

Making Aesthetic Choices

Kenneth M. Kahn
MIT Artificial Intelligence Laboratory
545 Technology Square
Cambridge, Massachusetts 02139, USA
Arpanet address: KEN@MIT-AI

A framework is presented for making choices that are primarily constrained by aesthetic, as opposed to pragmatic, considerations. An example of the application of this framework is a computer system called "Ani", capable of making simple computer animation in response to high-level incomplete story descriptions. Aesthetic choice is presented as a parallel computation in which each *choice point* gathers together and evaluates *suggestions*. When faced with difficulties these choices can be *postponed*. The order in which inter-dependent choices are made is influenced by the *focus* of the problem.

Introduction

People are often faced with choices that are under-constrained by considerations of utility, cost, simplicity and efficiency. In many of these cases, one can just choose arbitrarily between those alternatives which satisfy the pragmatic constraints. There remain many situations where this is inadequate however, where instead one wants to select the alternative that is the most beautiful, elegant, interesting, or that conforms to a particular style. Aesthetic considerations are important, sometimes even dominate, in tasks that vary from exploring mathematics to designing block diagrams, from writing an IJCAI paper to making an animated film.

An assumption of this paper is that making aesthetic choices is a knowledge intensive computational process. The interesting questions are what the knowledge is and how it is represented, organized, and used. These questions are addressed in [1]. This paper presents a general framework for how knowledge is *used* to create an aesthetic object (a detailed description of an object such as a film, a story or a proof) that is consistent and coherent. This means that ideally every choice should be justified by as much relevant knowledge combined in as reasonable a manner as possible.

Aesthetic choices cannot be made independently, since each decision *constrains* subsequent choices. The problem addressed here is really one of *generating* a reasonable set of constraints rather than trying to find a solution that *satisfies* a set of constraints. The control structure, the order in which choices are made, therefore becomes very important. Since early choices constrain the later ones, one should be careful that the early constraints are satisfiable and desirable.

The framework to be presented is based upon the notion of an *actor*. Actors are computational entities that communicate by passing messages. Each actor, containing both state and program, has the full power of a digital computer. Hewitt and others have argued for the usefulness of actors in the construction of large AI systems [2]. Actors provide a modular and convenient means of implementing systems built on the framework about to be presented.

The model presented here begins with the exploration of a set of related choices. Each *choice point* starts by gathering up *suggestions* by asking the elements of its "choice" for suggestions. The choice points evaluate the suggestions: combining closely related ones, classifying any conflicts, and noting any missing information. The choice points which run into trouble ask permission to *postpone* themselves until more has been decided, while the satisfied ones make their choices based on the suggestions gathered. Permission to postpone is granted depending upon the reason for desiring postponement, the state of other choice points, and whether the choice in question is part of the *focus* of the object being created. The hope is that when the choice point is awakened more information (or constraints) will be available and the choice will be easier. If permission to postpone is refused, more effort is expended despite the difficulty (e.g., more suggestions are gathered or conflicts are resolved). If no more progress is possible then a choice is made and its justification is recorded.

An Example

As part of my doctoral research I implemented a system called "Ani" that makes aesthetic choices ([1], [3]). It creates simple computer animation in response to high-level incomplete descriptions. The user describes the personalities of the characters (e.g. shy), their physical characteristics (e.g. ugly, powerful), the relationship between the characters (e.g. hates), and their interactions (e.g., one prevents another from meeting a third). The user also describes the style of film desired (e.g. varied, simple, obvious). Ani's job is to decide where the characters should be placed, how they should move, and what they should do.

Ani was given a simple version of the story of Cinderella to animate. One of many of Ani's problems is to choose typical speeds for Cinderella, her stepmother, the fairy godmother, and the Prince which are in keeping with their personalities, physical characteristics, and the desired style of the film. The relative speeds of the characters should, in turn, be in keeping with the relationships and comparisons of the characters. A simple example of aesthetic choice is the determination of these typical speeds. This choice is aesthetic because the character of movement is important in assessing the quality of animation, because it is interdependent with other choices, and because it is only weakly constrained by pragmatics.

Choice points are created to represent the process of picking typical speeds for each character. The choice point for the stepmother's speed, for example, begins by asking each of the descriptors of the stepmother for suggestions for her speed. Only the description "powerful" replies and suggests a high speed. The choice point is not happy with just that because there are not enough strong suggestions. So it asks permission to be postponed to wait for more information to become available and it is granted.

When the choice point for the stepmother's speed is reawakened, it inspects its record of its previous activations. It then asks the choice points for the relative speeds of the stepmother and the other characters for suggestions. These choice points are created in response to this request and they choose values (e.g., that the stepmother be faster than Cinderella because she dominates Cinderella and differs from her), but cannot make any concrete suggestions since none of the characters have speeds yet. The choice point for the stepmother's speed asks permission to postpone to wait for the speeds of the others to be determined and it is granted.

The choice points for the other characters also ask and are granted permission to postpone. This could potentially lead to a deadlock in which the four choice points wait for the others to make a decision. One of the reasons the choice points don't just postpone themselves, but instead ask permission first, is to avoid this type of situation. A postponement manager keeps track of the situation and will not grant someone permission to postpone for the same reason twice. A common exception to this is when the choice point is waiting for other choice points to finish and at least one of these is making progress. In this case, no one is making progress so the postponement manager must refuse permission to at least one of the choice points.

Ani is built upon the principle that as few decisions as possible be determined arbitrarily. The decision as to who should be refused permission to postpone has too many consequences to be determined by something like who asks first. Instead the postponement manager asks the *focus* which, in this case, indicates that conveying the personality of Cinderella is important. The choice point for Cinderella's speed is refused permission to postpone and the deadlock is broken. This means Cinderella's speed will be based on the description of Cinderella without being constrained to be faster or slower than the others.

The choice point for the stepmother's speed finally gets suggestions from the relative choice points. It discovers conflicts with one of these suggestions and the earlier suggestion it had received from "powerful" and postpones again. Upon being resumed the choice point asks the descriptions of the film's style for suggestions and receives them from the moderate variety level, high energy level, and low flashiness. Unfortunately they do not all agree and so the choice point postpones one more time.

When it is reawakened it discovers that there are no more sources of suggestions and proceeds with what it has. First it attempts to make compromises between the conflicting suggestions and makes one that in turn generates a new conflict. Excuses are found for rejecting some of the conflicting suggestions. The choice point finally picks a high speed for the stepmother and saves away a justification for this choice.

The Major Mechanisms

The model of aesthetic choice presented here consists of (a) suggestions (together with a means of gathering, combining, and resolving conflicts between them), (b) choice points to organize and record progress on choices, (c) a means of deciding when to postpone (and resume) work on a choice, and (d) a means of focusing upon the more critical elements.

Suggestions Suggestions are like rules, advice, or hypotheses that are rejected, modified, combined, and compromised. Suggestions differ from facts in that they typically conflict among themselves, are rejected or compromised, and possess degrees of reliability or strength. Suggestions, are often ignored or modified with minor consequences. Ani is constantly faced with contradictory suggestions and spends much time detecting, classifying, and resolving these conflicts. Conflicts are resolved by making compromises and by rejecting some of the suggestions.

Choice Points A choice point represents the exploration of a choice. Each choice point responds to messages asking it to accept suggestions, to combine suggestions, and to make choices. A choice point maintains in its database records of the best suggestions so far, conflicts between suggestions, postponements, and the sources of suggestions that have already been tapped. Choice points decide whether a set of suggestions is adequate, more work needs to be done, or permission to postpone should be sought.

Conflict Resolution If a choice point finds no conflicting suggestions for a value, the decision is simple and it just picks the value suggested. More typically there are many suggestions and they don't all agree. It is important that these conflicts be resolved as sensibly as possible if there is to be any coherence. The general approach is to use the heuristic with the strongest criterion of applicability. The criteria for making a choice between two conflicting suggestions are (a) the strength of the suggestions, (b) the degree of compatibility with the other suggestions, (c) the extent to which the sources of the conflicting suggestions have other suggestions followed, and (d) the number and kind of sources of the suggestions. If the difference between the conflicting suggestions is great along any of these dimensions, the associated rule is used.

Postponement Postponement is an important component in the making of choices that are interdependent. One choice adds new constraints that strongly influence later choices. Aesthetics only

emerges when there is a set of inter-related choices; one hesitates to call an isolated atomic choice aesthetic. Because of this the order in which subproblems are attacked is very important. We want the relatively straight-forward choices to be made first because they already have a strong consistent justification. We want the choices with the least basis for a decision to be postponed as long as possible in the hope that by the time a choice finally has to be made additional constraints will have been added. Otherwise the choice would have to be made on a relatively arbitrary basis and the consequences (the additional constraints), if chosen badly, will cause trouble. Additional constraints often make the choice between the conflicting alternatives easier. The result of this control strategy is that the final product is more coherent and less arbitrary.

The postponement mechanism is designed so that the choices with the most justification are made first, then those that are difficult due to conflicting suggestions, followed by those with the least to go on. The order in which postponed choice points are forced to continue is under-determined by these criteria. The unordered choices could be made in parallel, however the interdependencies between them are such that the order of execution could adversely affect some of the decisions.

Focus To help avoid this arbitrariness, we have a structure called the *focus* of the object being created. It describes the parts or aspects that are primary or deserve emphasis (e.g., Cinderella's personality, her relationship with her stepmother, the second scene and so on). Those choices that relate to the focus (e.g., Cinderella's speed) tend to be made first and so are more likely to be self-consistent and effective because they are not constrained by choices yet to be made.

Minimization of Arbitrariness and Maximization of Coherence The goals of this process of choosing is that the arbitrariness of each choice be minimal and the coherence of a set of choices be maximal. In Ani's filmmaking this means that each choice of any consequence should be compatible with the description of the film and Ani's general animation knowledge. Arbitrariness is minimized by use of knowledge in the form of suggestions and from the guidance of the focus. Coherence in this context means that the choices for the relative dynamics of the characters be satisfied and that the choices of activities of the characters be self-consistent and compatible with the choices of the character dynamics. Coherence results from the control structure that postpones troublesome choices and that focuses on the relatively more important ones. The aesthetics of Ani's films are a result of this striving for coherence, this minimization of arbitrariness, and the currently small amount of knowledge about animation and emotions that Ani brings to bear.

Relationship to Other Work

The most illuminating comparisons of this work with others are briefly described below. A more complete discussion can be found in [1].

Meehan's Talespin One system that creates objects that are usually judged primarily on an aesthetic basis is Meehan's "Talespin" [4]. It makes up fables about talking bears, birds, and so forth. Talespin is told the initial conditions (e.g., a bear is hungry and a bird in a tree is sleepy) and spins a tale based upon a high-level simulation. The characters of the story generate plans to satisfy their needs and desires. Executing these plans causes the characters to interact forming the substance of Talespin's stories. The resulting stories are plausible but typically not very interesting or aesthetic since there is no higher-level structure to the stories and no notion of style or focus. Meehan's research, despite the problem domain, is primarily concerned with plans and symbolic simulation, not with aesthetics.

Talespin is often faced with aesthetic choices: what name to give the bear, where the crow should be, what kind of food should be available, and so on. These story aspects are aesthetic in our

culture because they are aspects that we try to interpret as the result of purposeful choices made by the author. Not every aspect of a story is aesthetic. The number of letters in the name of a bear, the locations of the word "crow" on a printed page, and the amount of ink used in a description of food are not typically considered aesthetic choices. And writers rarely make deliberate choices regarding these aspects.

Talespin, however, rarely makes deliberate choices for those aspects that are normally considered aesthetic. Instead it either is told by the user, "chooses" randomly, chooses based upon a symbolic simulation of the characters involved, or chooses that which will help give the story a particular moral. The first two cases are not choice-making at all. The third is an interesting alternative to (or supplement of) the model presented herein. The difficulty is that simulation is concerned with plausibility, with having the components (the characters in the story) behave in a reasonable way. Aesthetics, on the other hand, is concerned with creating the world that such a simulation occurs in, with determining the rules of interaction, the goals of the components, and the initial conditions. The moral-fulfilling aspect of Talespin is more relevant to aesthetic choice. Unfortunately, it is a small, under-developed part of a large system.

Lenat's Artificial Mathematician One of the more creative AI systems of late is Lenat's AM [5]. AM starts with very elementary concepts such as sets, composition, and equality. AM creates many new concepts, makes conjectures, and discovers new aspects of the original concepts. AM's heuristics propose tasks to perform that are explorations of this space and these are placed upon an agenda. The tasks on the agenda that are most interesting are performed first.

AM is probably the research most related to the work described here. This may seem odd since AM's domain is elementary mathematics which is very formal and well understood --- almost the antithesis of aesthetics and art. Mathematics is formal and good models of it do exist, but, as Lenat points out, the exploration of mathematics, the heuristics that guide one in making conjectures, in constructing new concepts, and in evaluating them, in other words, the doing of mathematics is neither formal nor well-understood. The problems that Ani and AM address are both weakly specified: making good animation and discovering interesting mathematics. Both systems construct structures out of a very large space of possibilities. Since the creations of AM and Ani are not judged as right or wrong --- but as interesting or dull, plausible or implausible, good or bad --- aesthetic choice plays a crucial role. Both AM and Ani are knowledge-oriented, in contrast with other approaches which are simulation-based, search-oriented, or based upon a few very general pieces of knowledge. Both systems spend a considerable portion of their time deciding what to do, in addition to doing it.

There are many differences between the two systems, of course. The proposed tasks of AM can be viewed as suggestions from various heuristics as to what should be explored, however they are not treated as described above (e.g., combining, compromising, relating, and so on). AM has a focus of attention which tends to keep AM from jumping from topic to topic. Ani's focus instead influences the relative priority of the elaboration of the different parts. The control structures of the two programs are very different. AM keeps executing the most interesting task on its agenda, while Ani jumps from choice to choice on the basis of their past difficulties and the focus. AM executes its tasks in a fairly straight-forward manner, while Ani works on a choice by gathering, combining, and rejecting suggestions, noting and classifying conflicts, making compromises, and searching for more suggestions.

Discussion

Generality To some it is plausible that the model of aesthetic choice presented here is applicable to any set of choices that are not primarily constrained by pragmatics, whether or not

there are aesthetic considerations. Others think the framework is only appropriate for the kinds of choices that Ani makes in creating simple films and are skeptical of any claims of generality. This issue is separate from the question of how easy it is to apply the framework described here to a particular domain. It may be very difficult to apply it, say, to the production of oil paintings, not because it is a bad framework, but because it is very unclear how to structure the space of choices, what the sources of suggestions are, where to get the necessary knowledge in a complete and detailed enough form and so on.

Nonetheless, one should be suspicious of claims of generality when something has been applied only to one example. In progress is another test of the model of aesthetic choice. The model is being applied to the design of block diagrams that illustrate papers and lectures. The problem is to choose locations and dimensions for the boxes, to decide where links between boxes should originate and end, where to place labels and in what font, and so on. Preliminary results indicate that the notions of suggestions, postponement, and focus have their natural places in this application. This will be described in detail in a subsequent publication.

What is Missing The framework presented for making aesthetic choices is simple. Several components are missing for building high performance systems. No means of undoing poor choices is included, for example. Also the structure of the set of choices explored is given to the system, while ideally that too should be determined by the system. The examples discussed are limited to the selection of a single value for an element of a complex description. Many aesthetic choices are not the selection of values, but the generation of complex descriptions. Much of Ani's time, for example, is spent deciding what should happen in the scenes of the film. The problem is to create a description of the activities that should occur in a particular portion of a scene. While there are many similarities and parallels between this process of choosing activities and the selection of values, some special mechanisms were needed and are described in [1].

Where to Go from Here One avenue of future research is to attempt to apply the framework presented here to other domains to discover its shortcomings and strengths. Part of this involves extensive testing of computer systems like Ani. The focus, conflict resolution rules, or postponement criteria can be modified and the effects can be observed.

Another avenue of research is to attempt to fill in the missing components of the framework. The inability to undo previous decisions, for example, can be quite serious. The decision making process needs to be able to recognize when things are not right and to generate directed criticism to the choice points responsible. Choice points need to be extended to take such criticism into account and redo just that portion which needs changing.

Finally, one last avenue of further research is to formulate a general theory of aesthetics and incorporate it into the model. If formal aesthetic judgments can be made about the description and justification of a creation, then their discovery and inclusion in a model of aesthetic choice is a fascinating possibility.

Bibliography

- [1] Kahn, K. *Creation of Computer Animation from Story Descriptions*, MIT PhD thesis 1979 (AI Lab TR in progress)
- [2] Hewitt, C. "Viewing Control Structures as Patterns of Passing Messages", *Artif. Intell.*, 8:3 (1977) 323-364.
- [3] Kahn, K. "Ani: An Example of Computational Creativity", *Proceedings of the AISB/GI Conference*, Germany, July 1978
- [4] Meehan, J. *The Metanovel: Writing Stories by Computer*, Yale University Computer Science Research Report 74, 1976
- [5] Lenat, D. *AM: An AI Approach to Discovery in Mathematics as Heuristic Search*, Stanford AI Lab Memo AIM-286, 1976