

A Vision for the Design of AI in AI-based Gameplay

Ben Samuel

Computer Science
University of New Orleans
New Orleans, LA USA
bsamuel@cs.uno.edu

Mike Treanor

Computer Science and AU Game Lab
American University
Washington, DC USA
treanor@american.edu

Joshua McCoy

Computer Science and Cinema and Digital Media
University of California, Davis
Davis, CA USA
jamccoy@ucdavis.edu

Abstract—Artificial Intelligence has been applied to many facets of the game design and development process. Though this has led to many advances in games and AI research, there remain few examples of games in which play centers around engagement with AI processes: the design space of AI-based games remains underexplored. By examining a breadth of playful experiences through different lenses, it is determined that games which forefront AI are beneficial for players, designers, and for the field of game scholarship itself. Moreover, there is evidence that symbolic approaches (rather than statistical) lend themselves to experiences with more agency, greater human interpretability, and more controlled authorability.

Index Terms—AI-based game design, symbolic methods, statistical methods,

I. INTRODUCTION

Artificial Intelligence is used in many aspects of the game design and development process and is also a central part of gameplay. While there is much research in AI and games, there are aspects of game AI that are still not well understood and are critical to our understanding of games, gameplay, and design. Recently, the range of AI techniques used for game AI research clusters around machine learning and statistical approaches. Though this led to considerable advancements in areas such as player modeling, procedural content generation, natural language interfaces, and other areas, the use of AI as a core part of gameplay is not frequently addressed in depth.

Games that make AI available for the player to interact with and also understand have been referred to as AI-based games [1], [2]. These games can offer representations of complex phenomena such as social dynamics, have the potential to provide customizable narratives and events that adapt in real-time based on players' predilections, and can serve as tools and authoring systems that enable and inspire users to engage in self-expression and produce creative artifacts of their own. Frequently, the successful play of these games depends on the player developing a robust and somewhat accurate model of the underlying processes of their AI systems. We envision a world with more AI-based games that challenge existing gameplay conventions and help to explore the human experience. For this world to exist, there needs to be more technical, theoretical, and practice-based research.

This paper aims to promote AI-based game design research by laying out a roadmap to the techniques and challenges

within the space. We do so by first identifying where exactly AI-based game design exists within the space of AI research and application in game development. We make this step not to discourage work in the other areas of the space but to identify and emphasize that the space of AI-based game design is under-explored. In the following section, we describe several aspects of games—such as a game's agency and interpretability—in which AI influences the experience of both authors and players and how different approaches to AI influence them. The two general approaches we focus on are the broad areas of statistical and symbolic AI. We make these distinctions because we believe that symbolic approaches will promote good design, and we intend to push gently against the current focus on statistical methods. In the last section, we soften the dichotomy and present a comprehensive view of hybrid approaches that acknowledges the strengths and weaknesses of both approaches. When taken as a whole, these contributions present a vision for future research into AI-based game design.

II. A VIEW ON AI IN GAMES

As much as the authors might like to refer to “AI in Games” and leave it at that, that phrase inadequately captures the diverse stages of game development and the varied ways artificial intelligence is applied to them. To focus the conversation of this paper, the authors first present a novel view on artificial intelligence in games: two axes that form a graph featuring the placement of different aspects of game design and game development. This is an extension of previous writing on AI-based game design [1], [2], which we will soon see, has traditionally focused on the upper-right quadrant of this new model.

To understand this model, we must first understand the axes. One axis is the designer-centric Surface / Background axis. The other axis is the player-centric Mental Model axis. Upon this model, we place elements of the game development process that have had (or could have) artificial intelligence techniques applied to them. Note that any given individual game could occupy multiple places on this graph if AI was used in multiple aspects of its development.

A. The Surface / Background Axis

The Surface / Background Axis captures the notion of *where* in the artificial intelligence is applied to the game. If we think of game development as a pipeline, there are many stages to the process. Some of these stages include the initial brainstorming and ideation of pre-production, development of the game, marketing the game, and post-production upkeep and content drops (e.g., new DLC and other expansions in the currently popular “games as a service” model, maintaining servers for multiplayer components, etc.). Techniques that are considered “high-surface, low-background” are incorporated into the actual act of play itself and directly affect or influence the player experience while playing. Crucially, “high surface” techniques are also capable of being influenced by a player’s play as well—thus, they can influence and be influenced in turn. Application areas considered “high-background, low-surface” are peripheral to the act of play but are indeed no less critical to the process of game development.

To illustrate this, let us look at the game *Left 4 Dead* [3]. The AI systems of *Left 4 Dead* [4], and in particular, its AI Director, would rank highly on the surface dimension. Every time the player plays, it actively shapes their experience, modulating the number of enemies and equipment in the scenario to evoke a continuously tense-but-fair match. An element ranked lower on the surface axis is the game’s matchmaking system (or, indeed, matchmaking systems of multiplayer games in general). Certainly, the other players that a player is matched up with (or matched against, in competitive multiplayer games) will directly influence the game playing experience (i.e., ‘high surface’). However, it is a decision made at the outset of the match and then statically set in stone for its duration, outside of player influence (beyond the player’s accumulated MMR, ELO, or other scores that the player has garnered through previous play sessions [5]). A low surface / high background aspect of these games would be player telemetry [6]: capturing, logging, and analyzing button clicks, menu navigation, etc. This work has applications at several levels of the game development process, from user and player research for game design [7] to ensuring player retention for monetization [8]. Though the financial well-being of a developer undoubtedly indirectly affects and influences the player experience, these ‘behind the scenes’ measures do not directly affect the player’s gameplay experience.

B. The Mental Model Axis

The first axis roughly corresponds to applications of artificial intelligence on different parts of the game development production process. The Mental Model Axis shows how different forms of artificial intelligence might be present in the player’s mind. We refer to this as the mental model axis to speak to how much the player’s experience of play is affected—and, more specifically, rewarded—by devoting the energy required to construct a mental model of how the underlying AI systems of the gameplay experience work.

The presence of an AI system in a playable experience is insufficient to guarantee that a player will be incentivized to

develop this mental model. Wardrip-Fruin’s *Sim City Effect* [9] discusses the process of a player discovering the contours of a system through play. Though they may never fully learn all of the intricacies of the underlying system (nor learn enough practical city management for an actual mayoral bid), they might slowly piece together small connections of the greater system. Examples of these effects are industrial zoning and police stations impact on crime or the relative impact roads and railways have on pollution. These small revelations accumulate into the player being able to design cities that achieve their own particular goals. This is in contrast to Wardrip-Fruin’s *Tale-Spin effect*, which describes experiences with rich and complex underlying systems that are manifested via a surface layer presentation so simple it dissuades players from believing a sophisticated system even exists.

A visualization of these two axes is featured in Figure 1.

C. Quadrant Descriptions

The upper-left corner describes applications of artificial intelligence that directly affect the gameplay experience but require minimal mental modeling by the player. Many procedural content generation systems, in which the generator is run once and then left to explore, fall under this heading. Examples include *Minecraft* and other mining games [10], the map generation, procedurally generated sound, and many other applications [11].

The lower-left corner describes applications of artificial intelligence that do not affect the player, nor do they directly influence gameplay. Uses of AI in this quadrant are concerned mainly with the business of game development. As previously referenced, gathering user data for player retention purposes and monetization would fall into this quadrant [8]. This is undoubtedly an exciting and lucrative application area for many artificial intelligence techniques but is not the focus of this paper.

The lower-right quadrant contains application areas of artificial intelligence that could potentially reward the astute player but are designed not to be interactive. One example of this is the Matchmaking system that exists in many games. Systems of this type are a means to an end; they ferry the player to the next match. However, even an innocuous system can lead players to intentionally throw matches to artificially lower their matchmaking ranking, thus enabling them to face easier opponents. Similarly, MOBAs and other team-based games have had to find solutions to draft dodging. Telemetry and analysis of user play traces have been conducted to help inform the development of matchmaking services [12]. Non-AI mediated ranking measures in eSports face their own challenges [13].

The upper-right quadrant is the quadrant that the authors call for more people to explore; the type of game previously discussed as AI-based games. These are aspects of gameplay informed or enabled by artificial intelligence that players interact with directly, and for which the experience depends upon the player developing a mental model of these systems through play. Game experiences such as *The Sims* [14], *Dwarf*

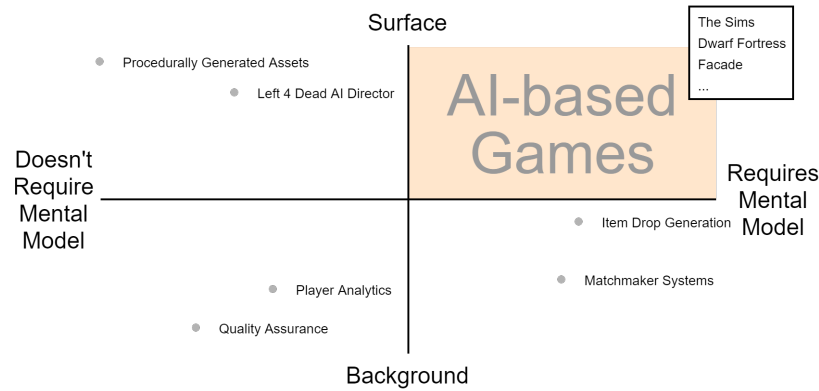


Fig. 1. A view on the use of artificial intelligence in games. In the top right quadrant are games that use AI to directly affect gameplay and also require a mental model of how the AI operates in order to play.

Fortress [15], and *Façade* [16] all reside here. Though previous quadrants represent applications of AI within a game, here we are referring to games as a whole. These are playable experiences whose AI Systems are core to their character; the line between the overall game and the underlying AI system is blurred. It is the vision of the authors that this quadrant needs significantly more research devoted to it to be fully explored.

The following section presents a comparative approach to different AI techniques and how the same 'underlying' experience might be radically different based on the AI technique employed to enable it.

III. COMPARATIVE ANALYSIS OF THE USE OF DIFFERENT APPROACHES TO AI IN AI-BASED GAMES

The overall goal of this paper is to advocate for and facilitate the creation of games that use AI within the surface-level/requires a mental model quadrant described above (i.e., AI-based games). Such games involve players directly engaging with the AI as part of their aesthetic experience, and this section lays out several ways that the use of AI influences that experience. Throughout the section, we compare two broad categories of AI: statistical and symbolic. We realize that these categories can be blurred (the next section goes into more detail), but we found it useful to consider each approach at this level. Overall, this section is meant to push against the current in AI research and explain why we believe there is ample opportunity to create novel and satisfying games by employing symbolic AI in AI-based games.

A. Interpretability

While both creating and playing games with targeted representations or experiences, the ability to interpret a game's behavior is an important factor. Game designers have framed this as the ability to learn the rules of the game system [17], creating Meaningful Play [18], and avoiding arbitrary and "meaningless" choices [19]. Each of these game designers advocates that a player should be able to reason about the operation of the game system and make informed choices toward their goals. This is especially important with AI-based

games, where all interaction is framed by the player needing to read meaning into the AI's actions [20], as the systems are often complex and dynamic.

In games that use AI that does not require mental models in order to intentionally interact with, interpretability is less important. For example, while understanding how the landscape of a planet was generated in *No Man's Sky* [21] is arguably not particularly relevant, understanding the connection between a *Sim's* [14] "needs meter" and their action is relevant.

1) *Statistical Approaches*: Interpretability is a well-recognized issue in machine learning [22]. Limited interpretability is inherent in most statistical methods as the content of a learned model is not represented in human-interpretable ways. However, statistical methods often still produce results that adhere to expectations and are very powerful for many applications in which the output is more important than process (e.g. texture generation, speech to text, computer vision, etc.).

However, this obliqueness can be an issue for AI-based games that prioritize the player's ability to build a mental model of the AI's operation. From a player's perspective, the interaction loop (listen, think, speak [23]) is a process of learning. Every output of a system gives the player material to build an understanding of how it operates. When a player makes a conclusion about the operation of a system that primarily relies on statistical methods, those conclusions necessarily cannot map to the true operation of the system, as the operation of the system is not internally represented in human-understandable terms.

That is not to say that players of these games can not build mental models that are useful to use toward desired ends. However, such mental models are limited in that their primary ability is to anticipate how a system will behave, but not why. Though, from a player-centric perspective, this will only be a problem if the output of the system significantly violates the player's model without providing enough evidence to update the model, or descriptions of why different behavior occurred (which, typically referred to as explainable AI [22], is one of the core challenges of statistical methods).

2) *Symbolic Approaches*: Given that symbols are the words used to refer to representational entities, symbolic approaches are highly interpretable as humans are able to ascribe meaning to them. When a theorem prover makes a proof, the step by step process is available for humans to inspect. Likewise, when a system enacts a plan for an agent, the system's rationale is available. While the "reasoning" of symbolic systems aren't always the most obvious to a human, they are valid and justified. This mechanical or "alien" manipulation of symbols might be considered aesthetically good or bad depending on the game, though it can negatively affect the believability of agent behavior.

Often, symbolic systems make use of sophisticated architectures to select content or choose agent behavior. Symbolic AI-based games involve the player acquiring a rough model of these architectures in order to achieve their desired ends. For example, in *The Sims 4* [14], the player is able to associate a Sim's need hygiene meters for the degree of hygiene a Sim has because to a human a low amount of "hygiene" is symbolically associated with not being clean. The system architecture, and the authoring, supports this with its behavior in other areas of the game. Furthermore, there can be a strong relationship between authored content and what is presented to the player.

However, it should be emphasized that a human's interpretation of a game (and its symbols), and even their perception of what symbols are available for interpretation can, and likely will, be different than what the system uses or the author intends. In order for a game to symbolically represent an author's intent, a game needs to be carefully designed and delivered to a receptive audience [24].

B. Agency

Most often, game designers strive to give the players of their games the experience of high agency. Generally, having a sense of agency involves the feeling of one having the power to take specific action in a situation toward a desirable end. In games, agency has been theorized to be achieved when the formal and material affordances of a system are in balance [25]. In other words, when the game leads a player to consider doing something (the formal), it is also able to respond to it (through the material). In many games, these criteria are not met, and games do not provide high agency experiences. This is most prevalent in narrative games. For example, the story and cutscenes in a game like *Grand Theft Auto V* [26] present detailed characters with personalities and things to say, but in gameplay characters mostly serve as targets and ragdolls, and do not respond to the player's actions or the resulting dynamics.

Much work has been done in the field of academic game and autonomous agent research to address this, and it is a difficult technical problem. However, it should be noted that agency can also be achieved through design choices. For example, lower fidelity character content (graphics, dialogue, etc.) arguably suggests to a player that the game is not meant to respond to everything that you would expect a human-like agent to respond to. Also, genre convention also informs what players

will expect, and thus their sense of agency. The degree to which a game can be said to provide agency results from a combination of technical, design, and cultural factors.

1) *Statistical Approaches*: A purely statistical approach to physical behavior in a game would inherently present a problem for agency, as part of what a player expects from a game is tied to their experience of reality (object permanence, consistent laws of physics, etc.). However, other areas of behavior are perceived as less static, and statistical methods can greatly alleviate authorial burden and provide surprising, and interesting experiences. For example, learned models can drive animation in unanticipated, or not specifically accounted for, environments or dialogue can be generated on the fly.

However, for AI-based games, the player is intended to have an idea of how and why the system performs as it does, and statistical methods are black boxes that the player has little access to understanding. Because agency is tied so closely to what a player perceives and expects, this is another area where a system's processes are more important than its output, and because statistical methods are generally not interpretable, they can impede how much agency a game can provide.

2) *Symbolic Approaches*: On the other hand, more symbolically oriented systems provide structure that is available to be discovered by players. Symbolic systems arguably have more consistent "hooks" that a player can ascribe meaning to and extrapolate about the system's overall operation.

As an example, consider the 2013 version of *SimCity* [27]. In order to grow their city as desired, the player must be roughly cognizant of the operation of the underlying Glassbox system [28]. Generally, Glassbox simulates the flow of resources between buildings via agents traversing player designed paths. Because the system is symbolic, the game is able to directly present select aspects of the Glassbox engine, and the player is able to take directed action and thus potentially achieve agency (depending on their level of understanding). This version of *SimCity* could be contrasted with an imagined city-building simulator that evolved based on a model based on large data sets from real world cities. In this game, the player would layout roads, zones, etc. (just as in *SimCity*), but how the city evolves would be generated from the model. While this game is arguably a more accurate simulation of a real world city and its growth, players would not have the same sense of agency, as the system's output would not map to a mental model, but rather the opaque learned model which would most likely invalidate the player's fledgling mental model.

Symbolically represented systems contain a theory of what they represent and are prescriptive, rather than just descriptive. Experientially, this lends symbolic systems a degree of legitimacy when the system doesn't behave as a player might expect. Without this, a player experiences a sense of helplessness and distrust that their actions or experience are meaningful to gameplay (i.e. a loss of agency).

C. Trust

We posit that a valuable lens through which to view these systems is one of trust. Though trust between a user and a system is difficult to measure, there exists some writing [29] [30] on how trust can be used as an evaluation metric of mixed-initiative and co-creative artificial intelligence systems [31]. In brief, trust is a form of measuring if a system is acting in “good faith”, where good faith is both a very human-centric value and is deeply contextually dependent on the experience in question. We discuss examples of both symbolic and statistical methods and the trust (or lack thereof) evoked.

1) *Statistical Approaches*: Trust—be it between humans or systems—is built over time. Imagine working together with a fellow human. Through repeated exposure, one learns the predilections of one’s partner; if they are reliable (i.e., can be trusted to complete work as expected), if they are genuine (i.e., what they say can be trusted), and if they are discrete (i.e., one can trust them with thoughts and feelings). Though humans are perhaps black boxes to each other in some ways, we also have a knack for divining each others’ moral codes; the guiding principles that dictate decision making. Though perhaps one might not describe a system as having a moral code, as discussed in the recent calls for Explainable AI [22], statistical approaches suffer from a human-understandable representation of the world, which directly translates into a difficulty to trust.

Take, for example, Tay.ai. [32] [33], a chatbot developed by Microsoft that was intended to emulate the speech patterns of a teenager. Though a chatbot might not be considered a “game” in the traditional sense, we posit it still adheres to the principles of expressive AI [20], in that system and user/player cyclically interact with, and enact change within, one another. In this spirit, Tay.ai adopted speech patterns based on interactions with users. Within a day, trained by those sending it words, it had managed to become a sexist, racist bigot. Though some damage control was maintained through a banned-word list, it was ultimately viewed as unsalvageable and quickly taken down. Its statistical approach to representation meant Tay.ai had no understanding that it was being offensive, and it is this naive, unknowing-yet-rapid descent into bigotry that rendered it untrustworthy. The trust with the system had been broken and deemed impossible to repair. Though outside the domain of games, Amazon’s statistical learning recruitment tool [34] was similarly naively bigoted, though far more subtle in its execution.

Though statistical approaches are enjoying a modern renaissance, this is not a recent phenomenon. Another example, with interesting if somewhat circuitous connections to play, is the Microsoft Word animated pedagogic assistant [35] affectionately known as Clippy. Clippy used Bayesian statistics [36] to identify the type of document the author was writing in real time and offer formatting suggestions. Like Tay.ai, there were reported issues of trust on an interpersonal level with Clippy. Women in initial focus groups voiced concerns that Clippy was “leering” at them, but these concerns were largely ignored [37]. Moreover, Clippy failed to present himself as a

particularly empathetic or human creative partner, as written about by Clifford Nass [38]. Interestingly enough, Nass writes of a prototype version of Clippy who would egg frustrated users on. This Clippy with attitude was much better received by test audiences, though it was never released.

It is interesting to note that to convert Clippy into his rude second form, users would have to directly interact with him. Rather than having him statistically infer truths via users editing the document, the user would specifically click a button to inform Clippy that he was being annoying. This is, in effect, the Microsoft Word version of a dialogue tree [39], and a symbolic approach to representing state. This could be considered an extremely simple example of a hybrid approach to artificial intelligence, and it is interesting to note that the statistical portion of this approach elicited frustration, while the symbolic provided the fun.

2) *Symbolic Approaches*: Earlier we mentioned the use of trust as a metric of success for co-creative or mixed-initiative systems. The goal of these systems vary, but they often provide an interesting blur between the roles of author and end-user. For example, Tanagra is a mixed-initiative tool for developing 2D platformer levels [31], which employed symbolic methods to facilitate and automate the construction of levels whose designs drew from music theory. Though ostensibly the goal for the user of such a tool is to develop levels that will ultimately be enjoyable to an end user (that is, a player), the act of engaging with such a tool can be considered a form of play unto itself. Let us consider, then, the level designer as a player, as well as the trust they might share with such a system.

Due to the affordances of the system, it is unlikely that the system would produce misogynistic or racist levels; the constrained affordances of the system of placing platforms and enemies ensures the expressive range of these levels focuses more on player challenge, and not challenging societal decorum [40]. However, the act of engaging with such a tool can also challenge the designer’s design sensibilities in a positive way. To form deep trust with a system is to enable it to affect you. Just as a collaborative writing partner could help one recognize tropes or patterns that one is becoming overly reliant on, a mixed-initiative tool could offer suggestions not based (solely) on creating a product for an end user, but rather offer suggestions to help the designer break new ground or develop their own style.

Though statistical approaches for such a system would be adept at recognizing the designer’s patterns of behavior, it is hard to envision how such a system would be trained to encourage meaningful deviation from these patterns. A more symbolic approach to this task could, for example, search a space of states; discover paths or through-lines that connect the author’s current designs to distant concepts, and present both the unexplored design space and a path to reach it. This enables the author to not only see new design territory but enables the author to understand the systems’ reasoning to get there. And with that, we have come full circle, as we stated at the beginning of this section: this understanding is key to

nurturing trust.

D. Consistency and Coherence

Whether a game aims to immerse the player in a fictional environment or engage players through strategic problem-solving scenarios, it is important that what players accurately learn about its operation are consistent. In representational worlds, common consistent features might include object permanence and gravity, and in abstract games, the features might be the game's rules, or win conditions. This isn't to say that players must always have perfect information about how the system operates, that is almost never the case, but very often good game design involves players feeling like they are learning about how to better play the game as they play (i.e. learning how the system operates). When a game presents players with information that invalidates player beliefs too often without convincing justifications, players can begin to feel that the game is choosing behavior randomly and ultimately feel that their input doesn't matter.

1) *Statistical Approaches*: AI Dungeon 2 is a text-based game that has the player control an avatar through typing any command into a prompt [41]. Using a variant of GPT-2, the system generates a description of what happens when the player takes that action and then gives the player a chance to give another command. While the powerful underlying natural language interface model generates surprising and amusing results, the game arguably doesn't give the player meaningful choice [18]. This is evidenced by the fact that most often a game ends when the player is killed by some unexpected event. This happens because the underlying learned model doesn't have the ability to maintain a consistent representation of the game world. Generally, statistical approaches only excel at maintaining local coherence (e.g. sentence level) and fail at broader levels of resolution (e.g. story). That said, the pleasure of AI Dungeon 2 is exploring its opaque model. While players may not be able to rely on the game to be fair, or consistent, it can still be amusing to wonder what is going on inside the black box.

2) *Symbolic Approaches*: In other types of games, the system's reaction to player choice is very important, and invalidating a player's mental model can ruin the experience. Different than statistical methods, symbolic approaches are able to explicitly track context and consistently react to it while in operation. This allows for the system to enforce hard preconditions and to better ensure that appropriate content is presented. A downside of this is that the higher order context needs to be supported through architecture, and often requires much more authored content. AI Dungeon 2 is an exception to most interactive fiction games, as most interactive fiction games make use of heavily symbolic systems such as Inform 7 [42]. Inform 7 roughly represents the fictional worlds as objects, containers, and actions that can be performed upon them. State change and descriptions of actions are chosen based on a rule system. As a result, games based on Inform 7 are deeply strategic and puzzle-like.

E. Authorability

An AI system's authorability depends on how well it supports the creation or addition of content. Just like the design and development of other parts of video games, authoring is an act of creation that is ultimately judged subjectively by players. It is also subject to hard constraints imposed by the game's systems and by soft constraints from existing and future content. Additionally, authors have a limited ability to create even when obeying those constraints. Not only do authors have to work with these constraints while being judged by players, creating content for AI-based games requires interfacing with complex technologies.

To create, authors need a channel to communicate design decisions with the AI system. Meaningful communication to the AI system is central to these aspects, varies widely, and includes text, reactive planners for character behavior [16], bespoke tools [43], data sets, or even mixed-initiative tools [31]. At a high level, this communication impacts the ease of content creation, the complexity of artifacts, creation with contextual appropriateness, amplification of authorial power, reduction of the authorial burden, and the ability to modify complex artifacts.

To explore how statistical and symbolic approaches intersect with authorability, we will focus on asset creation and narrative design. Each of these areas brings a different aspect of authorability into perspective. Asset creation, such as generating textures or geometry, focuses on ease of creation, the amplification of authorial power, and the reduction of authorial burden. Narrative design requires contextual appropriateness and creating and modifying complex artifacts.

1) *Statistical Approaches*: With plenty of examples and the algorithms explored by the computer vision community, statistical methods are well-suited to the needs of generating visual [44] and audio assets [45]. On the other hand, authoring stories or social behavior relies heavily on context, temporarily-disconnected information, and consistency as judged by the player. Even the current best of natural language generation models and dialogue systems have difficulty in keeping consistent context or answering questions that require small amounts of reasoning.

2) *Symbolic Approaches*: Asset authoring with symbolic approaches works well with tasks that can be abstracted into portions and annotated with descriptors like puzzle design [46] or procedural music composition [47] but perform poorly in generating textures and sound files. The contextual nature of, and through lines required for, social [48] and story games [49] plays to the strengths of symbolic approaches and its application to this space has a rich history.

F. Other Challenges

Aside from the comparisons made in the previous section, there are many intersections and differences to be explored each with their own impacts on game AI. One with the most impact on games is how the AI system encapsulates design. In data-driven approaches, the act of creating and employing the system is completely design-free and the design knowledge

is implicit within a trained model or policy. Other systems are dependent on design either in their construction [15] or in their deployment. Where data is abundant, solutions can often be found by applying engineering and computation. Symbolic approaches are less able to leverage data to find solutions. When data is scarce or absent, the ability of each area to solve the task is reversed: statistical approaches tend to not find acceptable solutions and symbolic ones can leverage human-derived domain knowledge to solve problems.

Environmental impact and reproducibility are problematic for both approaches to varying degrees. Deepmind’s AlphaStar Final serves as an example for both because it required an enormous amount of power to train (though a large portion of that energy was offset renewably [50]) and consequently expensive to reproduce¹. Additionally, work done on this scale by corporations often contains proprietary techniques that are not sufficiently described and the data required for recreation is not made available. As these costly results are extremely impressive, they set a strong example for future research which cannot be sustainably or practically adopted by independent and academic institutions. On the other hand, symbolic approaches have yet to achieve success at this scale and leave their environmental impacts and reproducibility unknown.

There are many other challenges ready for future analysis including using learning as a hammer when not every problem is a nail; impacts of AI-based game design on the knowledge level [52]; and the function of hierarchy in reasoning and learning on game design.

IV. HYBRID APPROACHES AND FUTURE DIRECTIONS

As argued above, symbolic methods in AI-based games tend to help achieve meaningful play, trust, and other important factors that involve the player interacting with the AI processes themselves. Such methods provide the interpretative affordances needed to build actionable mental models and are conducive to high agency experiences. Statistical methods tend to be successful when AI is used to produce content where the process involved to create it is not central to the player experience, and authoring. When statistical methods are used as part of the gameplay loop, such as with generating dialogue, there is a risk that the system will present information that invalidates a player’s mental model. However, this is not always a problem when surprise is central to the desired aesthetic (as with AI Dungeon 2). Ultimately, designers should carefully choose what techniques they employ for what task they are trying to solve.

That said, these two approaches need not be considered as entirely separate. Connections between symbolic and statistical systems (i.e. hybrid intelligent systems) have taken forms like fuzzy logic [53] and fuzzy neural networks [54].

¹Estimated training cost of approximately \$13,000,000 based on Google Cloud prices taken from 21 May 2020 from <https://cloud.google.com/tpu/pricing> and 44 days each for training 12 agents on 32 3rd generation TPUs [51].

These systems represent a pipe-lined approach to problem-solving where a neural network and a fuzzy logic system send information between themselves but exist within their own islands of algorithmic and knowledge representation. While useful for some tasks, connecting these systems together does not advance our collective understanding of AI techniques in the way a true melding of two approaches. Techniques such as probabilistic soft logic [55], Markov logic networks [56], and Bayesian logic [57] represent aspects of symbolic logic and connectionist approaches in a single algorithm and mode of knowledge representation. Even with these capable mixed systems, AI-based game design presents many AI challenges simultaneous with mixed representations, scales, and evaluation criteria. To comprehensively address these challenges in a way that systems can reason about and impact one another, a systematic approach or framework that allows for the types of reasoning necessary for each of the related challenges while allowing for inter-system communication may be necessary. While other fields have developed promising solutions to this problem [58], game AI could leverage integrated systems [59], [60] or perhaps artificial general intelligence [61] as potential solutions.

While evocative, comprehensive solutions to game AI may not be possible. Each subproblem has its own domain and representation needs that may not generalize or be legible by other problems. Because they all could have this quality, designers should deeply consider their approaches in each subproblem and use the technique that is the best fit for that space.

V. CONCLUSION

We believe that games that forefront AI, by requiring the player to build a mental model of their operation, are beneficial as they expand our conception of what games can be, and are meaningful cultural artifacts in themselves. This paper was written to promote this AI-based game design research and to lay out a roadmap to the techniques and challenges within the space. We specifically compared and contrasted statistical and symbolic approaches to using AI in this space, and tended to conclude that symbolic approaches are more conducive to player understanding in AI-based games.

In the end, making games, especially with AI, is a complicated process with many technical challenges. While symbolic approaches may tend to assist with mental model building, the design considerations addressed above can certainly be broken into subproblems where statistical problems can be very useful. Furthermore, hybrid systems show great promise.

REFERENCES

- [1] M. P. Eladhari, A. Sullivan, G. Smith, and J. McCoy, “Ai-based game design: Enabling new playable experiences,” *UC Santa Cruz Baskin School of Engineering, Santa Cruz, CA*, 2011.
- [2] M. Treanor, A. Zook, M. P. Eladhari, J. Togelius, G. Smith, M. Cook, T. Thompson, B. Magerko, J. Levine, and A. Smith, “Ai-based game design patterns,” 2015.
- [3] V. South, “Left 4 dead [video game],” *Bellevue, Washington: Valve Corporation*, 2008.

- [4] M. Booth, "The ai systems of left 4 dead," in *Artificial Intelligence and Interactive Digital Entertainment Conference at Stanford, 2009*, 2009.
- [5] M. Suznjecic, M. Matijasevic, and J. Konfic, "Application context based algorithm for player skill evaluation in moba games," in *2015 International Workshop on Network and Systems Support for Games (NetGames)*. IEEE, 2015, pp. 1–6.
- [6] T.-H. D. Nguyen, Z. Chen, and M. S. El-Nasr, "Analytics-based ai techniques for a better gaming experience," *Game AI Pro*, vol. 2, pp. 481–500, 2015.
- [7] A. R. Gagné, M. S. El-Nasr, and C. D. Shaw, "A deeper look at the use of telemetry for analysis of player behavior in rts games," in *International Conference on Entertainment Computing*. Springer, 2011, pp. 247–257.
- [8] B. G. Weber, M. John, M. Mateas, and A. Jhala, "Modeling player retention in madden nfl 11," in *Twenty-Third IAAI Conference*, 2011.
- [9] N. Wardrip-Fruin, *Expressive Processing: Digital fictions, computer games, and software studies*. MIT press, 2009.
- [10] M. Kreminski and N. Wardrip-Fruin, "Gardening games: an alternative philosophy of pcg in games," in *PCG Workshop*, 2018.
- [11] J. Togelius, E. Kastbjerg, D. Schedl, and G. N. Yannakakis, "What is procedural content generation? mario on the borderline," in *Proceedings of the 2nd international workshop on procedural content generation in games*, 2011, pp. 1–6.
- [12] M. Véron, O. Marin, and S. Monnet, "Matchmaking in multi-player on-line games: studying user traces to improve the user experience," in *Proceedings of Network and Operating System Support on Digital Audio and Video Workshop*, 2014, pp. 7–12.
- [13] M. Carter and M. R. Gibbs, "esports in eve online: Skullduggery, fair play and acceptability in an unbounded competition." *FDG*, vol. 2013, no. May, pp. 47–54, 2013.
- [14] Maxis, "Sims 4 [video game]," *Redwood Shores, CA: Electronic Arts*, 2008.
- [15] T. Adams and Z. Adams, "Slaves to Armok: God of Blood Chapter II: Dwarf Fortress," 2006.
- [16] M. Mateas and A. Stern, "A behavior language for story-based believable agents," *IEEE Intelligent Systems*, vol. 17, no. 4, pp. 39–47, jul 2002.
- [17] R. Koster, *A Theory of Fun*. O'Reilly Media, Inc., 2004.
- [18] E. Zimmerman and K. Salen Tekinbaş, *Rules of Play*. MIT Press, 2003.
- [19] B. Romero and I. Schreiber, *Challenges for Game Designers*. Charles River Media, 2008.
- [20] M. Mateas, "Expressive ai: Games and artificial intelligence." in *DiGRA Conference*, 2003.
- [21] H. Games, "No man's sky [video game]," *Guildford, United Kingdom: Hello Games*, 2016.
- [22] D. Gunning, "Explainable artificial intelligence (xai)," *Defense Advanced Research Projects Agency (DARPA), nd Web*, vol. 2, 2017.
- [23] C. Crawford, *On Game Design*. New Riders Publishing, 2003.
- [24] M. Treanor and M. Mateas, "An account of proceduralist meaning," in *Proceedings of the 2013 DiGRA International Conference: DeFragging GameStudies*, 2013.
- [25] M. Mateas, "A preliminary poetics for interactive drama and games," *Digital Creativity*, vol. 12, pp. 140–152, 2001.
- [26] R. Games, "Grand theft auto v [video game]," *Edinburgh, United Kingdom: Rockstar North*, 2013.
- [27] Maxis, "Simcity [video game]," *Redwood Shores, CA: Electronic Arts*, 2013.
- [28] A. Willmott, "Inside the glassbox," *Game Developer's Conference*, 2012.
- [29] B. Samuel, "Crafting stories through play," Ph.D. dissertation, UC Santa Cruz, 2016.
- [30] A. Reed, "Changeful tales: design-driven approaches toward more expressive storygames," Ph.D. dissertation, UC Santa Cruz, 2017.
- [31] G. Smith, J. Whitehead, and M. Mateas, "Tanagra: A mixed-initiative level design tool," in *Proceedings of the Fifth International Conference on the Foundations of Digital Games*, 2010, pp. 209–216.
- [32] G. Neff and P. Nagy, "Automation, algorithms, and politics—talking to bots: Symbiotic agency and the case of tay," *International Journal of Communication*, vol. 10, p. 17, 2016.
- [33] A. Schlesinger, K. P. O'Hara, and A. S. Taylor, "Let's talk about race: Identity, chatbots, and ai," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 2018, pp. 1–14.
- [34] J. Dastin, "Amazon scraps secret ai recruiting tool that showed bias against women," *San Fransico, CA: Reuters. Retrieved on October*, vol. 9, 2018.
- [35] G. Veletsianos, "Cognitive and affective benefits of an animated pedagogical agent: Considering contextual relevance and aesthetics," *Journal of Educational Computing Research*, vol. 36, no. 4, pp. 373–377, 2007.
- [36] L. Swartz, "Why people hate the paperclip: Labels, appearance, behavior, and social responses to user interface agents," Ph.D. dissertation, Stanford University Palo Alto, CA, 2003.
- [37] S. Cass, "Code: Debugging the gender gap," *IEEE Spectrum*, vol. 52, no. 7, pp. 25–25, 2015.
- [38] C. Nass and C. Yen, *The man who lied to his laptop: What we can learn about ourselves from our machines*. Penguin, 2010.
- [39] M. Treanor, J. McCoy, and A. Sullivan, "Social play in non-player character dialog," in *Eleventh Artificial Intelligence and Interactive Digital Entertainment Conference*, 2015.
- [40] G. Smith and J. Whitehead, "Analyzing the expressive range of a level generator," in *Proceedings of the 2010 Workshop on Procedural Content Generation in Games*, 2010, pp. 1–7.
- [41] N. Walton, "Ai dungeon 2 [video game]," 2019.
- [42] G. Nelson, "Inform 7," 2006.
- [43] D. Isla, "Handling complexity in the halo 2," *Game Developers Conference*, 2005.
- [44] C. Li and M. Wand, "Precomputed real-time texture synthesis with markovian generative adversarial networks," in *Computer Vision – ECCV 2016*, ser. Lecture Notes in Computer Science. Springer International Publishing, 2016, pp. 702–716.
- [45] A. McDonagh, J. Lemley, R. Cassidy, and P. Corcoran, "Synthesizing game audio using deep neural networks," in *2018 IEEE Games, Entertainment, Media Conference (GEM)*, 2018, pp. 1–9.
- [46] A. M. Smith, E. Andersen, M. Mateas, and Z. Popović, "A case study of expressively constrainable level design automation tools for a puzzle game," in *Proceedings of the International Conference on the Foundations of Digital Games*, ser. FDG '12. Association for Computing Machinery, 2012, pp. 156–163.
- [47] D. L. Brown, "Mezzo: An adaptive, real-time composition program for game soundtracks," in *Eighth Artificial Intelligence and Interactive Digital Entertainment Conference*, ser. FDG '12, 2012.
- [48] J. McCoy, M. Treanor, B. Samuel, A. A. Reed, M. Mateas, and N. Wardrip-Fruin, "Social story worlds with comme il faut," *Computational Intelligence and AI in Games, IEEE Transactions on*, vol. 6, no. 2, pp. 97–112, 2014.
- [49] A. A. Reed, J. Garbe, N. Wardrip-Fruin, and M. Mateas, "Ice-bound: Combining richly-realized story with expressive gameplay," in *In Proceedings of the 9th International Conference on the Foundations of Digital Games*, 2014, p. 8.
- [50] R. Porat and U. Hölzle, "Google environmental report 2019," p. 64, 2019.
- [51] O. Vinyals, I. Babuschkin, W. M. Czarnecki, M. Mathieu, A. Dudzik, J. Chung, D. H. Choi, R. Powell, T. Ewalds, P. Georgiev *et al.*, "Grandmaster level in starcraft ii using multi-agent reinforcement learning," *Nature*, vol. 575, no. 7782, pp. 350–354, 2019.
- [52] A. Newell, "The knowledge level," *Artificial Intelligence*, vol. 18, pp. 81–132, 1982.
- [53] L. Zadeh, "Fuzzy logic," *Computer*, vol. 21, no. 4, pp. 83–93, 1988.
- [54] C.-T. Lin and C. S. G. Lee, "Neural-network-based fuzzy logic control and decision system," *IEEE Transactions on Computers*, vol. 40, no. 12, pp. 1320–1336, 1991.
- [55] S. H. Bach, M. Broecheler, B. Huang, and L. Getoor, "Hinge-loss markov random fields and probabilistic soft logic," *The Journal of Machine Learning Research*, vol. 18, no. 1, pp. 3846–3912, 2017.
- [56] M. Richardson and P. Domingos, "Markov logic networks," *Machine Learning*, vol. 62, no. 1, pp. 107–136, 2006.
- [57] K. A. Andersen and J. N. Hooker, "Bayesian logic," *Decision Support Systems*, vol. 11, no. 2, pp. 191–210, 1994.
- [58] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, "Ros: an open-source robot operating system," in *ICRA workshop on open source software*, vol. 3, no. 3.2. Kobe, Japan, 2009, p. 5.
- [59] A. Aamodt and M. Nygård, "Different roles and mutual dependencies of data, information, and knowledge — an AI perspective on their integration," *Data & Knowledge Engineering*, vol. 16, no. 3, 1995.
- [60] P. Langley, "Cognitive architectures and general intelligent systems," *AI Magazine*, vol. 27, no. 2, pp. 33–33, 2006, number: 2.
- [61] D. Perez-Liebana, S. Samothrakis, J. Togelius, T. Schaul, and S. M. Lucas, "General video game AI: Competition, challenges and opportunities," in *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.