of Common Elements (Example B)

	Schell	Bates	Adams	S & S	Blow	Teikari	Swift
Make the rules of the puzzle clear, concise, understandable and concrete.		✓	✓	✓	✓	✓	✓
Puzzles should have short problem statements that explicitly lay out the rules and goal.		✓					
The rules of the puzzle must be clear.		✓	✓				
Define the rules. Be explicit about the board, pieces, moves and goal.			✓	✓			
Ensure the player cannot break the puzzle's rules.			✓				
Use simple or familiar concepts to create the core ruleset to make your game more approachable. Things that exist in the real world often work very well.						✓	
Keep design simple – select one core piece of gameplay and iterate on it.							✓
If the core of your game is strong, you may not need to add new, novel elements to make it more interesting. Keep the focus on the core of the gameplay.					✓		
Attempt to keep the game concise (this does not mean short). Don't pad the game with filler, repetition and red herrings.					✓		
Present the results you found **cleanly**.					✓		
Lightest Contrivance – Contrivance is proportional to yield. If two mechanics yield the same amount of interesting consequences, the simpler/more elegant					✓		

[42] Directive 2.8 – "Define the rules and mechanics of your puzzles, ensuring that they are clear, concise, understandable, and concrete."

one will feel the best (especially if both are seen together). Keep your rules simple and elegant.							
The clearer and simpler a puzzle is, the more beautiful and powerful the feeling of epiphany gained from it.					✔		
A good puzzle is concise, elegant and often minimal.					✔		
Edit out mechanics, rules and puzzles that fail to surprise, don't say anything interesting or overlap with each other.					✔		
Remove mechanics that: do not provide enough surprise, are not rich relative to others, do not fit the definition of what the game is about, create space of consequences that are not focussed, or overlap too much.					✔		
Present the results you found cleanly. Do this as naturally as possible – avoid contrivance to come closer to the truth.					✔		
Create a system and explore its consequences in a math-like way (thoughtful exploration of the unknown, looking for the simplest equations with the deepest consequences). Then, present the results in a way that lets the player discover the same truths.					✔		
The player should understand the crux/hook of your game within 5 minutes.							✔
The visuals and sounds in a game are fundamental in getting the player to understand the game better.					✔		

While this addressed the majority of the information we'd gathered, there would undoubtedly be some topics that arose from our *Dynamic Informal Works* which were not present in the *Foundational Literature*. To ensure that nothing was lost, I kept a record of all the elementary concepts that had been transferred into amalgamated tables like the one above. Once the simplified selection sort was complete, I went through the *Dynamic Informal Design Model* to find any headings that were unique to it, and so had yet to go through the amalgamation process. Using these last few straggling headings as themes, and now treating the *Foundational Design Model* as the data set, I performed one final bout of selection sorting. *Example C* was formulated from an amalgamated table found in this reversed phase of the thematic analysis.

This put the final bow on what was now our *Amalgamated Design Model*. Note that the various Tables of Common Elements that made up this model were all left separate, as each would be handled individually in the Synthesis step to follow. Additionally, they were not yet organized in any particular order, as this would be more efficiently handled once said synthesis was complete.

Divide and Conquer – A Brief, but Interesting, Tangent

It turns out that Selection Sorting was not the only algorithmic approach we used to organise this data (Gaurav, 2022). While we initially decided to create two separate models because of the timely need for a foundation for *Allegory,* as well as the notable differences in our information sources[43], this decision accidentally led us into using the Divide and Conquer sorting algorithm (GeeksforGeeks, 2023). This algorithm involves breaking up a large problem into smaller sub-problems, solving each of the simpler sub-problems individually, and then joining all the separate solutions together[44]. While this algorithm's efficiency can vary based on the way it is implemented and the data being sorted, in our case, it significantly reduced the complexity of the problem at hand, thus improving the efficiency with which we solved it.

[43] See *A Tale of two Models* for details.

[44] This is usually a recursive process, in which another sorting algorithm is applied multiple times to smaller scale data sets. In our case, we used a selection sort algorithm to organize the sub-problems, then used a significantly simplified selection sort to join the two solutions together. This second sorting process was simplified by the fact that the data had already been partially sorted.

Synthesis (Creating Guide Passages)

With our *Amalgamated Design Model* ready, we now had to refine it into a format that was actually usable – specifically, the one we discussed earlier in the *Finalised Format* section (*Figure 2* from this section has been included below again for your convenience). The first step to reformatting each amalgamated table was to settle on a final directive that encompassed all of the major elementary principles within it (① in *Figure 2*). This process was essentially identical to *Step 3* of developing the *Foundational Design Model*, but some additional steps became necessary to accommodate the increased amount of information that needed reformatting. This additional processing involved categorizing each elementary concept into one of three sub-groups – either *directly connected*, *adjacent*, or *irrelevant* to the final directive.

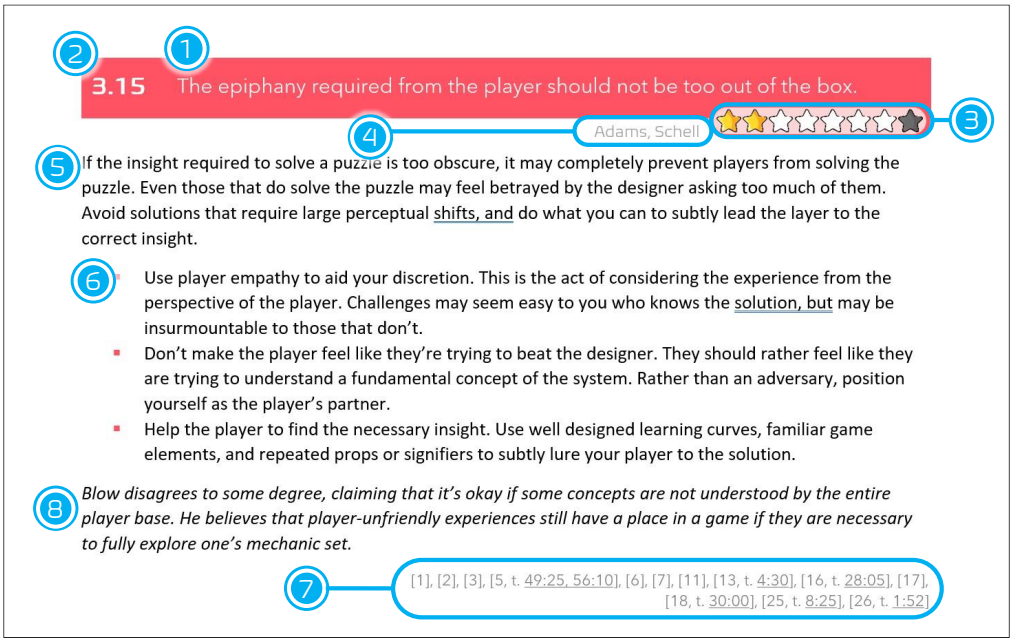

Figure 1 (again): Format of a point from the Final Guide

Directly connected concepts were those that clearly showcased an author or designer's opinion about our directive, and so were always considered when calculating the priority rating of a point in the guide (Area ③). The contents of those that aligned with the sentiment of the main directive would be reworked into a single, explanatory paragraph designed to improve the reader's understanding of the directive (Area ⑤). On the other hand, those that highlighted when an author or designer disagreed with the directive would be summarized into a brief paragraph explaining the context behind the disagreement (Area ⑧). Note that the core directive was always constructed to maximise the number of elementary concepts that were *directly connected* to it.

Adjacent concepts were those that weren't integral to the understanding of the central directive, but still felt natural under its umbrella. Some of these were more actionable instructions that helped achieve the goal set out by the directive, others offered advice on how one could more effectively integrate the directive into their puzzles, while others still would showcase the connection between this directive and other points in the guide. These were all reformatted to have strong, opening thesis statements, before being sorted based on their relevance and priority and placed in Area ⑥. These points were only considered for the priority rating (Area ③) if their underlying sentiment clearly portrayed an author or designer's views on the core directive.

Finally, the *irrelevant* concepts were those that fell into the same general sphere as the directive, but offered no contribution to its explanation or execution. These were often remnants of *Step 3* of developing the *Dynamic Informal Design Model*, and could simply be removed from this section of the guide. Once this process of categorization and reformatting was complete, our amalgamated table of

common elements was now a condensed and neatly organized collection of information, much like the example shown below. Once all the tables had been reworked, all that remained was to organize them into a cohesive, ordered document.

Synthesised Guide Passage (Example A)

4.14 Give struggling players the option to learn the solution.

If a player has become well and truly stuck on a problem, you can give them the option to simply see the solution. While this choice is controversial, it does have some notable advantages. The most prominent being that it allows your player to experience at least some semblance of the euphoria that goes hand in hand with the moment of insight. While they didn't stumble into themselves, seeing the answer will still bring them some satisfaction.

- **This must be *optional***. Not all players want help, some enjoy overcoming these hurdles themselves.
- **Save your players the trouble**. In the age of the internet, the solution to your toughest puzzles will quickly become available online if the game is released on a large enough scale. If players can access the answer anyway, one may as well save them the effort of having to look for it.

Blow and Selinker and Snyder argue against this. Blow warns that assisting the player too frequently reduces their ability to solve puzzles and prevents them from understanding the deeper meaning in your system. Selinker and Snyder suggest that much of the enjoyment of puzzle solving comes from being frustrated. Rather than simply offering the solution, one should give the solver a subtle hint or an acknowledgement of their progress to get them back on track.

Organisation

Our original plan for organizing the points of the guide was to use a healthy mix of chronological and priority-based ordering. As described in the *Order of Presentation* section, this would first involve sorting the guide's directives into chronological groups based on Scott Kim's Eight Steps to puzzle design (Kim, as cited by Adams, 2015, p. 12), and Selinker and Snyder's guide to crafting a videogame puzzle (Selinker and Snyder, 2018, p. 202). Once all the points were categorized, each group could be sorted based on the priority rating of the points within it. Unfortunately, neither of these goals worked in practice, and so a new approach became necessary. This section expands on why each of these ideas was not practically feasible, and describes the organization system I ultimately settled on for the final guide.

Issues with Categorisation

When organizing the guide's directives into the initially planned categories, three inherent issues became apparent. Firstly, the headings available in Kim (as cited by Adams, 2015) and Selinker and Snyder's (2018) work didn't appropriately house all the directives we had constructed. For instance, many of the goals surrounding documentation and planning in the first section of the guide didn't naturally fit into any of the available categories. Their inclusion felt forced no matter where I placed them.

Secondly, the weighting of each section was notably imbalanced. This probably stemmed from the issues outlined earlier in the *Consistent Level of Detail* and *Generalisability* sections, as the more vague categories like *Construct the Puzzles* ended up housing a significantly larger portion of the guide than more specific sections. While not an inherently serious problem, this did draw attention to the final issue with this method of categorization.

Basing the guide sections on the work of Kim (as cited by Adams, 2015) and Selinker and Snyder (2018) felt *arbitrary* in the context of our guide. While it was useful to split the guide into smaller portions to ensure it was sufficiently chronologically organized, it felt as if there was no valuable explanation as to why certain groups of directives were separated.

These concerns led to the decision to devise my own categorization system, which appropriately housed all the directives of the guide and had clear justification for the partitions between its sections. The four categories in this system are:

1. *Planning and Preparation*
2. *Systems Design*
3. *Creating Puzzle Levels*
4. *Experience Design*

The separation of points was now designed to allow the guide user to repeat the use of particular sections in order to accomplish specific goals or address distinct issues that may arise in their design. These goals and potential issues, as well as how to address them, are outlined in the diagram below.

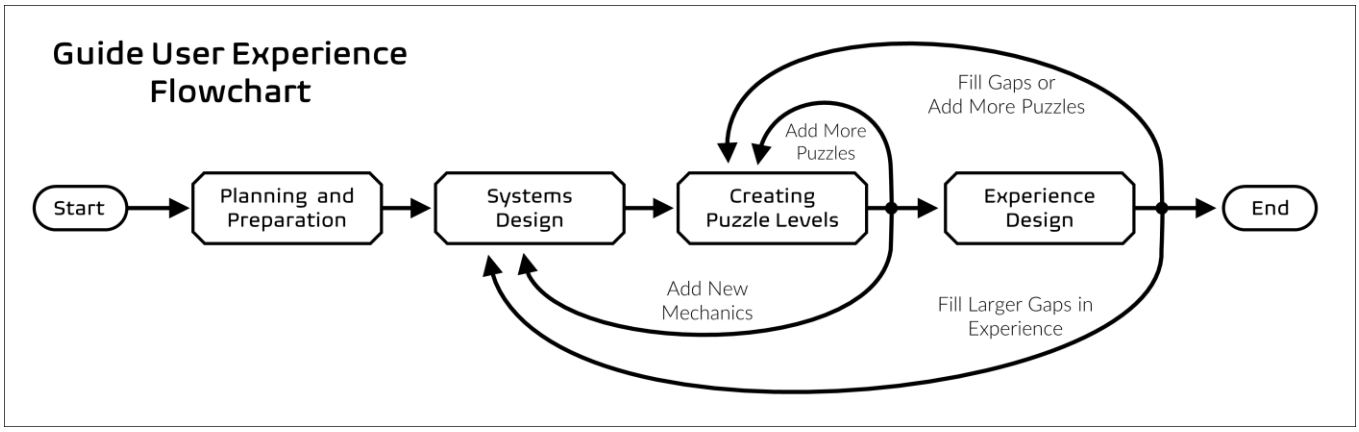

Figure 3: User Experience Flowchart for the Guide to Puzzle Design

Issues with Priority-Based Sorting

While our new categorization system was well justified, its design meant that guide directives were now sorted into relatively large groups. This heavily reduced the prominence of the chronological sorting present in the guide, which led to situations in which the priority-based sorting did more harm than good.

For instance, a directive like 2.3 (*Try and make the core idea of your game something innovative and unique*) would be placed right at the end of the section, due to its relatively low priority rating. While this makes conceptual sense, as it is not as imperative to consider this point as some of the more prominent ones, it doesn't make any practical sense, as the execution of this directive can only happen right at the beginning of this phase. By using priority sorting to organize these sections, we were indirectly forcing our users to keep the entirety of each section in mind at once. Additionally, we were leaving them with the burden of determining the optimal order in which the directives that interested them should have been implemented.

As discussed in the *Order of Presentation* section above, there is value in leaving some of this critical thinking to the user – but the larger our set of information is, the more complex this critical analysis becomes. At some point, it is no longer reasonable to expect a beginner designer to make the necessary deductions and justifications. As such, the directives in the final guide were all ordered chronologically, based on when they should be implemented in the design process.

That being said, the priority rating of each directive was still clearly communicated, which meant that the user would still hone their critical thinking skills when deciding which directives were imperative to their project. While the need for justification and understanding was less significant with this design, I still feel the current format adequately challenges a reader at a beginner's level.

Cementing the Order

Once all of the organization discussed above was complete, the collection of directives could be numbered, thus setting their order in stone. Initially, the numbering of each point in the guide would be used to imply whether it was a Goal (an identifier beginning with 'G') or a Directive (an identifier beginning with 'D'). However, after numbering the guide in this initial fashion, it became clear that less than ten percent of the guide fell into the latter category. Additionally, the contents of this group showcased that the different point types could essentially be handled in the same way. Thus, there was little practical value in differentiating between the two concepts, and so the letter was removed from the unique identifier of each point for the sake of simplicity.

The Salad

Playtesting Process

Up to this point, we've managed to justify all of our decisions pretty well by picking apart the structures of foundational texts, critically reflecting on each phase of our process as we complete it, and meticulously predicting what will prove important in the steps yet to come. While all well and good for justifying choices in an academic paper, I feel that these critical analyses and references do not sufficiently prove the effectiveness of our guide. At the end of the day, this guide was made to be handled by *people*, and so I feel that we can't draw any final conclusions on its value until we've exposed *people* to the fruit of our labour.

Thus, a playtest was designed to assess the response of players to puzzles created via the guide's instruction. This chapter will discuss the development process used to create this playtest, examine why certain design decisions were made, present a breakdown of the puzzle levels used in the testing, convey a summary of the most interesting observations found in the responses, and finally touch on how our playtesting could have been made even more valuable, had our scope been larger.

Palate Cleanser

At this point, you have the **option** of playing one of the two **Puzzle Design Playtests** submitted with this document.

Please be aware that this is *completely optional* – the content of the playtest has been adequately described throughout this chapter. That being said, actually engaging with the playtest will likely improve your understanding of the discussions to follow.

If you do wish to play the playtest experience, open up the **Playtesting.zip** file included in the submission, and choose either **Test A.zip** or **Test B.zip**. You'll find an executable in each of these folders that will start the playtest.

The survey questions have also been made available, should you wish to examine them. Take a look at **PlaytestingQuestionnaire.pdf** in the **Playtesting.zip** folder if you are interested.

Playtesting Design

The ultimate purpose of this playtest was to investigate the value and validity of the *Beginner's Guide to Puzzle Design*. To accomplish this, five puzzle levels of varying complexity were created using the guide. While valuable data could have been obtained simply by examining the response of players to these five levels, the conclusions drawn from these results would be severely limited without a control to compare them to. As such, a control group of five puzzle levels was retrospectively created to be tested alongside the original levels.

This control group was meant to represent the original five puzzle levels, had they been designed without using the contents of the guide. To simulate this, the original levels were analysed and purposefully stripped of their compliance to specific directives within the guide[45]. Note that this group of stripped-down levels *had* to be created using this retrospective process, as the playtest was

[45] The *Level Showcase* to follow explores this process with more specificity.

designed at a point where I was already intricately familiar with the contents of the guide. This understanding naturally framed the way I thought about puzzle design, making it nigh impossible to simply create levels without referencing the guide.

With the control group prepared, both collections of levels could be distributed to playtesters, whose responses to each group could be directly compared to one another. This would allow me to observe the effects of the specific directives that were stripped from each level, as well as draw more general conclusions about the value of the guide as a whole. But we weren't out of the woods just yet, as the way in which we presented these levels would have an integral effect on the results the tests would yield.

There were two key things to be aware of when giving the levels to the players. The first was the potential of unsuccessful level design having a compounding effect. Assuming that the guide truly does have a positive effect on the quality of puzzles designed with it, the stripped-down levels may end up being poorly received. The problem with this is the fact that puzzles are innately designed to create some form of understanding in the player (Game Maker's Toolkit, 2016; Writing on Games, 2016). This implies that the effects of a poorly designed puzzle are not isolated to the level in question – if the concepts within this level resurface in later levels, the player's ability to engage with them may be directly affected.

The second concept to be aware of is that of player bias. Even the subtlest of changes in presentation can drastically affect the way a player thinks about a puzzle (Blow, ten Bosch and Martin, 2013). If players were made aware of the fact that not all the puzzles involved in the test were designed equally, they may make assumptions about which category the levels they are playing fall into, which could heavily bias their answers one way or another. It goes without saying that players were not informed about the comparative nature of the test, but additional steps were taken to prevent them from noticing this nature on their own.

In an attempt to address both these potential issues, the two tests each contained a healthy portion of original and stripped-down levels. The specific distribution of level types is shown in the table below. This pattern was chosen to reduce the potential impact of poor level design on subsequent levels, and to keep the player guessing about whether the test they were doing was a control test or not (should they have pieced together the comparative structure of the test).

Level Type Distribution

	Level 1	**Level 2**	**Level 3**	**Level 4**	**Level 5**
Test A	Original	Stripped	Original	Stripped	Original
Test B	Stripped	Original	Stripped	Original	Stripped

Finally, specific care was taken in the construction of the questionnaire to ensure that the questions themselves did not imply the comparative nature of the test. The questions presented for every level were identical, and sought to imply that this was a straightforward test to assess the general quality of the game. Even though the structure of these questions appeared simple, valuable data was collected in the player responses, such as perceived difficulty, enjoyment and frustration levels, and their understanding of each level's solution[46].

[46] It's worth noting that all of these decisions were made with the insights of Directive 2.12 of the guide in mind (*Test, test, and test again. Incorporate Playtesting as a major part of your design process*).

While the results would all be self-reported, thus skewing the test into the qualitative sphere, the numerical data we were collecting would still allow for the use of some quantitative data analysis techniques. This, in combination with an examination of the open-ended responses, would hopefully give us a well-rounded understanding of the players' experiences.

Level Showcase

This section will focus on the actual puzzle design in each of the five playtest levels. It will comment on how these puzzles were conceptualized and created, how the guide assisted with this process, how they were each augmented into stripped-down versions, as well as the specific goals from the guide that each pair of levels focused on testing.

Our Visual Language

This chapter makes heavy use of Puzzle Level Diagrams to communicate the configuration and design of our playtest levels. As such, having a strong understanding of the visual language within these diagrams will ensure that you can fully engage with the content to follow. This section will briefly introduce you to the iconography used in the diagrams, as well as the gameplay mechanics associated with these representations.

Note that engaging with this chapter will be significantly easier if you have already played through *Allegory* before reading it. You will probably find that engaging with either *Test A* or *Test B* beforehand will improve your understanding even further – but this additional step is by no means necessary. All the content of these two tests is adequately described within this chapter.

The icons presented in this legend are all explained in more detail in the remainder of this section. Those that are self-explanatory have been omitted from the explanation for the sake of conciseness.

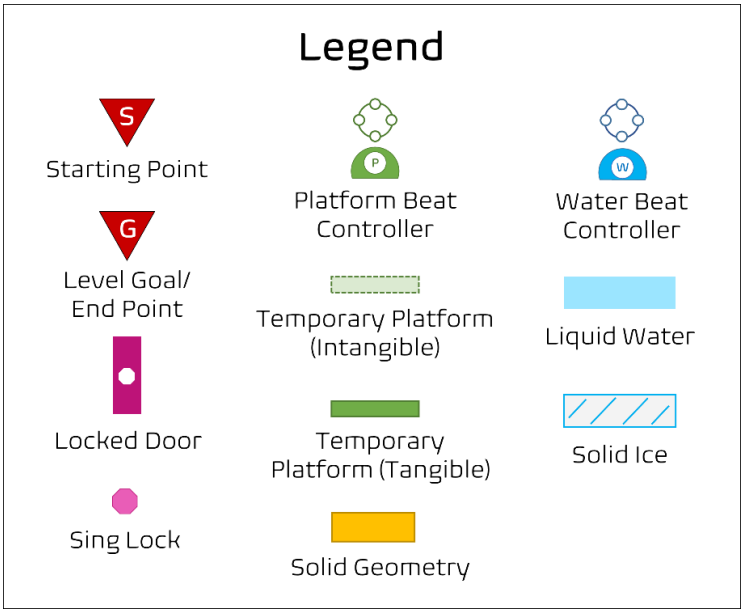

Figure 4: Legend of Icons used in Level Design Diagrams

- **Locked Door:** These are barriers that will impede the player's progression. To unlock a door, the player must interact with each of the Sing Locks connected to it. The number of holes in the icon represents the number of Sing Locks necessary to open that particular door.

- **Sing Lock:** These will be used to open Locked Doors. Sing Locks become active for a short period of time if the player *sings* when next to them. If all the Sing Locks associated with a single door are active at the same time, they will remain active permanently and the door will open.
- **Beat Controllers:** A player can *sing* next to a Beat Controller to activate and deactivate the musical beats associated with it. Active beats will play distinct sounds in harmony with the background music, while inactive beats will be skipped. Each time an active beat is played, it will cause a specific change in the surrounding environment.
- **Platform Beat Controller:** Each time an active beat associated with this Beat Controller is played, all the Temporary Platforms in the level will simultaneously change state. They will toggle between being tangible and intangible.
- **Temporary Platform (Tangible):** These are solid platforms that the player can stand on. If an active beat associated with a Platform Beat Controller is played, this Temporary Platform will become Intangible.
- **Temporary Platform (Intangible):** These are spectral platforms that the player will pass through. If an active beat associated with a Platform Beat Controller is played, this Temporary Platform will become Tangible.
- **Water Beat Controller:** Each time an active beat associated with this Beat Controller is played, all the bodies of water in a level will simultaneously change state. They will toggle between Liquid Water and Solid Ice.
- **Liquid Water:** The player can swim through liquid water, automatically rising up to the water's surface. This means that they can pass through the bottom of a water source to emerge at its top. If a player jumps into a water source, the depth they sink to will depend on how high they fall from. This means the player can pass downwards through water bodies if they jump from a high enough point.
  If an active beat associated with a Water Beat Controller is played, this Liquid Water will become Solid Ice.
- **Solid Ice:** Behaves almost identically to Solid Geometry. The player cannot pass through it and can wall-jump off of it. If an active beat associated with a Water Beat Controller is played, this Solid Ice will become Liquid Water.

With this initial understanding of the mechanics used in the levels, as well as the icons that represent them, we should be ready to discuss the design of each of the playtest levels. The subsections that follow break down the levels in pairs, directly contrasting the complete design with its stripped-down version.

Level 1: Vocal Exercises

This first level was designed to teach the player the core movement and interaction mechanics they'd need throughout the rest of the playtest. The fundamental mechanics the player will learn here include moving, jumping, activating Sing Locks, opening Locked Doors, and floating. Both activating Sing Locks and Floating are accomplished by using the Sing mechanic in the correct way (when next to a Sing Lock or Airborne, respectively). This simple tutorial level finished off with a relatively easy puzzle (shown in the figure below), designed to teach the floating mechanic.

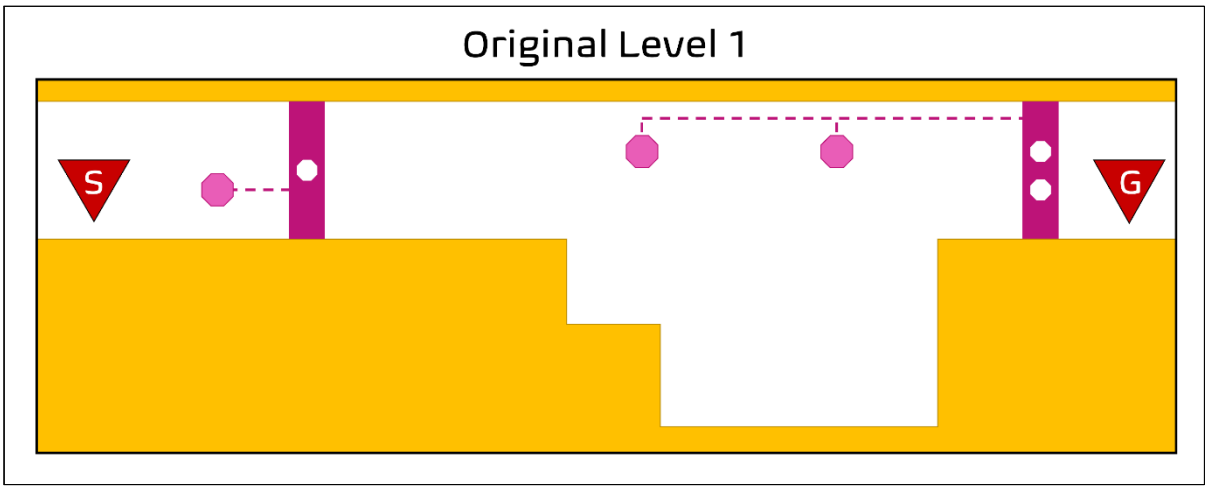

Figure 5: Original version of Level 1.

As this level is dedicated to teaching basic mechanics, it would probably be fair to simply offer the player a collection of direct tutorials. That being said, various points of the guide show that there is potentially a more valuable way of going about this. The directives that had the most prominent effect on the design seen here were points **4.1** (*Teach the player with gameplay*) and **3.10** (*Guide the player through the level towards the solution*).

Point 4.1 essentially implies that as much of our tutorialisation as possible should be *implicit*, rather than explicit (Adams, 2015; Davies, 2015; Game Maker's Toolkit, 2016; Game Maker's Toolkit, 2019; Swift, 2007; The Game Overanalyser, 2019; Writing on Games, 2016). This will theoretically make the learning process more engaging and improve the player's ability to recall taught mechanics (Blow, ten Bosch and Martin, 2013; Swift, 2007). As such, the tutorial for the floating mechanic was designed to teach via the moment of insight of an undemanding puzzle (see diagram above). [47]

In the puzzle above, the player is taught (through the locked door on the left) that they must sing when near Sing Locks to interact with them. As soon as they exit this door, they are shown another Sing Lock. But this time, the lock cannot be reached from a standing position. The natural reaction should be to jump towards the lock and sing – which organically shows off the floating mechanic. From here, point 3.4 comes into play.

The crux of this directive lies in the idea that we designers should draw attention to the path we want the player to take through our levels (Bates, 2004; Game Maker's Toolkit, 2019; Menzel, 2016; Selinker and Snyder, 2018). This level attempts to do this by incorporating some of the specific suggestions related to point 3.4, including clearly highlighting elements in the environment that are relevant to the puzzle (Bates, 2004; Menzel, 2016), breaking up the challenge into smaller problems (Game Maker's Toolkit, 2019), and cleverly using patterns and repetition to direct the player (Game Maker's Toolkit, 2016; Menzel, 2016; Selinker and Snyder, 2018; Talk and Dev, 2020).

To cut a long story short, the positions of the two right-most Sing Locks not only entice the player into floating naturally, they also point the player in the direction they should travel once they've understood this mechanic. This concept is finally cemented by the Locked Door on the far right, which the player is led to assume must both be opened and travelled through.

These two points from the guide were the focus of our testing in level 1. In order to ascertain how much value they added to the design of this puzzle, they were removed from the stripped-down version of the level. As the pair of Sing Locks and the Locked Door on the right of the level were all

[47] For more details on any of the guide points being referenced here, please take a look at the point in the guide itself.

imperative to the incorporation of points 4.1 and 3.10, removing these points was as simple as deleting these elements from the level. The result of this process is shown in the figure below.

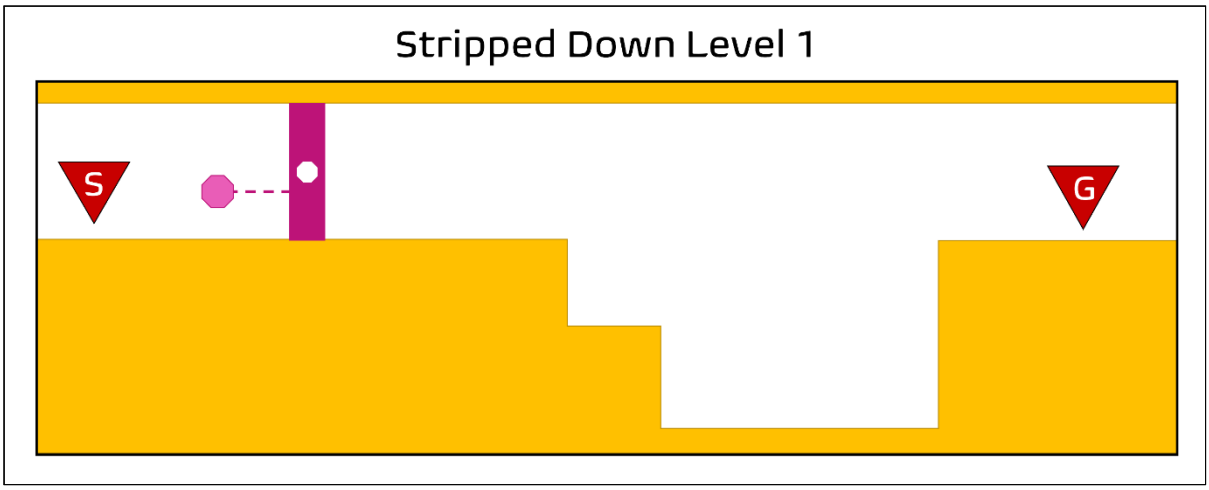

Figure 6: Stripped-Down version of Level 1.

By using the Original and Stripped-Down versions of this level in Tests A and B respectively, the playtester responses to each version of level 1 can be directly compared to one another. Any differences in how these two versions are received by playtesters will have clear implications on the effect that these two guide directives have on puzzle quality.

Note that while directives 4.1 and 3.10 were the *most prominent* guide points used in this design, they were not the only ones at play. The following table shows various other directives that were clearly evident in this level, while also highlighting how important each of them was to the playtest. In addition to this, the table shows exactly which directives were removed from the stripped-down level – these are indicated by the grey 'cancel' symbol behind the point numbers in question.

Notable	Core Focus	Negligible
~~3.7~~ ~~3.12~~ ~~4.2~~ 4.8	~~3.10~~ ~~4.1~~	3.4 3.9

Table 1: Guide directives that had a noteworthy presence in Level 1.

As pointed out by the table, other particularly notable points in this level were **3.7** (*Make the goal of the puzzle easily understood and explicit*), **3.12** (*A puzzle should clearly communicate when it is solved*), **4.2** (*Use repetition wisely*) and **4.8** (*Accommodate the introduction of new mechanics and ideas with specifically designed puzzles*). For clarity, both 3.7 and 3.12 were addressed by the door at the puzzle's end, which both gave the player a clear goal to reach as well as an unambiguous signal of the puzzle's completion. While very clearly present in the level, these directives didn't quite tick all the necessary boxes to be part of the level's core focus.

Generally speaking, the core focus would be made up of points that could be directly tested, either by looking for specific signals in the open-ended playtester responses, or by removing them from the stripped-down version of the puzzle. There were several reasons that specific directives may not have met the necessary criteria, the most common of which are listed below.

- Points may have been too integral to the experience, and thus couldn't be comfortably removed (points 4.8 and 3.4 in this case).
- Points may not have been unique to a level, resulting in them being tested across multiple levels (points 3.7 and 3.12).
- There may be another level in the playtest that was better suited to testing a specific point directly (point 4.2).
- The points chosen as the core focus may have had such a prominent effect on the observed results that the effect of other noteworthy points may not have been clearly observable (point 3.9).
- Points may have been ingrained in the design of the system, which couldn't be adjusted at the time of the playtest's creation (this is why so few points from section 2 are highlighted in this chapter).

Finally, I should expand on the difference between a Notable point and a Negligible one. A point would be marked as Negligible if there was no clear way of observing its effects in the testing results. This would occur if the effects were outshined by a particularly prominent point with similar effects (as is the case with point 3.9 here), or if there was simply no effect on the value of this point caused by the changes made in the stripped-down level (point 3.4 in this case). Regardless of the reason, there would be little chance of discerning any information about the effectiveness of these points from the levels in which they are flagged as Negligible.

Since this was the first of the five level analyses, I've gone into a significant amount of depth breaking down the visual representation system and playtest design, and explaining the effects of various guide directives on the level. However, due to space limitations, I will not be able to go into such depth in the four level subsections to follow. That being said, I will still give you a sufficient amount of information to extrapolate a similar in-depth analysis yourself. To assist with this process, I will also include additional diagrams (such as the one seen below) that highlight which areas within the level are most important to the specific guide points being tested.

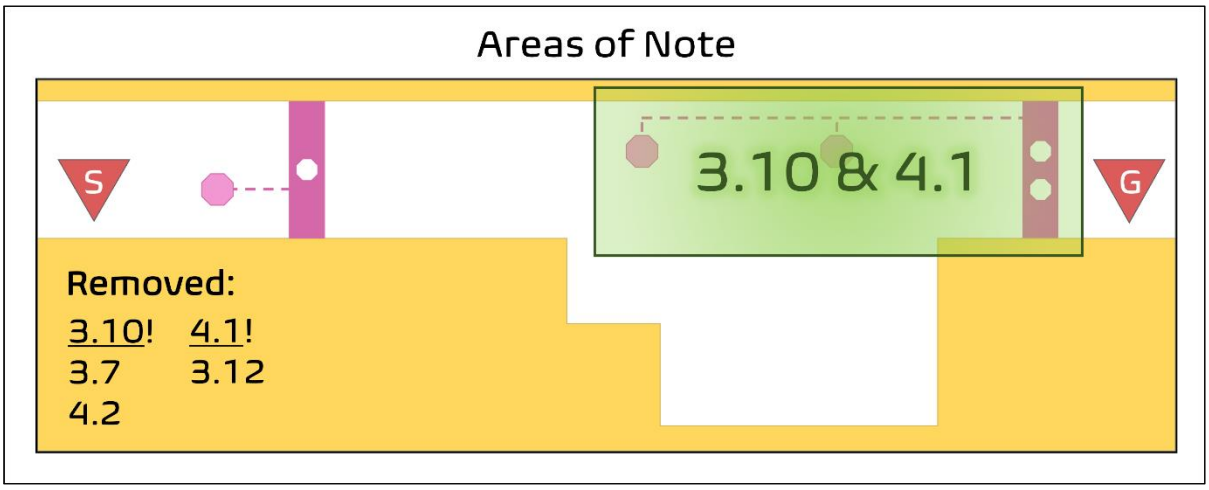

Figure 7: Important focal points in Level 1.

Here you can clearly see that the Sing Lock and Locked Door system in the top right was imperative to the functioning of points 3.10 and 4.1. These diagrams will also give you a quick representation of all the points that were removed in the stripping process (as seen in the bottom left here). Note that the underlined points are part of the core focus, while the rest are the other notable directives that were removed.

Level 2: Hard to Beat

The main goal of Level 2 was to teach the player how to interact with beat controllers. In a larger product (such as the full version of *Allegory*), I would have preferred to expose the player to these concepts more gradually, but the scope of the playtest demanded it be taught quickly. This limitation, however, was turned into an opportunity to compare the value of explicit and implicit tutorialisation.

In the original version of Level 2, the player is expected to figure out the workings of beat controllers on their own, having only subtle contextual clues to go on. In the stripped-down version, points 4.1 and 4.2 are replaced using an explicit tutorial message. Both versions of the level can be seen below.

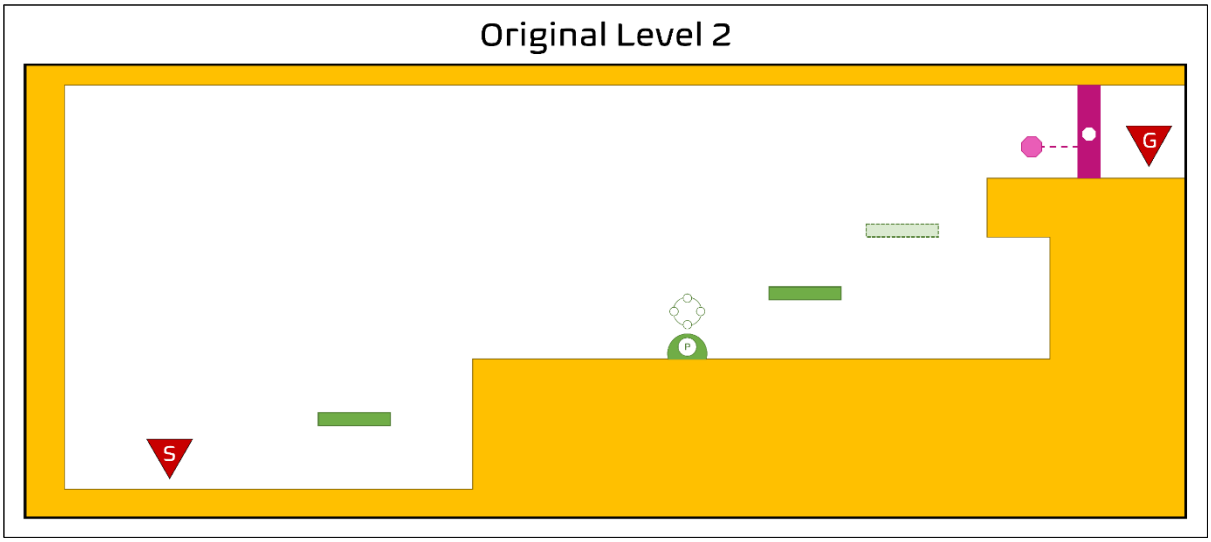

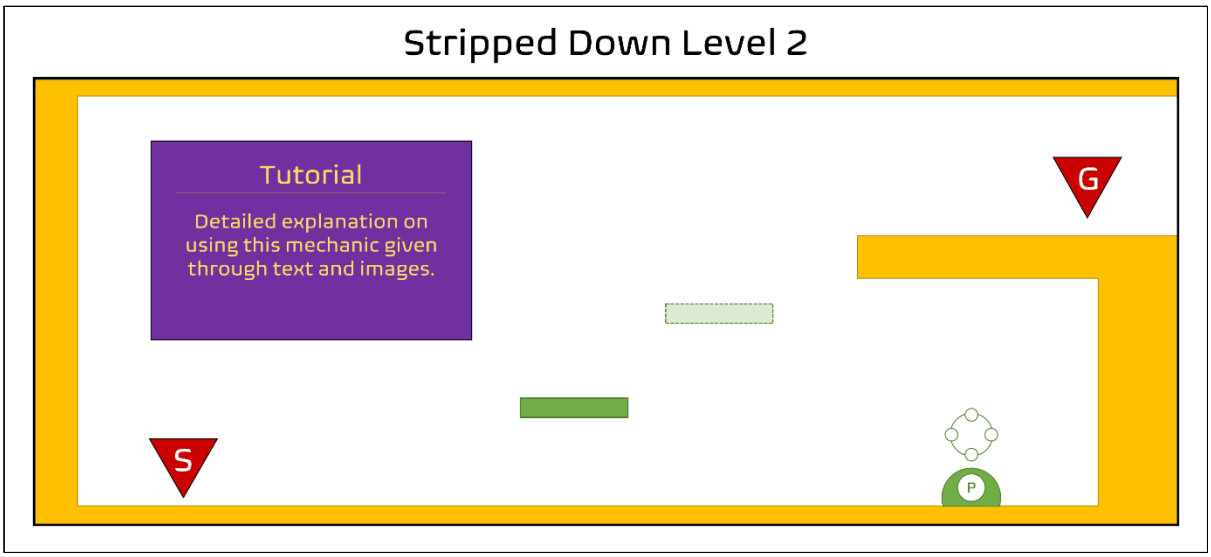

In the same fashion as before, the following table shows all the directives that are somewhat relevant to this level, highlighting both their relative prominence and whether or not they remained after the stripping process.

Notable	Core Focus	Negligible
~~3.7~~ ~~3.10~~	~~4.1~~ ~~4.2~~	2.6 3.4

Finally, the following figure once again shows the areas in the level that were most heavily affected by the stripping process. Here we can see that the player is originally asked to use a Temporary Platform to overcome an obstacle, using repetition to imply the purpose of these platforms. This subtle clue was replaced with an explicit explanation in the Stripped-Down level. Additionally, a door was used in the Original level to signal the end point to the player.

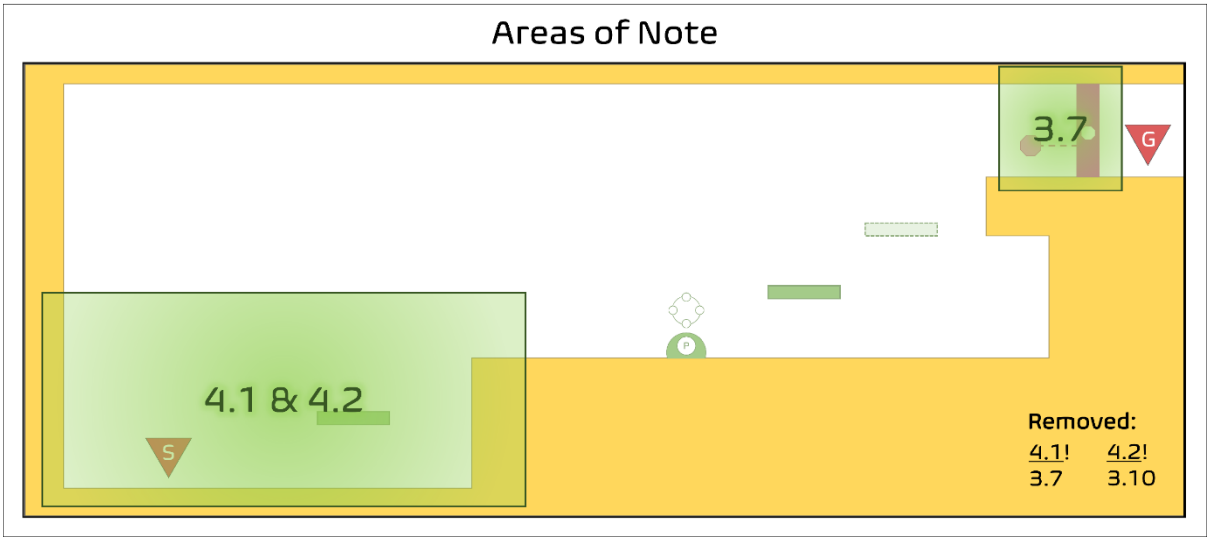

Level 3: Swimming Lessons

I used the Reverse Engineering[48] process to design this particular level, starting the level's creation from a specific moment of insight that I wanted to capture (Adams, 2015; Blow, ten Bosch and Martin, 2013; Couture, 2018; Game Maker's Toolkit, 2019; Gamelab Conference, 2019; Selinker and Snyder, 2018; Talk & Dev, 2020; The Game Overanalyser, 2019). The epiphany I wanted to capture was that the player could squeeze through low gaps by swimming under them. Since the conceptual leap was relatively simple, I felt this puzzle was well suited to teaching the player the properties of Bodies of Water, as well as the use of Water Beat Controllers. Below you can see the Original version of Level 3.

[48] See directive 3.6 in the Guide.

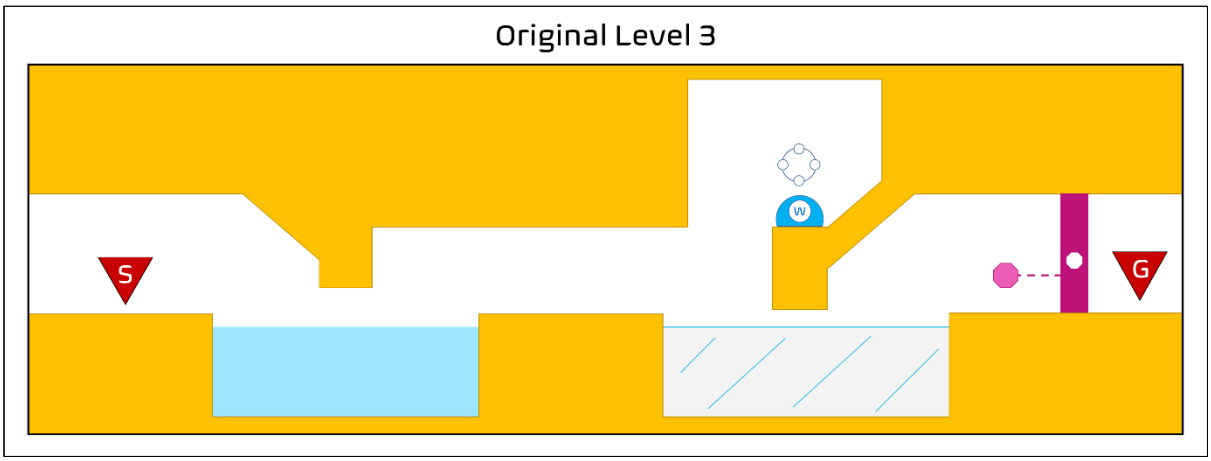

In order to teach the player that water is safe, I force them to swim through the first body of water using a pillar hanging from the ceiling. I then repeat a similar structure in the second half of the level (playing on directive 4.2), but this time the player must realise that they can turn the ice into water to progress in the same fashion. In the stripped-down version (shown below), I've removed the repeating structure, and introduced the player to the Beat Controller before they've encountered the crux of the problem, making it more likely that they solve the problem accidentally (see directive 3.16).

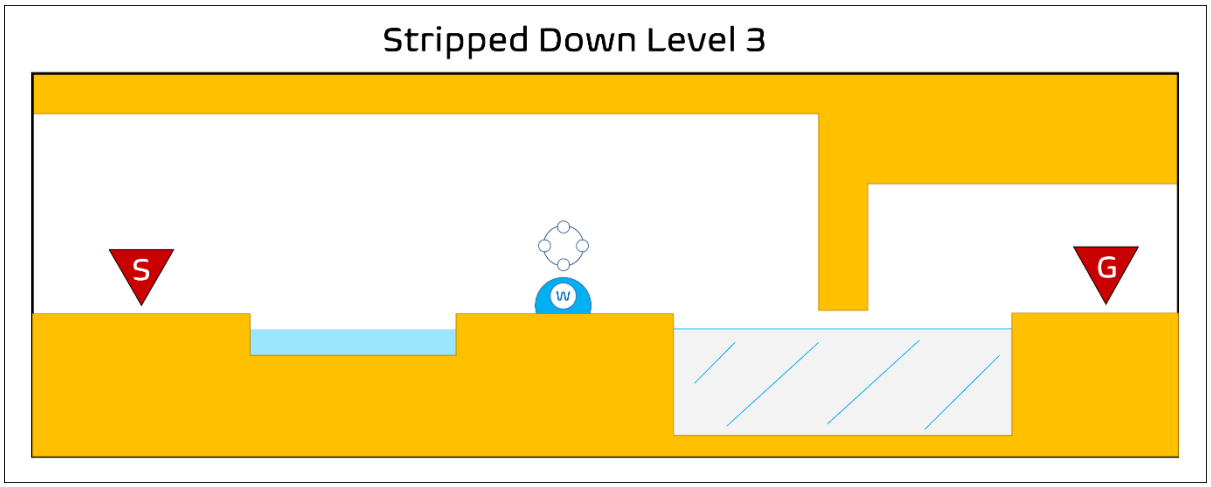

The following table shows all the guide points that played a noticeable part in this level, as well as which of these were removed via the stripping process.

Notable	Core Focus	Negligible
2.9 ~~3.7~~ ~~3.14~~	2.14 ~~3.16~~ ~~4.2~~	3.4 3.17 4.1

In the supplementary diagram below, you can clearly see the previously mentioned repeated structure, as well as how the Beat Controller was repositioned to reduce the likelihood of accidental solutions. Once again, a Locked Door was originally used to signal the end of the level to the player.

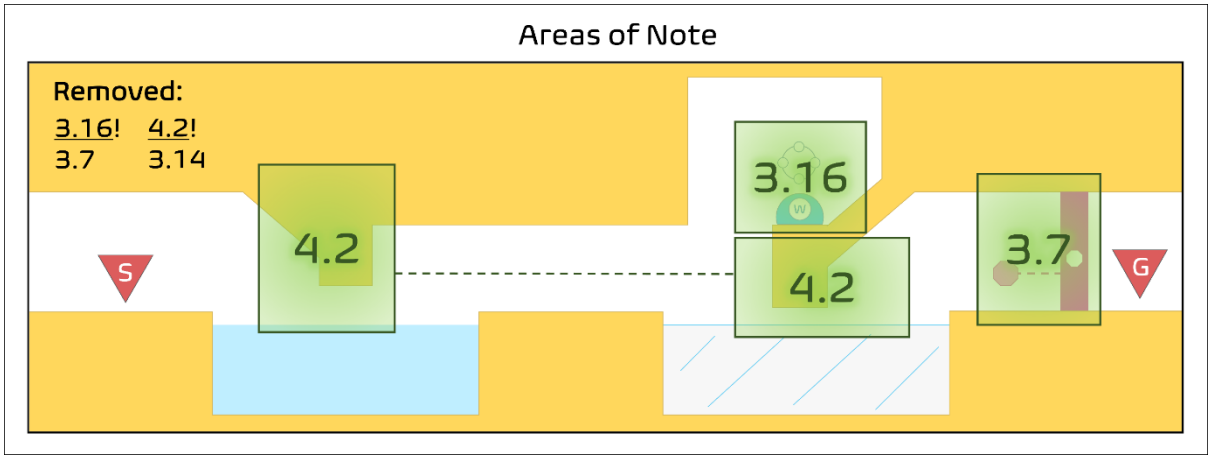

Level 4: Leap of Faith

This level once again used the Reverse Engineering model[49] to kick off the design process (ibid.). The epiphany I wanted to focus on was the fact that the player could pass downward through thick bodies of water, provided they jumped from a high enough point. The level generated from this (shown below) housed a Sing Lock hidden beneath a prominent body of water. I made sure that the existence of said Sing Lock was made clear by forcing the player to pass it on their way into the main puzzle (directives 2.18 and 3.10), and used repetition to teach players how to dive through water (directive 4.2).

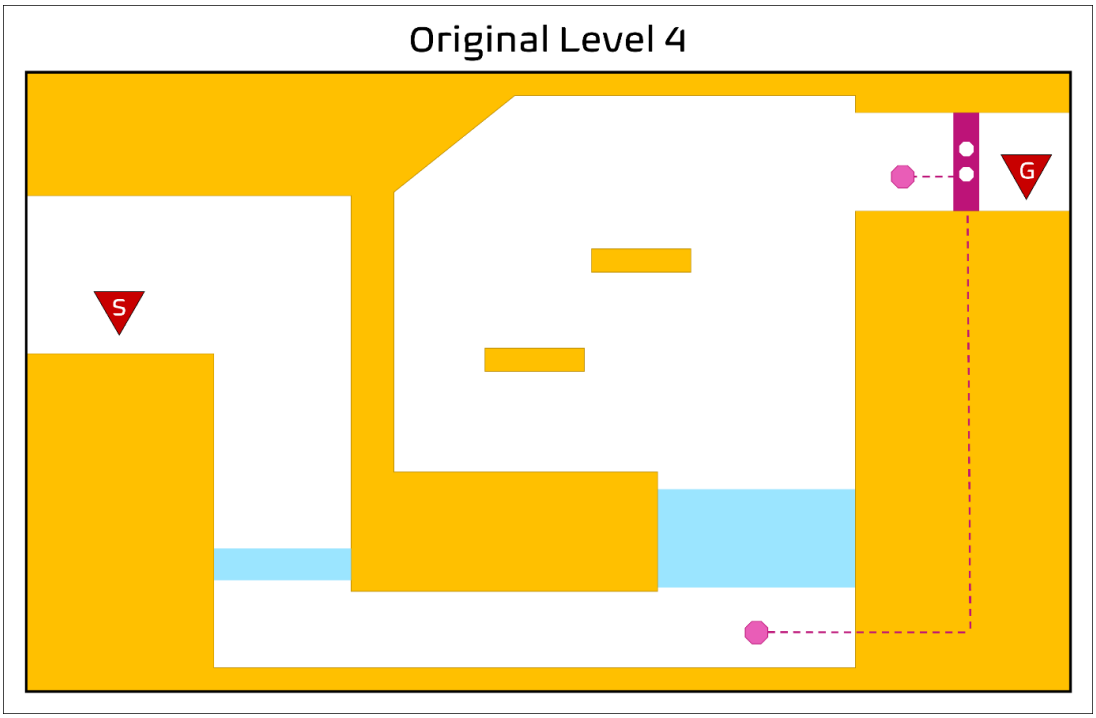

To strip this level of its compliance with the guide, I removed the repetition of the dive and added in further complexity to the puzzle through the use of a Water Beat Controller. As shown in the diagram below, the player would now have to make the water in the level intermittently turn to ice, allowing

[49] Directive 3.6 in the guide.

them to wall jump off the frozen waterfall up to the higher platform[50]. They would then need to time their dive with the moment the Body of Water turned back into a liquid. This added various layers of unnecessary steps to the puzzle, obscuring the epiphany.

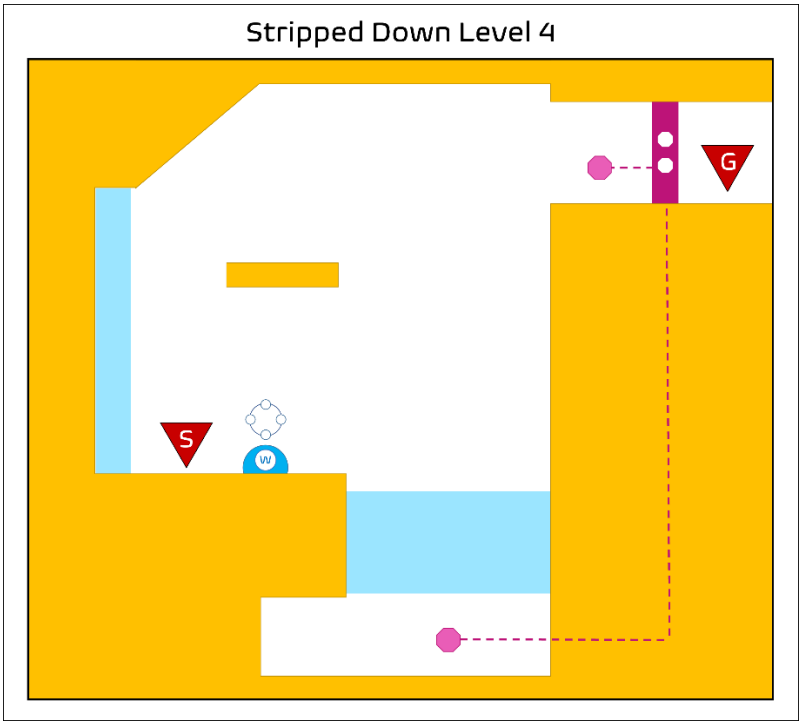

All of the guide directives that were relevant to this level are shown in the table below, as are the indications of those that were removed during the stripping process.

Notable	Core Focus	Negligible
~~2.18~~ 3.10 3.16	~~3.4~~ ~~3.17~~ ~~4.2~~	~~2.6~~ 4.1 ~~3.19~~

Finally, the supplementary diagram below clearly showcases the use of repetition and the guided path through the level that were included in its original form, as well as the layers of complexity that were added on to the stripped-down version.

[50] Note that the player in unable to wall jump when in a waterfall, such as the one seen on the left of the level.

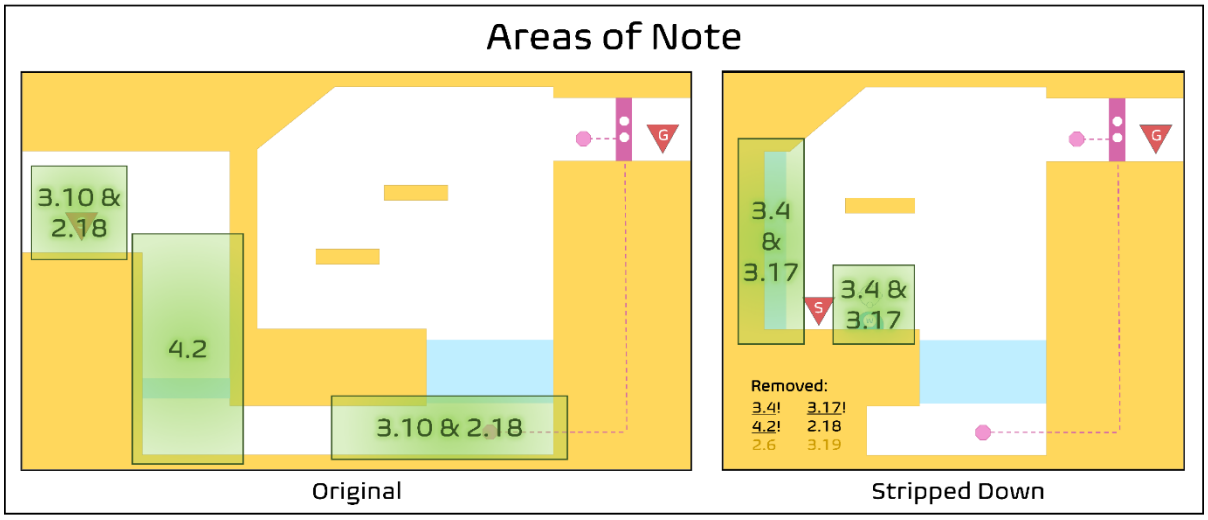

Level 5: The Structural Integrity of Water may not be Suitable for Bridge Construction

Yet again, this level was born from the process of Reverse Engineering (ibid.). This time, the epiphany in question was the corollary of that learnt in the previous level – if the player is forced to fall from a height, they'd have no means of stopping themselves from falling through a thin Body of Liquid Water. Thus, the player would need to turn the water into Solid Ice to create a stable place to land. The diagram below shows how this concept was turned into a puzzle level.

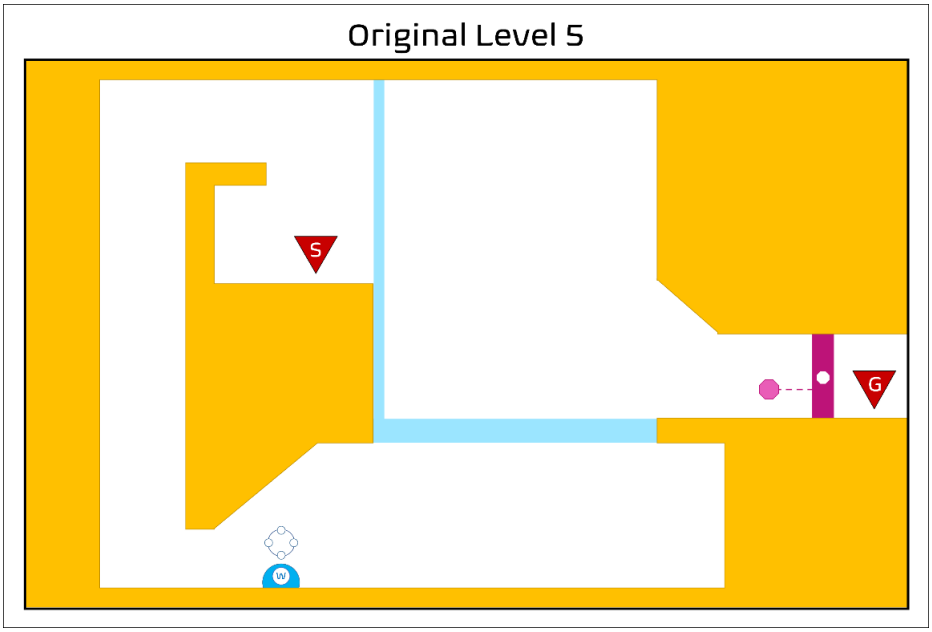

Here, the player starts at such a point that the only way to progress is to fall into the body of water. Because the player cannot sing for a moment after passing under a waterfall, they cannot engage the Float mechanic, and so are guaranteed to fall with enough speed to break through the water. This immediately introduces them to the crux of the level (directives 3.4, 3.7, and 3.8), and leaves them in a position that naturally filters them towards the Water Beat Controller, which is the core element of the puzzle's solution (directive 3.10).

In the stripped version of the level (shown below), the player's starting position does little to direct them in the space. Additionally, the player is given access to the Beat Controller before they are shown

the crux of the problem at hand. This should make it harder for the player to get started on the puzzle and increase the size of the conceptual leap necessary to solve the puzzle.

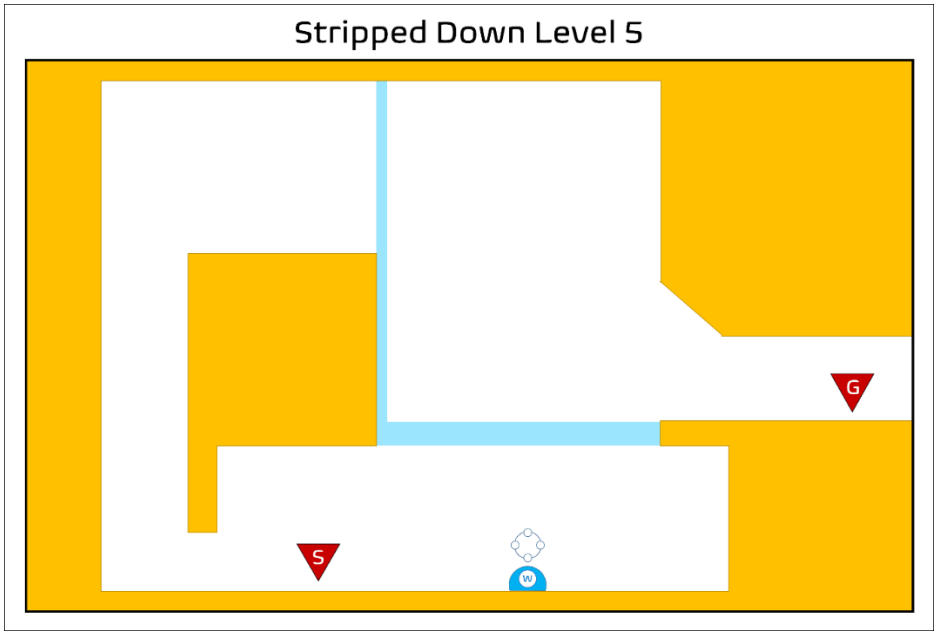

As it has before, the following table highlights which points from the guide were involved in these test levels, and indicates which of them were removed in the stripping process.

Notable		Core Focus		Negligible		
3.4	~~3.7~~	~~3.8~~	~~3.10~~	2.14	~~3.14~~	4.1

Finally, this supplementary diagram clearly outlines how the geometry of the level was adjusted to create the stripped-down version of the level. Pay special attention to how the player's starting position was changed.

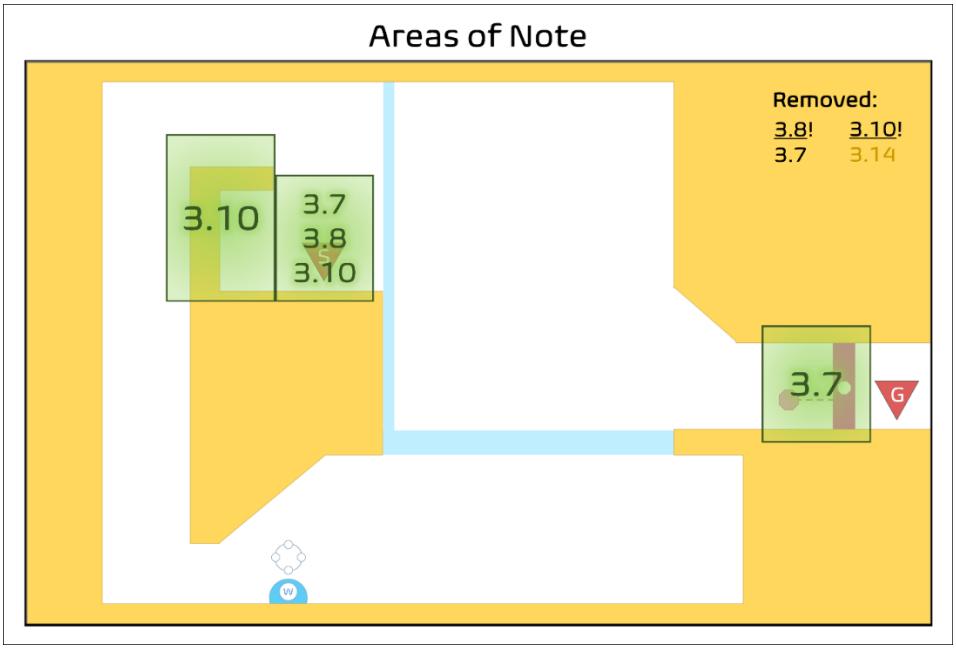

Dinner Conversation

Testing Observations

The playtest described above was distributed to a group of 28 voluntary playtesters. Both versions of the test were run simultaneously, with half of the group being given Test A and the other half Test B. Of these 28, only 13 submitted a response, but this data was still sufficient to draw some useful conclusions about the featured guide directives, as well as the usefulness of the guide as a whole.

This section discusses the most notable trends in playtester responses across all of the levels. At the point of release, I had hypothesized that the results would show a *general decrease in frustration and difficulty*, as well as an *increase in enjoyment*, in the *Original Levels*. Let's begin by looking at the general trends of difficulty, frustration and enjoyment over the course of the playtest, as showcased by the diagrams to follow.

Important Note Regarding Level 5

As you examine the data to follow, please bear in mind that three of the playtesters were unable to play Level 5 in their respective tests. Two of these testers were simply unable to complete Level 4 (Test B), while the third experienced an unexplained game crash that they couldn't work around (Test A). As such, all the data shown in connection to Level 5 has been calculated using only the responses of the ten playtesters that actually attempted it.

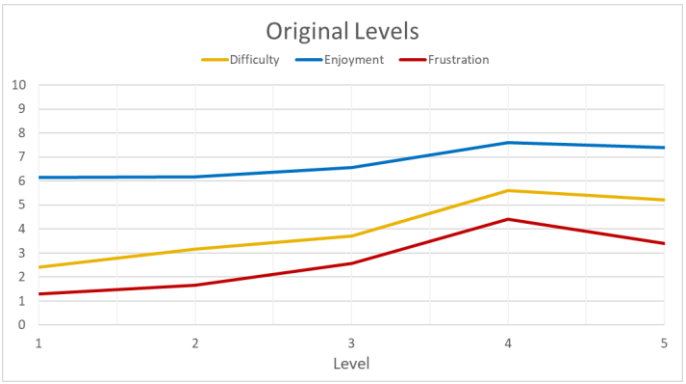

Figure 8: Experience over Playtest in Original Levels

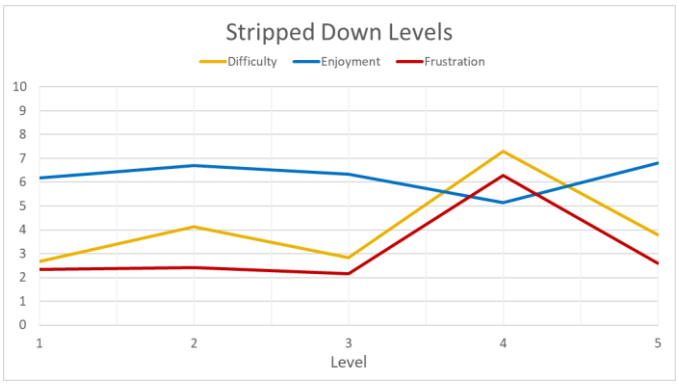

Figure 9: Experience over Playtest in Stripped Levels

It's clear from the Original Levels' graph above that there is a trend for the three highlighted properties to rise, in general, over the course of the experience. There is also a clear correlation between the difficulty, frustration and enjoyment throughout the original levels. The same correlations and trends can be seen in the stripped-down levels, with the exception of the clear outlier, Level 4. Through some critical analysis, we can delve deeper into where these patterns may be coming from, and draw some conclusions about what they may imply.

The general trend for the difficulty to rise over time was to be expected. The order of the playtest was purposefully arranged such that the complexity of the levels increased over time. The only exception to this design was Level 4, which was identified as the most difficult level. I decided to place Level 4 before Level 5 to better test the effect of directive 4.2 (*Use Repetition Wisely*) within it[51]. To some degree, the correlation between difficulty and frustration was also to be expected. What could not be predicted was the strong correlation between enjoyment and the other two properties.

This correlation was particularly interesting, as much of the research we've compiled on puzzle design insinuates that one should avoid frustration in their design to as large a degree as possible (see guide

[51] In Level 5, the player is naturally exposed to the idea of passing through water sources from above. Seeing this before Level 4 would severely affect the results that could be obtained by removing the repeated water interactions.

directives 2.15, 2.19, 3.10, 3.11, 4.7, and 4.11), and that there is no correlation between difficulty and the quality of a puzzle (see guide directives 2.12, 2.13, 2.15, 3.1, 3.16, 3.17, 4.5, 4.6, 4.7, and 4.9). Seeing these results has improved my understanding of the nuance behind this information.

Firstly, with regard to difficulty, I suspect that the correlation to enjoyment may stem from the challenges being appropriately matched to the skill level of players. The first three levels of this playtest were particularly simple, as they were designed to teach the base mechanics of the system. I suspect that these levels weren't quite challenging enough to be engaging for this specific set of players.

I also believe that the type of difficulty matters. In the stripped-down version of Level 2, we can see a clear increase in average difficulty that is not reflected in the frustration and enjoyment. Given that Level 2 is where the first Beat Controller is introduced, this difficulty may have come from learning to use the system. In the stripped-down level, players are taught to use the Beat Controller with an explicit tutorial, which may have framed this system as a tool to engage with the puzzle, rather than an integral step in the solution itself. In the original version, where learning to edit the beats is framed as part of the puzzle, there is still a strong correlation between the three properties. This contrast in responses supports the idea that not all difficulty adjustments will have a valuable effect on player experience.

Regarding the correlation between frustration and enjoyment, I suspect that this implies that there are also both valuable and detrimental types of frustration. Frustration brought on by a meaningful challenge can add value and enjoyment to the experience, as the moment of catharsis is enhanced by it. This must be the frustration that Selinker and Snyder refer to when they describe "The puzzlemaker's most important skill" (2018, p. 8). On the other hand, frustration caused by friction and clunkiness in the system will likely reduce player enjoyment overall. I suspect the latter form is the focus of the guide directives that call for the avoidance of frustration.

Finally, I feel that these figures imply that some form of threshold exists, at which the frustration or difficulty of a puzzle is greater than what the player is willing to engage with. At this point, unlike the most common trend seen in these diagrams, any increase in difficult or frustration will likely further reduce the enjoyment of the player. I suspect that this is what caused the inversion seen in the stripped-down version of Level 4, which was notably more challenging (both conceptually and physically) than its original counterpart[52]. There is also an implication that this threshold is different for every player, as three of the seven playtesters that engaged with the stripped-down Level 4 gave it a reasonably high score for enjoyment.

Looking at these results in contrast also yields some valuable insights. The following diagrams focus on highlighting the differences in the responses to the Original and Stripped-Down levels.

[52] To the extent that two of the playtesters were unable to solve the puzzle.

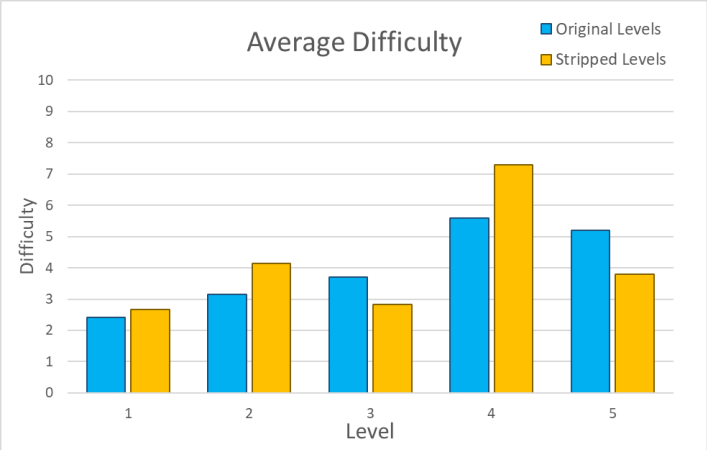

Figure 10: Average Difficulty in each Playtest Level

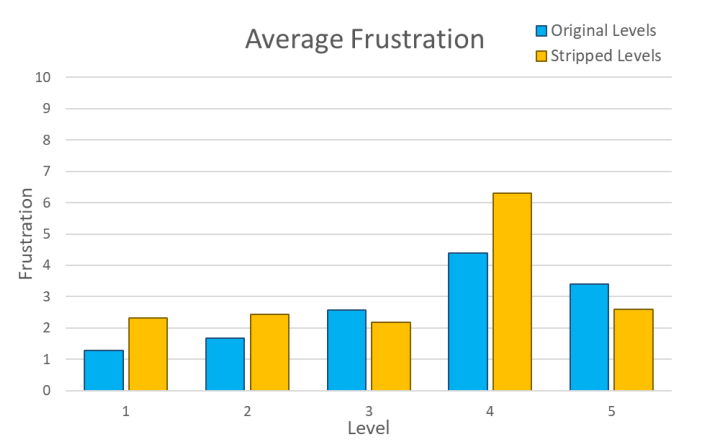

Figure 11: Average Frustration in each Playtest Level

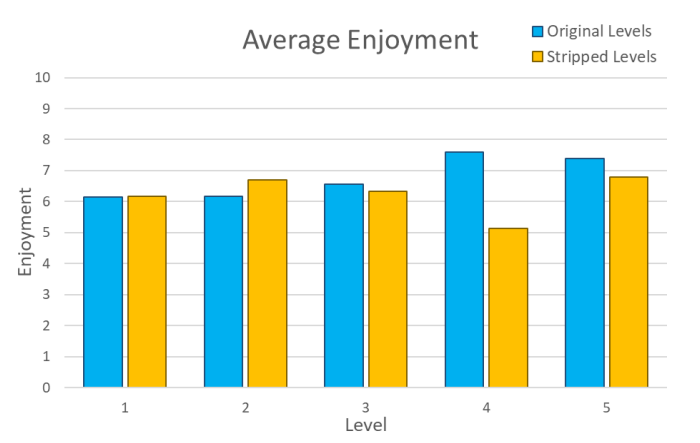

Figure 12: Average Enjoyment in each Playtest Level

In Levels 1, 2 and 4, there is a clear increase in difficulty and frustration in the stripped-down version of each level. Given the changes that were made, this is not unexpected. A more interesting contrast can be seen in Levels 3 and 5, in which the difficulty and frustration noticeably decrease in the stripped-down levels. I suspect this has to do with a common thread between the levels, related to guide directive 3.16 (*Be wary of brute force and accidental solutions*). In both stripped-down levels, the player is exposed to the Beat Controller before they have understood the crux of the puzzle. This implies that they could accidentally solve the puzzle by entering a random beat pattern, resulting in them finishing the level without completely understanding how they overcame it.

Another observation worth noting is that the enjoyment in these two stripped-down levels also decreased on average, reinforcing the earlier observed correlations between difficulty, frustration and enjoyment. This behaviour is also visible in the results of Levels 1 and 2, for which the more challenging and frustrating stripped-down versions yielded slightly higher levels of enjoyment.

Testing Shortcomings

While there were several interesting trends and correlations in the data obtained from playtesting, it is difficult to make any concrete conclusions based on these. This is because of various limitations that affected the scope and quality of the playtesting process. This section will briefly discuss these limitations, as well as postulate on some more effective means of testing our guide, had the scope of the project been larger.

As the purpose of this guide was to give beginner puzzle designers a user-friendly and approachable means of entering the sphere, the most effective testing methodology would have involved distributing it to a group of people interested in puzzle design. Armed with the guide, these designers could have created a series of puzzles. By assessing the quality of the puzzles created, along with the experience of those that used the guide, we'd have been able to draw conclusions that were more specific to the actual use case of the product.

As this was not feasible in the scope of this project, I took on the burden of the design myself. But even then, there was only so much design work I could perform in the time available. To truly understand the value of each individual guide directive, a specific control test would need to be performed for every point in the guide. This would involve creating a series of puzzles with said guide point, and testing them in contrast to nearly identical puzzles in which the focal directive had been removed. Essentially, it would be performing our entire testing process once for each distinct point in the guide – an undertaking clearly too large to fall within the extent of this project.

These are mere postulations about what could be accomplished if we had an infinite scope and significantly more time. Regarding the approach that was actually taken, there were two main problem areas. The first comes down to the size of the test, which was relatively small both in terms of content and in terms of sample size. With only five levels in the playtest, much of the time players spent engaging with the experience was used for tutorialisation. While these results are still useful, it is hard to say if they mirror what we would see with a full array of complex puzzle levels. In addition to this, only thirteen individuals responded to the playtest. With a group this small, it is unreasonable to conclude that any quantitative analysis showcases the average player experience accurately.

Outside of logistical limitations, the very design of the levels could have been improved. Firstly, within our specific context, it was difficult to test any directives relating to Systems Design, as the gameplay had already been cemented for *Allegory*, the final product. In addition to this, the directives being tested in each level were selected *after* the original version of the level was completed. A more effective process would have been to highlight the directives to be tested in each level and design a specific configuration for testing said directives. This would have afforded more control over which directives were ultimately examined in detail.

Playtesting Conclusions

Despite the shortcomings listed above, I feel that the testing process added significant value to the project overall. While we may not be able to make any definitive statements about the quality of the guide, the implications of the patterns we've observed are still worth considering. Firstly, the testing has alluded to the existence of a direct correlation between frustration, difficulty and enjoyment within specific difficulty ranges. It also seems to imply that there is some threshold, unique to the individual player, at which this correlation becomes an inverse relationship. At this theoretical point, any increase in difficulty or frustration will result in less enjoyment and engagement overall.

Regarding my initial hypothesis[53], there was a common increase in both difficulty and frustration in all the playtest levels that accounted for accidental solutions (see guide directive 3.16). As such, Levels 1, 2 and 4 all aligned with my expectations, while Levels 3 and 5 surprised me. That being said, the average enjoyment levels seen across the playtest were generally higher in the Original Levels (with the exception of Level 2), which aligned well with my predictions. It is clear that my initial hypothesis had oversimplified the effects of using the guide, and that there is a more complex relationship between the variables that change in response to its use.

Apart from the data obtained, the testing process was valuable as it gave me an opportunity to use the finished guide in practice[54]. I found that the final product behaved as expected, giving me the inspiration and instruction necessary to quickly develop these five levels, as well as a reliable means of refining them over various iterations. Overall, it significantly improved my efficiency as a puzzle designer. In addition to this, developing this testing process proved that there is a means of testing the content of the guide, and created a foundation for future testing processes to be built upon. Considering how rich the initial observations already were, I suspect that more in-depth testing may reveal a whole host of valuable new understanding in the sphere of puzzle design.

[53] The results will show a *general decrease in frustration and difficulty*, as well as an *increase in enjoyment*, in the *Original Levels* (relative to the Stripped-Down levels).

[54] Recall that for major portions of the development of *Allegory*, the guide was also still in development. This is discussed in more detail in the *Overall Reflections* section.

Dessert

Overall Reflections

This section is dedicated to exploring the key lessons that were learnt throughout this entire process. It will unpack all that was achieved, starting from the moment we began constructing a framework using our *Foundational Literature*, moving through to the intense *Content Analysis* that gave us our guide, breezing past the targeted *Playtesting* process, until finally settling on the culmination of it all – *Allegory*.

Theoretical Framework

Before we could create a guide, we needed a solid idea of what that would look like. By meticulously breaking down the *Foundational Literature* of Adams (2015), Bates (2004), Schell (2015) and Selinker and Snyder (2018), we found that there was no single correct way to portray information. That being said, it also became clear that every design choice in a written work had a notable effect on how it would be received and engaged with. By clearly defining the goals and intended audience of our guide, we could establish a firm set of criteria that its format would have to fulfil. This format not only ensured the quality and applicability of the final guide, it gave direction to the next steps in the process, thus making the *Content Analysis* more targeted and achievable.

The lessons embedded in the creation of this framework are valuable to anyone hoping to create literature with an educational or instructional purpose. They have also proven valuable to me as a researcher, author, and designer, as I feel this critical analysis has better equipped me to communicate my experience and knowledge with others.

Content Analysis

Armed with the compass that was our final format, we now had the direction necessary to set a course through the vast sea of puzzle design literature. Outside of the development of *Allegory*, this was easily the largest venture undertaken in the span of this project. Through carefully examining the *Foundational Literature*, as well as a plethora of *Dynamic Informal Sources*, we had hoped to collate a library of the most potent and valuable teachings regarding the design of puzzles.

As the task at hand was colossal, we required an efficient and clever approach to complete it. The unique form of content analysis we used, namely *Divide and Conquer Sorting* paired with *Thematic Selection Sorting*, showed that inspiration for design could come from any field. In this case, techniques and algorithms borrowed from the computer science sphere significantly improved the efficiency of our analysis. As The Game Overanalyser (2019) points out, thinking outside the box and making use of one's entire knowledge base can often lead to new and innovative solutions.

Our guide to puzzle design was a direct product of this Content Analysis. This clearly cements that this endeavour was not in vain, as a massive collection of information had been sorted into a far more palatable format for future designers to use. But aside from the obvious product of this process, the analysis was also particularly valuable to me personally. After engaging with such a vast collection of puzzle design doctrines and critically examining the merits of each, I feel I have grown significantly as a game, level and puzzle designer.

The Guide to Puzzle Design

In Search of Eureka – The Beginner's Guide to Puzzle Design in Videogames is easily the most valuable and important outcome of this project. While *Allegory* stands proud as the most conspicuous product of this work, its value is largely personal to the development team. In contrast, this guide is completely saturated with rich information that future designers may benefit from. Its worth, as well as its identity as an achievement, hardly requires explanation.

With the development of *Allegory,* as well as the design and implementation of a targeted *Playtest*, I had the opportunity to test the guide in practice. Through the application of the guide in my work, I found that it functioned well not only as a source of inspiration and instruction for puzzle creation, but also as an orderly reference bank that connected the reader to an immense archive of puzzle design wisdom.

For those that were just starting out in the field, the guide offered a reasonable level of depth to give the user a basic understanding of the concepts it described. For those that wanted to delve deeper, the guide also showcased some of the intricate connections between various directives, as well as availed access to a plethora of references in a relatively user-friendly manner. Admittedly, more could have been done to make these references easier to navigate[55], but these areas of the guide's format would likely not be the focus of our intended audience, and so were somewhat compromised due to the time constraints of the project.

While the content of the guide may be finite, and the interface integrated into it could be refined, the *Playtesting* has shown that it may have the potential to improve the enjoyment elicited by one's puzzles. This, in conjunction with my own time spent with it, has led me to believe that it holds value in the puzzle design sphere.

Allegory

Allegory is certainly the most striking product developed through this project. Its unique, eye-catching art style and rich, emergent soundscape make it stand out within the submission. With reference to this investigation, the finished game exists as a showcase of my newly obtained puzzle design knowledge, as well as a testament to my growth as a game developer.

The game was created by a team of three master's students and designed in such a way that all of the topics they were investigating could be explored within it. I was the first of the three, with my focus being the puzzle and level design that made up the gameplay experience. With their investigation concentrating on the inclusion and effect of emergent music in videogames, Rodwin Malinga was the second contributing member. Finally, Daniël Hartwell was third, with the intention to explore the relationship between colour theory and affect.

Working as a team undoubtedly allowed us to create a more well-rounded, polished experience, as each of our unique skill sets could be compounded into a single, prominent product. That being said, the need to intricately combine our three unique goals into a single system resulted in there being less time for each of us to explore these individual goals. Blow draws attention to this phenomenon, stating that the intersection of a game's various goals will often have a notable effect on each of its distinct goals, potentially limiting the amount of exploration that is possible in each individual sphere (IndieCade, 2011).

As such, the scope of the final product was slightly smaller than we'd all initially expected it to be. In terms of my area of focus, the final product contained approximately ten distinct puzzles that the player could engage with. This was not a serious issue though, as these puzzles were further

[55] Ideas such as descriptive tooltips, or links to specific moments in the audiovisual references could have been incorporated with more time.

supplemented by the five levels released in the *Playtest* at a later point. In any case, this playtest was probably a more appropriate representation of the effects of the guide and my improved skills as a puzzle designer. This is because it was created *after* the completion of the guide, in contrast to the game, which had to be built in tandem with the guide. As *Allegory* was primarily designed using the *Foundational Design Model*, it is now evident in reflection that there are several areas that could be improved in the final game, had there been more time to iterate on it further.

All in all, I still consider *Allegory* a monumental achievement, given the time frame and scope of the project at hand. It manages to showcase the strengths of all the designers that worked on it, while directly giving each of them a space to explore their relevant topics. On top of this, even while facilitating various forms of investigative and experimental research, I feel that the game did not compromise when it came to the quality of the player experience.

Yet Another Palate Cleanser

At this point, you **may** want to briefly engage with ***Allegory*** again.

Please be aware that this is *completely optional*. Now that you have all the context with which this videogame was created, playing it again will likely reveal new insights that weren't apparent during your initial playthrough.

Mignardise

Conclusion

A veritable library of techniques to improve the experience of solving puzzles exists, but its contents are strewn throughout disparate texts, the doctrines of varying experts in the field, and a mountain of anecdotal irrelevancies. This research sought to collate this vast collection of knowledge into an easily accessible and readily applicable format.

To this end, four key pieces of *Foundational Literature* were critically analysed to construct a framework for the new format (Adams, 2015; Bates, 2004; Schell, 2015; Selinker and Snyder, 2018). With this framework in hand (*R.O.* D) [56], a thematic analysis of both the *Foundational Literature*, as well as a broad collection of *Informal Dynamic Works,* was performed using a custom *Thematic Selection Sort* algorithm. This ultimately led to the creation of two condensed models for puzzle design – one based on the *Foundational* work (*R.O*. A), and the other (*R.O*. B) showcasing the design philosophies of three notable professionals in the field of puzzle design (Blow, Swift, and Teikari). By amalgamating these models into one, *In Search of Eureka – The Beginner's Guide to Puzzle Design in Videogames* was created (*R.O*. C).

The knowledge obtained through this process was further put into practice through the creation of *Allegory*, a puzzle platforming videogame (*R.O*. E). This game, along with a targeted process of playtesting, was used to test the validity and usefulness of the guide. The guide proved to be practical, while housing a significant collection of puzzle design wisdom – more than enough to get a beginner puzzle designer started. In addition to this, *Allegory* was well received by players, offered a valuable space in which to practice and test puzzle design principles, and stood as a manifestation of all that I had learnt as a puzzle designer (*R.O*. F). Thus, I believe that both the final version of the guide, as well as the videogame created through it, were notably successful.

In conclusion, this research has clearly shown me that there is no way to define a single, definite process for creating a puzzle. Each puzzle is unique, and thus calls for its own bespoke implementation. Additionally, each designer is unique, and so has their own beliefs on what is most important in the art of puzzle design. With that being said, there are a number of techniques that can be used to construct a very specific puzzle experience. By broadening one's knowledge of these techniques, and clearly identifying the goals they want their puzzle to achieve, a puzzle designer can notably increase their chances of success in the field. Hopefully, with the use of a tool like *In Search of Eureka*, even novice designers can begin engaging with this concept.

Future Research

Due to the scope of this project, this research cannot truly be considered complete. With more time and resources, it could be significantly improved by future researchers. For instance, the research could be elevated by the collation of additional *Foundational Literature* (such as Rogers' *Level Up! The Guide to Great Video Game Design* (2010) and Crawford's *The Art of Computer Game Design* (1997)), or through the analysis of the philosophies of additional professional puzzle designers.

Outside of expanding the guide, the value of the guide could be more effectively tested. Larger groups of test subjects, and levels and controls designed specifically for selected points in the guide would allow future researchers to draw more concrete conclusions, as well as examine the portions of the guide that were overlooked in our brief testing procedures.

All of these additions have the potential to reveal new and valuable information about puzzle design, and so are worth investigating.

[56] R.O. stands for Research Objective, as outlined in the *Aim* section.

Bibliography

This section contains all the sources used to prepare this research. I have sorted them based on their priority and format to make this list more comprehensive, but a complete alphabetical reference list is also available at the end of the section, if you prefer.

Categorised References

Foundational Literature

Core Works

Adams, E., 2015. *Fundamentals of Puzzle and Casual Game Design*. New Riders.

Bates, B., 2004. *Game design*. 2nd ed. Boston, Mass: Thomson, pp.84-88, 119-134.

Selinker, M. and Snyder, T., 2018. *Puzzlecraft: How to Make Every Kind of Puzzle.* Humble Bundle ed. Renton, WA: Lone Shark Games.

Schell, J., 2015. *The art of game design*. 2nd ed. Pittsburgh: Carnegie Mellon University and Schell Games, pp.239-252.

Additional Books

Crawford, C., 1997. *The Art of Computer Game Design.* Vancouver (Washington): Washington State University. Available at: <https://www.digitpress.com/library/books/book_art_of_computer_game_design.pdf> [Accessed 10 March 2021].

Danesi, M., 2002. *The puzzle instinct*. Bloomington (Indiana): Indiana University Press, pp.1-35.

Rogers, s., 2010. *Level Up! The Guide to Great Video Game Design*. Chichester, West Sussex: John Wiley & Sons, Ltd.

Salen, K. and Zimmerman, E., 2004, 'Chapter 7: Defining Games', in *Rules of play*. 1st ed. Cambridge, Mass: The MIT Press, pp. 83-96.

Selinker, M. and Snyder, T., 2013. *Puzzlecraft: The Ultimate Guide on How to Construct Every Kind of Puzzle*. Puzzlewright.

Additional Articles, Journals and Papers

Asbell-Clarke, J., Barnes, T., Eagle, M., Edwards, T., Hicks, D. and Rowe, E., 2016. Using game analytics to evaluate puzzle design and level progression in a serious game. *Proceedings of the Sixth International Conference on Learning Analytics & Knowledge - LAK '16*, [online] Available at: <https://dl.acm.org/doi/10.1145/2883851.2883953> [Accessed 10 March 2021].

Fernandez-Vara, C. and Osterwil, S., 2010. *The Key to Adventure Game Design: Insight and Sense-making*. [online] Hdl.handle.net. Available at: <http://hdl.handle.net/1721.1/100238> [Accessed 10 March 2021].

Kim, S., 2006. *What is a Puzzle?* [eBook] pp.1-5. Available at: <http://cs.wellesley.edu/~cs215/Lectures/L17-IntroGamesJigsawPuzzle/ScottKim-What_is_a_Puzzle.pdf> [Accessed 10 March 2021].

Linehan, C., Bellord, G., Kirman, B., Morford, Z. and Roche, B., 2014. Learning Curves: Analysing Pace and Challenge in Four Successful Puzzle Games. *Proceedings of the first ACM SIGCHI annual symposium on Computer-human interaction in play*, [online] Available at: <https://www.researchgate.net/publication/268389769_Learning_Curves_Analysing_Pace_and_Challenge_in_Four_Successful_Puzzle_Games> [Accessed 18 April 2021].

Dynamic Informal Sources

Talks and Conferences

Blow, J., ten Bosch, M. and Martin, A., 2013. *Empuzzlement*.

CS50, 2018. *Portal Problems - Lecture 11 - CS50's Introduction to Game Development 2018*. [video] Available at: <https://www.youtube.com/watch?v=ivyseNMVt-4> [Accessed 12 July 2021].

Gamelab Conference, 2019. *#GAMELAB19: Making Baba Is You, what is this? by Arvi Teikari Hempuli*. [video] Available at: <https://www.youtube.com/watch?v=8eTFea_UZcc> [Accessed 3 May 2021].

Indie Cade, 2011. *IndieCade 2011: Jonathan Blow & Marc Ten Bosch*. [video] Available at: <https://www.youtube.com/watch?v=OGSeLSmOALU> [Accessed 12 July 2021].

Kim, S., 2008. *The art of puzzles*.

Menzel, J., 2019. *Level Design Workshop: Solving Puzzle Design*.

Swift, K., 2007. Our Journey From Narbacular Drop To Portal. In: *Game Developers Conference 2007*. [online] Informa PLC. Available at: <https://www.youtube.com/watch?v=2F0SVZA9fIo> [Accessed 12 July 2021].

Talk & Dev, 2020. *Talk & Dev: The making of Baba Is You - Arvi Teikari (Independent)*. [video] Available at: <https://www.youtube.com/watch?v=v7ICYHL5MD4> [Accessed 12 July 2021].

Video Essays and Documentaries

Adam Millard - The Architect of Games, 2021. *What We Can Learn From A 2000 Year Old Puzzle*. [video] Available at: <https://www.youtube.com/watch?v=KAd3lCQFhPc> [Accessed 3 May 2021].

Game Maker's Toolkit, 2019. *How Baba Is You Works*. [video] Available at: <https://www.youtube.com/watch?v=7zLwa4bztWs> [Accessed 3 May 2021].

Game Maker's Toolkit, 2018. *What Makes a Good Puzzle?* [video] Available at: <https://www.youtube.com/watch?v=zsjC6fa_YBg> [Accessed 20 January 2021].

Game Maker's Toolkit, 2016. *How Jonathan Blow Designs a Puzzle*. [video] Available at: <https://www.youtube.com/watch?v=2zK8ItePe3Y> [Accessed 22 January 2021].

O'Dwyer, D. and Jayne, J., 2017. *The Witness Documentary*. [video] Available at: <https://www.youtube.com/watch?v=YdSdvIRkkDY> [Accessed 10 May 2021].

The Game Overanalyser, 2019. *The Art of Puzzle Design | How Game Designers Explore Ideas and Themes with Puzzles and Problems*. [video] Available at: <https://www.youtube.com/watch?v=hCOHjTX4GYE> [Accessed 12 May 2021].

Writing on Games, 2016. *The Witness and the Rejection of Immediacy (Analysis/Review) - Writing on Games*. [video] Available at: <https://www.youtube.com/watch?v=p9w5GFC5ii8> [Accessed 30 June 2021].

Interviews

Alexander, L., 2014. *The Witness: Modeling epiphany*. [online] Gamasutra. Available at: <https://gamasutra.com/view/news/218953/The_Witness_Modeling_epiphany.php> [Accessed 10 March 2021].

Couture, J., 2018. *Road to the IGF: Hempuli Oy's Baba is You*. [online] Gamasutra. Available at: <https://www.gamasutra.com/view/news/314548/Road_to_the_IGF_Hempuli_Oys_Baba_Is_You.php> [Accessed 12 July 2021].

Parkin, S., 2008. *Jonathan Blow: The Path to Braid*. [online] Gamasutra. Available at: <https://www.gamasutra.com/view/feature/132180/jonathan_blow_the_path_to_braid.php> [Accessed 10 March 2021].

Rigney, R., 2012. *Portal Designer Kim Swift Won't Let You Take Her Toys Away*. [online] Wired. Available at: <https://www.wired.com/2012/06/kim-swift-interview/> [Accessed 12 July 2021].

Articles

Davies, M., 2015. *A Good Puzzle Game Is Hard To Build*. [online] Rock Paper Shotgun. Available at: <https://www.rockpapershotgun.com/how-to-make-a-puzzle-game> [Accessed 10 March 2021].

Gaurav, S., 2022. *What is the Time Complexity of Selection Sort?* [online] Scaler. Available at: < https://www.scaler.com/topics/time-complexity-of-selection-sort/> [Accessed 5 March 2023]

GeeksforGeeks, 2023. *Introduction to Divide and Conquer Algorithm – Data Structure and Algorithm Tutorials.* [online] GeeksforGeeks. Available at: < https://www.geeksforgeeks.org/introduction-to-divide-and-conquer-algorithm-data-structure-and-algorithm-tutorials/> [Accessed 5 March 2023]

Games

Teikari, A (2019) *Baba is You* [Video Game] Teikari, A.

Thekla Inc. (2016) *The Witness* [Video Game] Thekla Inc.

Valve (2007) *Portal* [Video Game] Valve.

Full Reference List

Adam Millard - The Architect of Games, 2021. *What We Can Learn From A 2000 Year Old Puzzle*. [video] Available at: <https://www.youtube.com/watch?v=KAd3lCQFhPc> [Accessed 3 May 2021].

Adams, E., 2015. *Fundamentals of Puzzle and Casual Game Design*. New Riders.

Alexander, L., 2014. *The Witness: Modeling epiphany*. [online] Gamasutra. Available at: <https://gamasutra.com/view/news/218953/The_Witness_Modeling_epiphany.php> [Accessed 10 March 2021].

Asbell-Clarke, J., Barnes, T., Eagle, M., Edwards, T., Hicks, D. and Rowe, E., 2016. Using game analytics to evaluate puzzle design and level progression in a serious game. *Proceedings of the Sixth International Conference on Learning Analytics & Knowledge - LAK '16*, [online] Available at: <https://dl.acm.org/doi/10.1145/2883851.2883953> [Accessed 10 March 2021].

Bates, B., 2004. *Game design*. 2nd ed. Boston, Mass: Thomson, pp.84-88, 119-134.

Blow, J., ten Bosch, M. and Martin, A., 2013. *Empuzzlement*.

Couture, J., 2018. *Road to the IGF: Hempuli Oy's Baba is You*. [online] Gamasutra. Available at: <https://www.gamasutra.com/view/news/314548/Road_to_the_IGF_Hempuli_Oys_Baba_Is_You.php> [Accessed 12 July 2021].

Crawford, C., 1997. *The Art of Computer Game Design.* Vancouver (Washington): Washington State University. Available at: <https://www.digitpress.com/library/books/book_art_of_computer_game_design.pdf> [Accessed 10 March 2021].

CS50, 2018. *Portal Problems - Lecture 11 - CS50's Introduction to Game Development 2018*. [video] Available at: <https://www.youtube.com/watch?v=ivyseNMVt-4> [Accessed 12 July 2021].

Danesi, M., 2002. *The puzzle instinct*. Bloomington (Indiana): Indiana University Press, pp.1-35.

Davies, M., 2015. *A Good Puzzle Game Is Hard To Build*. [online] Rock Paper Shotgun. Available at: <https://www.rockpapershotgun.com/how-to-make-a-puzzle-game> [Accessed 10 March 2021].

Fernandez-Vara, C. and Osterwil, S., 2010. *The Key to Adventure Game Design: Insight and Sense-making*. [online] Hdl.handle.net. Available at: <http://hdl.handle.net/1721.1/100238> [Accessed 10 March 2021].

Gamelab Conference, 2019. *#GAMELAB19: Making Baba Is You, what is this? by Arvi Teikari Hempuli*. [video] Available at: <https://www.youtube.com/watch?v=8eTFea_UZcc> [Accessed 3 May 2021].

Game Maker's Toolkit, 2019. *How Baba Is You Works*. [video] Available at: <https://www.youtube.com/watch?v=7zLwa4bztWs> [Accessed 3 May 2021].

Game Maker's Toolkit, 2018. *What Makes a Good Puzzle?* [video] Available at: <https://www.youtube.com/watch?v=zsjC6fa_YBg> [Accessed 20 January 2021].

Game Maker's Toolkit, 2016. *How Jonathan Blow Designs a Puzzle*. [video] Available at: <https://www.youtube.com/watch?v=2zK8ItePe3Y> [Accessed 22 January 2021].

Gaurav, S., 2022. *What is the Time Complexity of Selection Sort?* [online] Scaler. Available at: < https://www.scaler.com/topics/time-complexity-of-selection-sort/> [Accessed 5 March 2023]

GeeksforGeeks, 2023. *Introduction to Divide and Conquer Algorithm – Data Structure and Algorithm Tutorials.* [online] GeeksforGeeks. Available at: < https://www.geeksforgeeks.org/introduction-to-divide-and-conquer-algorithm-data-structure-and-algorithm-tutorials/> [Accessed 5 March 2023]

Kim, S., 2006. *What is a Puzzle?* [eBook] pp.1-5. Available at: <http://cs.wellesley.edu/~cs215/Lectures/L17-IntroGamesJigsawPuzzle/ScottKim-What_is_a_Puzzle.pdf> [Accessed 10 March 2021].

Kim, S., 2008. *The art of puzzles*.

Linehan, C., Bellord, G., Kirman, B., Morford, Z. and Roche, B., 2014. Learning Curves: Analysing Pace and Challenge in Four Successful Puzzle Games. *Proceedings of the first ACM SIGCHI annual symposium on Computer-human interaction in play*, [online] Available at: <https://www.researchgate.net/publication/268389769_Learning_Curves_Analysing_Pace_and_Challenge_in_Four_Successful_Puzzle_Games> [Accessed 18 April 2021].

Menzel, J., 2019. *Level Design Workshop: Solving Puzzle Design*.

O'Dwyer, D. and Jayne, J., 2017. *The Witness Documentary*. [video] Available at: <https://www.youtube.com/watch?v=YdSdvIRkkDY> [Accessed 10 May 2021].

Parkin, S., 2008. *Jonathan Blow: The Path to Braid*. [online] Gamasutra. Available at: <https://www.gamasutra.com/view/feature/132180/jonathan_blow_the_path_to_braid.php> [Accessed 10 March 2021].

Rigney, R., 2012. *Portal Designer Kim Swift Won't Let You Take Her Toys Away*. [online] Wired. Available at: <https://www.wired.com/2012/06/kim-swift-interview/> [Accessed 12 July 2021].

Rogers, s., 2010. *Level Up! The Guide to Great Video Game Design*. Chichester, West Sussex: John Wiley & Sons, Ltd.

Salen, K. and Zimmerman, E., 2004, 'Chapter 7: Defining Games', in *Rules of play*. 1st ed. Cambridge, Mass: The MIT Press, pp. 83-96.

Schell, J., 2015. *The art of game design*. 2nd ed. Pittsburgh: Carnegie Mellon University and Schell Games, pp.239-252.

Selinker, M. and Snyder, T., 2018. *Puzzlecraft: How to Make Every Kind of Puzzle.* Humble Bundle ed. Renton, WA: Lone Shark Games.

Selinker, M. and Snyder, T., 2013. *Puzzlecraft: The Ultimate Guide on How to Construct Every Kind of Puzzle*. Puzzlewright.

Swift, K., 2007. Our Journey From Narbacular Drop To Portal. In: *Game Developers Conference 2007*. [online] Informa PLC. Available at: <https://www.youtube.com/watch?v=2F0SVZA9fIo> [Accessed 12 July 2021].

Talk & Dev, 2020. *Talk & Dev: The making of Baba Is You - Arvi Teikari (Independent)*. [video] Available at: <https://www.youtube.com/watch?v=v7ICYHL5MD4> [Accessed 12 July 2021].

The Game Overanalyser, 2019. *The Art of Puzzle Design | How Game Designers Explore Ideas and Themes with Puzzles and Problems*. [video] Available at: <https://www.youtube.com/watch?v=hCOHjTX4GYE> [Accessed 12 May 2021].

Thekla Inc. (2016) *The Witness* [Video Game] Thekla Inc.

Teikari, A (2019) *Baba is You* [Video Game] Teikari, A.

Valve (2007) *Portal* [Video Game] Valve.

Writing on Games, 2016. *The Witness and the Rejection of Immediacy (Analysis/Review) - Writing on Games*. [video] Available at: <https://www.youtube.com/watch?v=p9w5GFC5ii8> [Accessed 30 June 2021].