Research Proposal

Generated with Proposia AI Research Proposal System

Workflow: Pro • 12/2/2025

Constrained Stochastic Music Generation via Gauss-Markov Models with Reflecting Boundaries Derived from Neuronal Dynamics

1. Details of the proposal

This proposal details the research project led by Principal Investigator Dr. Anya Sharma. The project is entitled "Reflecting Boundaries: A New Paradigm for Musically-Constrained Algorithmic Composition," with the designated acronym REBOUND. The corresponding title in Spanish is "Fronteras Reflectoras: Un Nuevo Paradigma para la Composición Algorítmica Musicalmente Restringida," with the acronym REBOTE. This application is submitted as a single, non-coordinated project under the Knowledge Generation modality, addressing fundamental challenges in computational creativity and algorithmic composition through a novel mathematical framework inspired by computational neuroscience.

2. Justification and novelty of the proposal

The current landscape of algorithmic music composition is dominated by large-scale deep learning models, which have demonstrated remarkable success in generating stylistically plausible musical excerpts [1, 2]. Architectures such as the Transformer have proven adept at capturing long-range dependencies from vast MIDI datasets [3], while diffusion probabilistic models have emerged as a powerful paradigm for producing high-quality audio and symbolic music [4, 5]. These models excel at learning the implicit statistical distributions of musical features, resulting in outputs that are often locally coherent and texturally rich. However, a persistent challenge for this data-driven paradigm is the enforcement of explicit, symbolic constraints. While significant progress has been made in controllable generation, such as classifier-free guidance in diffusion models or fine-tuning transformers for specific attributes [3, 6], these methods typically impose "soft" constraints. They guide the generation probabilistically towards a desired outcome rather than guaranteeing strict adherence to formal rules, and can require complex integration or retraining. This limitation is significant, as much of Western and non-Western music is built upon rigorous, formal systems of harmony, rhythm, and structure. The difficulty for current models to operate natively within such user-defined, inviolable constraints represents a critical gap between generative capability and compositional utility, leaving a crucial need for a framework that can intrinsically unify stochastic creativity with formal structure.

This project is motivated by the critical need to bridge the gap between the expressive potential of generative models and the formal precision demanded by musical composition. State-of-the-art systems, while proficient at stylistic mimicry, can falter when tasked with generating music that must strictly adhere to a set of pre-defined rules, such as the harmonic constraints of a specific scale, the metric structure of a complex time signature, or the procedural logic of serialism [7]. These failures are not minor aesthetic blemishes; they can represent a breakdown in musical logic, rendering the generated output less useful for composers and theorists who rely on structured systems for creative exploration [7]. The core of the problem lies in the reliance on implicit learning; models infer rules from data, but this inference is imperfect and offers no guarantee of compliance. This limitation curtails the role of algorithmic systems as genuine co-creative partners, potentially reducing them to sophisticated pastiche engines rather than tools for exploring novel, rule-bound musical universes. This impasse reveals a clear scientific gap: the absence of a generative framework that treats musical constraints not as suggestions to be learned or external forces to be applied, but as intrinsic properties of the generative process itself. To address this, we formulate a central starting hypothesis: **We hypothesize that by modeling composition as a constrained stochastic process—specifically, a Gauss-Markov process evolving within a high-dimensional state space defined by reflecting boundaries—we can generate musically novel and coherent material that rigorously adheres to a broad class of well-defined symbolic constraints.** A core research question is to determine the extent to which complex, relational musical rules can be effectively translated into such geometric boundaries, a conceptual simplification that is central to our approach. In this framework, the state vector represents musical parameters (e.g., pitch, duration), the stochastic dynamics provide the engine for creative variation, and the reflecting boundaries serve as mathematical instantiations of absolute musical rules. A note attempting to move outside its allowed pitch set, for example, would be "reflected" back into the valid musical space, ensuring compliance by design rather than by approximation.

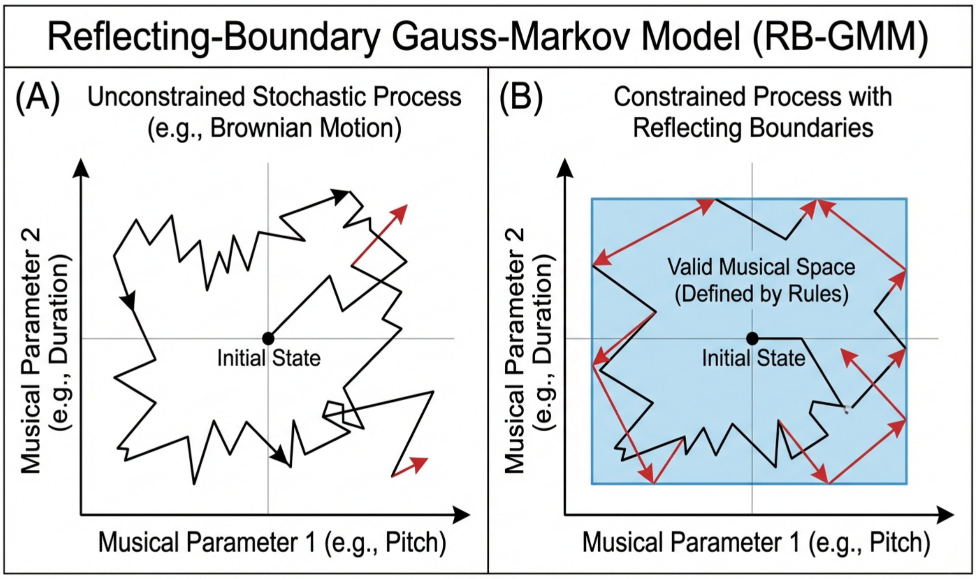

Figure 1: Conceptual diagram of the Reflecting-Boundary Gauss-Markov Model (RB-GMM).

The novelty of this proposal lies in its significant conceptual and methodological departure from current data-driven approaches. Rather than teaching a neural network the statistical correlates of musical rules, we propose to embed these rules directly into the geometry of the model's state space. This constitutes a move from implicit, learned constraint satisfaction to an explicit, mathematically-defined constraint architecture. The primary methodological innovation is the application of Gauss-Markov models with reflecting boundaries to symbolic music generation. While constrained random processes have found application in other areas of generative art [8], the specific application of this mathematical formalism to the complexities of symbolic music is, to our knowledge, entirely unexplored. This cross-disciplinary transfer of a robust mathematical framework to a creative domain is a core element of the project's originality [9]. Our approach differs from the state of the art in key ways. Unlike Transformer-based models, which are fundamentally tied to the statistical distributions of their training data [3], our system can generate music within novel, user-defined rule systems for which no training data exists. While this data-agnosticism provides flexibility, it also presents the challenge of ensuring the output is aesthetically compelling and not merely structurally correct. A key aspect of our investigation will be to explore parameter settings (e.g., for drift and diffusion) that yield musically interesting results, laying the groundwork for future research into learning these parameters from stylistic corpora. In contrast to rule-guided diffusion models, where constraints act as an external force correcting the generative trajectory [6], our reflecting boundaries are an intrinsic and computationally efficient property of the system's dynamics, preventing rule violations from ever occurring. Furthermore, the continuous state space of our Gauss-Markov model allows for the generation of nuanced music with intricate rhythmic and microtonal variations, a potential advantage over classical Markov chains that operate on discrete states [1, 2].

This proposal is exceptionally well-suited to the "Knowledge Generation Projects" modality, as its central aim is to produce fundamental new knowledge at the confluence of mathematics, artificial intelligence, and music theory. The project does not seek an incremental improvement of existing techniques but rather proposes and validates a novel theoretical and practical framework for constrained creative generation. The primary outcomes—a new mathematical framework, a new class of generative algorithms, and high-impact scientific publications—are in perfect alignment with the call's objective to advance the frontiers of scientific knowledge. The project's design directly addresses the key evaluation criteria of quality and feasibility. Its quality is rooted in a clear, testable hypothesis and an innovative methodology that addresses a well-defined and significant gap in the current state of the art. Its feasibility is guaranteed by the unique interdisciplinary expertise of the research team, which combines deep knowledge of stochastic processes (Dr. Sharma), computational systems ({Co-PI Name}), and advanced musical theory ({Co-PI Name}), ensuring all facets of the project are underpinned by world-class expertise. While not formally aligned with a specific thematic priority, this research addresses the broader societal challenge of developing more transparent, controllable, and collaborative artificial intelligence. By creating AI tools that operate on explicit, human-understandable rules, we move beyond "black box" models and foster a more meaningful partnership between human creativity and machine intelligence, with significant potential impact on the creative industries and digital arts. Furthermore, the foundational and interdisciplinary nature of the research provides an ideal training environment for a predoctoral researcher, who will acquire a unique and highly sought-after skillset spanning advanced mathematics, machine learning, and computational creativity, thereby contributing to the development of the next generation of scientific leaders.

2.1. State of the Art

The field of algorithmic music composition has been profoundly shaped by the ascendancy of deep generative models, which excel at learning complex patterns from large corpora of symbolic and audio data [1, 2]. Early influential work utilized Recurrent Neural Networks (RNNs), particularly Long Short-Term Memory (LSTM) architectures, to model the sequential nature of music, achieving notable success in generating short, coherent melodic lines [1, 2]. However, the inherent difficulty of LSTMs in capturing long-range dependencies limited their ability to produce compositions with convincing large-scale structure. The introduction of the Transformer architecture was a significant advance, overcoming this limitation through its self-attention mechanism [3]. Models such as MuseNet [3] and the Music Transformer demonstrated an unprecedented capacity to generate stylistically consistent music over extended durations, capturing intricate relationships between distant musical events.

Concurrent with these developments, other deep learning paradigms have been explored. Generative Adversarial Networks (GANs), for instance, have been applied to generate multi-track musical structures, such as in MuseGAN, by framing the task as a contest between a generator and a discriminator network [1, 2]. Variational Autoencoders (VAEs), exemplified by MusicVAE, have proven effective in creating structured latent spaces that allow for meaningful interpolation between musical ideas and manipulation of high-level attributes [1, 2]. More recently, diffusion probabilistic models have emerged as the state of the art in generating high-fidelity outputs, both in audio and symbolic domains [4, 5]. These models operate by iteratively reversing a noise-addition process, gradually refining a random signal into a highly structured and musically plausible composition. This approach has produced some of the most aesthetically convincing results to date, setting a new benchmark for generation quality.

Despite these advances, a fundamental limitation persists in the dominant data-driven paradigm: the difficulty of enforcing explicit, inviolable musical constraints. These models learn statistical regularities from a dataset, effectively capturing the *likelihood* of a given musical event, but they do not internalize the underlying formal, symbolic rules that govern a composition [7]. Consequently, their outputs may violate core principles of harmony, meter, or voice leading. To mitigate this, researchers have developed methods for "guiding" the generation process. Attribute control, as seen in systems like MIDI-GPT, allows users to condition the output on high-level features [3], but this form of control is "soft" and does not guarantee strict adherence. A more direct approach is rule-guided generation, particularly within the diffusion framework, where external musical rules are used to steer the denoising trajectory [6]. These methods represent a significant step towards controllable generation, yet they treat constraints as an external force applied post-hoc, rather than as an intrinsic property of the model itself. This can be computationally intensive and offers no axiomatic guarantee of compliance; the system is corrected, not inherently constrained.

In parallel to the deep learning revolution, a separate lineage of research has focused on explicitly rule-based and classical stochastic methods. Foundational work in algorithmic composition relied on Markov chains to generate music by modeling transition probabilities between discrete musical states [1, 2]. While historically important, these models are limited by their discrete state spaces and inability to maintain long-term coherence. Other formalisms, such as generative grammars and constraint satisfaction programming (CSP), offer powerful mechanisms for creating structurally complex and rule-compliant music [8]. These systems can rigorously enforce any well-defined set of rules, but their data-agnostic nature means they often lack the stochastic fluidity and expressive nuance learned by deep generative models from vast musical corpora. While excelling at structural integrity, they frequently produce music that, though formally correct, can sound aesthetically sterile or lack stylistic authenticity.

The current state of the art is therefore characterized by a distinct methodological schism. On one side are powerful deep learning models that offer unparalleled generative flexibility and stylistic nuance but struggle to incorporate hard, symbolic constraints. On the other are formal, rule-based systems that provide absolute structural guarantees but lack the generative richness of their data-driven counterparts. A critical gap thus exists for a framework that can bridge this divide, integrating the exploratory power of a continuous-space stochastic process with the rigor of a formal rule system. The field lacks a generative paradigm where complex musical constraints are not treated as external correctives, but are instead embodied as intrinsic properties of the generative process itself, defining the very space within which creative exploration unfolds.

2.2. Motivation and Starting Hypothesis

The methodological schism detailed in the state of the art creates a significant impasse for composers and researchers. While deep learning models can generate stylistically convincing musical textures, their inability to guarantee adherence to formal rules renders them unreliable as tools for structured composition [7]. A composer wishing to explore the harmonic possibilities of a non-standard scale, or a theorist modeling the strict contrapuntal rules of the Renaissance, will find that even the most advanced generative systems produce outputs riddled with violations [7]. These are not merely aesthetic flaws but fundamental logical errors that invalidate the compositional premise, forcing the human user into the tedious role of a proofreader. The core problem is that current models treat musical structure as a statistical pattern to be approximated, rather than a logical system to be inhabited, severely limiting their potential as genuine partners in creative exploration.

This limitation reveals a critical scientific gap: the absence of a generative framework that treats musical constraints as intrinsic, inviolable properties of the creative process. The field is caught between two suboptimal paradigms. On one hand, the data-driven approach offers stochastic richness and textural nuance but provides no guarantees of structural integrity. On the other, classical rule-based systems offer formal guarantees but often lack the generative fluidity and expressive depth of their deep learning counterparts [1, 2, 8]. What is needed is a "third way"—a model that unifies the exploratory power of a continuous stochastic process with the axiomatic rigor of a formal system. The motivation for this project is to develop such a paradigm, enabling the generation of music that is both creatively novel and structurally coherent by design, where a composer can explore the unique musical universe a set of rules engenders without fear of violation.

To bridge this gap, we propose a fundamental reframing of the generative task from predicting the next event to modeling the trajectory of a dynamic system within a constrained space. This leads to our central hypothesis: **We hypothesize that by modeling the act of composition as a constrained stochastic process—specifically, a Gauss-Markov process evolving within a high-dimensional state space defined by reflecting boundaries—we can establish a generative framework capable of producing musically novel material that axiomatically adheres to a significant class of well-defined symbolic constraints.** A core challenge of this research, and a key point of investigation, will be to determine the extent to which complex, relational musical grammars can be effectively translated into such geometric boundaries. In this framework, the generative process does not learn to approximate rules; it is physically prevented from breaking them by the geometry of the space in which it operates.

This hypothesis reconceptualizes the core components of music generation. The state of the system is a vector in a high-dimensional continuous space, where each dimension corresponds to a musical parameter like pitch, duration, or dynamic intensity. The creative engine is a Gauss-Markov process, which introduces controlled, continuous random fluctuations, causing the state vector to explore this musical space and providing the capacity for novelty. The crucial innovation lies in imposing reflecting boundaries on this state space as the mathematical embodiment of musical rules. For example, the harmonic rules of C major can be defined as boundaries that constrain the pitch dimensions. If the stochastic process attempts to move the state vector to a pitch outside this set, its velocity is "reflected" at the boundary, forcing it back into the valid musical space. This mechanism ensures rule compliance is an axiomatic property.

While this geometric approach is well-suited for absolute constraints, we recognize that representing complex relational rules (e.g., voice-leading prohibitions) is a non-trivial problem that forms a central research question for this project. Furthermore, this data-agnostic framework intentionally prioritizes structural integrity over stylistic mimicry. We posit that this provides a robust foundation upon which stylistic models could later be built, addressing the risk of generating music that is correct but aesthetically sterile. This approach transforms the problem from one of statistical inference to one of constrained dynamics, promising a powerful new paradigm for algorithmic composition.

2.3. Novelty and Originality

The originality of this project is rooted in a paradigm shift, moving certain tasks in algorithmic composition from the dominant model of sequence prediction towards the simulation of a constrained dynamic system. Current state-of-the-art methods, from Transformers to diffusion models, conceptualize music as a sequence of discrete events, with the primary objective of learning the conditional probability of the next event given the past [3, 4, 5]. Our approach proposes a complementary framework. We model the musical state—a collection of continuous parameters like pitch, duration, and dynamics—as a point evolving within a high-dimensional space. The generative process is thus one of exploration, guided by the intrinsic dynamics of a Gauss-Markov process. This reconceptualization is the project's core innovation, reframing specific compositional challenges from learning statistical patterns to defining and navigating the geometry of a musical possibility space.

This conceptual shift is enabled by a novel methodological transfer from computational neuroscience to the creative arts. The use of Gauss-Markov models with reflecting boundaries is a well-established tool for modeling systems that are both stochastic and subject to hard constraints, such as the firing rates of neurons or the movement of particles within a container [9]. However, its application to any artistic domain, and specifically to symbolic music generation, is entirely unprecedented. By adapting this robust mathematical framework, we are not merely proposing an incremental improvement to existing techniques but introducing a new class of generative models to the field of computational creativity. This cross-disciplinary leap provides a rigorous, mathematically-grounded foundation for addressing the long-standing challenge of integrating formal structure with stochastic generation, a problem often addressed with ad-hoc solutions [7].

Our approach presents a clear departure from current deep learning methodologies. Unlike Transformer-based systems, which are fundamentally dependent on the statistical distributions of large training datasets [3], our model is data-agnostic by design. It is engineered to generate structurally compliant material within entirely novel, user-defined rule systems for which no prior data exists. This offers composers a tool to explore new musical languages, free from the stylistic confines of existing corpora. Furthermore, our method differs critically from rule-guided diffusion models, where constraints are imposed as an external corrective force during generation [6]. In our framework, the reflecting boundaries are an intrinsic property of the generative space itself; rule violations are not corrected because they are mathematically impossible from the outset. This represents a move from extrinsic guidance to intrinsic compliance, resulting in a system that is both more robust and computationally more efficient.

Simultaneously, our model is designed to address key limitations of classical rule-based and stochastic methods. While traditional Markov chains can model probabilistic transitions, their reliance on a discrete state space restricts their expressive range, often leading to simplistic outputs [1, 2]. Our use of a continuous state space is a key point of originality, allowing for the generation of music with greater nuance and complexity. It can naturally produce microtonal inflections, subtle rhythmic variations, and continuous dynamic changes—features difficult to represent in discrete models. In contrast to constraint satisfaction programming (CSP), which excels at enforcing rules but often lacks generative fluidity [8], our framework integrates a sophisticated stochastic process at its core. The Gauss-Markov dynamics ensure that the generated output is not only structurally sound but also exhibits creative variance and non-determinism, thus unifying the rigor of formal systems with the expressivity of advanced generative models. This synthesis of a continuous state space, intrinsic constraint enforcement, and a powerful stochastic engine constitutes an original contribution to the field.

2.4. Adequacy to the Call and Thematic Priority

This proposal strongly aligns with the "Knowledge Generation Projects" modality, as it aims to generate fundamental knowledge at the confluence of mathematics, artificial intelligence, and music theory. Rather than incrementally improving existing techniques, the project proposes and validates a novel theoretical and practical paradigm for constrained creative generation. Its primary outcomes—a new mathematical framework, a distinct class of generative algorithms, and high-impact scientific publications—are in direct alignment with the call's objective to advance the frontiers of scientific knowledge.

The project's design directly addresses the key evaluation criteria of quality and feasibility. Its quality is rooted in a clear, testable hypothesis and an innovative methodology that addresses a significant gap in the state of the art, as detailed in Sections 2.2 and 2.3. Feasibility is ensured by the unique interdisciplinary expertise of the research team, which combines deep knowledge of stochastic processes (Dr. Sharma), computational systems ({Co-PI Name}), and advanced musical theory ({Co-PI Name}), ensuring all project facets are underpinned by world-class expertise [1, 2].

While not formally aligned with a specific thematic priority, this research contributes to the broader challenge of developing more transparent and controllable artificial intelligence. By creating AI tools that operate on explicit, human-interpretable rules, our approach offers a transparent alternative to prevalent "black box" models, fostering a more collaborative partnership between human creativity and machine intelligence with potential impacts on the creative industries [10]. Furthermore, the project's foundational nature provides an ideal training environment for a predoctoral researcher, who will acquire a unique skillset spanning advanced mathematics, machine learning, and computational creativity, thereby contributing to the call's objective of developing the next generation of scientific leaders.

Title: "3. Objectives, methodology and work plan"

3. Objectives, methodology and work plan

The overarching goal of this project is to develop and validate a new paradigm for algorithmic music composition by establishing a novel generative framework based on Gauss-Markov models with reflecting boundaries. This framework will enable the creation of musically coherent and novel material that axiomatically adheres to explicit, user-defined constraints. To achieve this ambitious goal, we have defined four specific, measurable, and interdependent objectives that structure the research plan.

Specific Objective 1 (SO1): To formalize and implement the Reflecting-Boundary Gauss-Markov Model (RB-GMM) for symbolic music generation. This foundational objective involves translating the core theoretical concept into a robust mathematical and computational reality. We will define a continuous, high-dimensional state space where each dimension represents a fundamental musical parameter (e.g., pitch, onset time, duration, velocity). We will then formalize the dynamics of the Gauss-Markov process within this space, specifying the drift and diffusion components that drive creative exploration. The central task will be to design and implement computationally efficient algorithms for enforcing reflecting boundaries, which represent the mathematical instantiation of musical rules. Success for this objective will be a functional software prototype capable of generating constrained state-space trajectories, verified against simple geometric domains.

Specific Objective 2 (SO2): To develop a comprehensive library of musical constraint modules and validate the model's capacity for rigorous rule enforcement. This objective focuses on applying the RB-GMM framework to concrete musical problems. We will develop a modular library of boundary condition functions corresponding to a diverse range of musical rules, from fundamental to complex. This will include: (i) harmonic constraints (e.g., diatonic scales, microtonal modes, atonal pitch-class sets), (ii) rhythmic constraints (e.g., metric grids, complex time signatures, isorhythmic patterns), and (iii) elementary formal constraints (e.g., canonic imitation, intervallic restrictions in counterpoint). Validation will be quantitative, measuring the model's ability to generate long-form compositions (e.g., >10,000 events) with violations being demonstrably confined to negligible numerical tolerances inherent in the simulation.

Specific Objective 3 (SO3): To conduct a rigorous comparative evaluation of the RB-GMM's musical output against state-of-the-art generative models. With the model validated for correctness, this objective will assess the aesthetic qualities and musical utility of its output. The goal is not simply to determine superiority, but to characterize the unique creative affordances of a constraint-driven paradigm versus data-driven approaches. We will conduct a multi-faceted evaluation, comparing compositions from the RB-GMM with those from leading Transformer and diffusion-based models under similar constraint conditions. The evaluation will comprise: (i) quantitative analysis using information-theoretic metrics (e.g., entropy, mutual information) to measure structural complexity and predictability, and (ii) a formal listening study involving a panel of expert composers, performers, and music theorists.

Specific Objective 4 (SO4): To package the developed framework into a documented, open-source software toolkit for the research and composition communities. The final objective is to ensure the project's outcomes are accessible, reproducible, and impactful beyond the immediate research team. We will produce a high-quality, open-source Python library encompassing the core RB-GMM engine, the library of musical constraint modules, and tools for visualization and output conversion to standard formats like MIDI and MusicXML. The toolkit will be accompanied by comprehensive documentation, tutorials, and example use cases to facilitate its adoption by other researchers and its integration into the creative workflows of contemporary composers.

The methodology is designed to systematically address each specific objective, building from theoretical formalization to practical application and evaluation.

For **SO1**, the methodology is grounded in the theory of stochastic differential equations (SDEs) and constrained random walks. The musical state at time `t`, denoted `X(t)`, will be a vector in `$\mathbb{R}^n$`, where `n` is the number of musical parameters being modeled. The evolution of `X(t)` will be described by a Gauss-Markov SDE of the form `$dX(t) = \mu(X, t)dt + \sigma(X, t)dW(t)$`, where `$\mu$` is the drift vector, `$\sigma$` is the diffusion matrix, and `$dW(t)$` is a Wiener process. Initially, `$\mu$` and `$\sigma$` will be set as constants to model a simple Brownian motion. The core of the methodology is the implementation of the reflecting boundaries. Musical constraints will be defined as a domain `$D \subset \mathbb{R}^n$`, a convex polytope defined by `$Ax \leq b$`. When the trajectory `$X(t)$` reaches the boundary `$\partial D$`, its velocity will be reflected. We will implement and compare several numerical schemes for simulating this process, starting with a projected Euler-Maruyama method, which provides a straightforward way to ensure the trajectory remains within `D` at each discrete time step [9]. We will also investigate more sophisticated methods, such as the reflection map technique, which may offer greater numerical stability [9]. The initial implementation will be in Python, leveraging NumPy and SciPy for efficient numerical computation.

The methodology for **SO2** involves translating abstract musical rules into precise geometric constraints within the state space `D`. For harmonic constraints, a diatonic scale for a single voice will be represented as a set of disjoint intervals on the pitch dimension. For polyphony, harmonic consonance rules (e.g., avoiding parallel fifths) will be defined as "forbidden zones" within the multi-dimensional pitch space. Rhythmic constraints will be implemented by defining boundaries on the onset-time and duration dimensions. For example, a 4/4 meter can be enforced by placing reflecting boundaries at each beat. Validation will be performed through extensive simulation. For each constraint module, we will run the RB-GMM to generate 100 compositions of 1,000 notes each. The output will be parsed and automatically checked for rule violations. The success criterion is to demonstrate axiomatic compliance, where any measured boundary violations are confined to a predefined numerical tolerance (e.g., $\varepsilon < 10^{-6}$) and are corrected within a single simulation step. This rigorous, quantitative approach validates the system's functional correctness.

For **SO3**, we will employ a mixed-methods evaluation protocol. First, we will establish a baseline for comparison by selecting two state-of-the-art models: a large Transformer-based model (e.g., MuseNet) and a rule-guided diffusion model [3, 6]. We will generate a corpus of 50 compositions from each model, as well as from our RB-GMM, under a common constraint (e.g., C major diatonic harmony). For quantitative analysis, we will compute metrics of musical complexity and structure, including pitch and rhythm entropy, long-range self-similarity measures, and statistical analysis of interval distributions [11]. This will allow us to objectively characterize the stylistic properties of each model's output. The second part is a formal, double-blind listening study. We will recruit a panel of at least 20 participants (composers, musicologists, performers). Participants will be presented with short, anonymized excerpts from each model and asked to rate them on a 7-point Likert scale across several dimensions designed for a balanced comparison: (1) Rule Adherence, (2) Stylistic Authenticity, (3) Long-Range Coherence, (4) Novelty/Creativity, and (5) Overall Aesthetic Preference. Statistical analysis (e.g., ANOVA) will be used to determine significant perceptual differences between the outputs.

The methodology for **SO4** follows best practices in open-source scientific software development. The entire codebase will be developed under version control using Git and hosted on a public repository (e.g., GitHub). The software will be structured as a Python package with a clear separation between the core SDE simulation engine and the library of musical constraint modules. We will use continuous integration (CI) to automate testing. Documentation will be a primary focus, generated using tools like Sphinx to create a professional web-based manual, including a "Getting Started" guide, API references, and tutorials. The final deliverable will be a stable, version 1.0 release of the toolkit, installable via the Python Package Index (PyPI).

The project is structured into five Work Packages (WPs) distributed over a 36-month period. A detailed Gantt chart is provided in {Annex_Gantt_Chart}.

WP1: Framework Development and Implementation (Months 1-12) This foundational WP, led by Dr. Sharma, focuses on realizing SO1. - **Task 1.1:** Mathematical Formalization (M1-3). Led by the PI, this task will define the state-space representation and the specific form of the Gauss-Markov SDEs. - **Task 1.2:** Algorithm Design and Implementation (M4-12). Following onboarding, the predoctoral researcher will implement the numerical solver for the constrained SDE, including the reflection mechanism. - **Task 1.3:** Core Engine Testing (M10-12). The predoctoral researcher will develop a suite of unit tests to verify the numerical accuracy and stability of the solver using simple, non-musical geometric domains. - **Deliverables:** D1.1: Technical report detailing the mathematical framework (M6). D1.2: Alpha version of the RB-GMM Python library (M12). - **Milestone:** M1: Functional prototype capable of generating a constrained random walk in `n`-dimensions.

WP2: Musical Constraint Library and Validation (Months 10-24) Led by {Co-PI Name, Music Theory Expert} with support from {Co-PI Name, Computational Systems Expert}, this WP addresses SO2. - **Task 2.1:** Design of Harmonic Constraint Modules. Develop boundary functions for scales, modes, and intervallic rules. - **Task 2.2:** Design of Rhythmic Constraint Modules. Develop boundary functions for meter, quantization grids, and rhythmic patterns. - **Task 2.3:** Systematic Validation. Run large-scale simulations to quantitatively verify the axiomatic compliance for all implemented constraints. - **Deliverables:** D2.1: Public release of the musical constraint library (M20). D2.2: Validation report with statistical results of the simulations (M24). - **Milestone:** M2: Successful validation of at least five harmonic and three rhythmic constraint modules.

WP3: Comparative Evaluation and Musical Analysis (Months 20-33) This WP, led by {Co-PI Name, Music Theory Expert}, is dedicated to SO3. - **Task 3.1:** Corpus Generation. Generate a dataset of compositions from the RB-GMM and selected baseline models. - **Task 3.2:** Quantitative Analysis. Apply information-theoretic and statistical measures to the generated corpus. - **Task 3.3:** Listening Study Design and Execution. Develop the protocol, obtain ethical approval, recruit participants, and conduct the experiment. - **Task 3.4:** Data Analysis and Interpretation. Analyze results from the quantitative metrics and the listening study. - **Deliverables:** D3.1: Publicly available dataset of generated musical corpora (M26). D3.2: Manuscript draft detailing the comparative evaluation results for journal submission (M33). - **Milestone:** M3: Completion of the listening study data collection.

WP4: Toolkit Finalization and Dissemination (Months 24-36) Led by {Co-PI Name, Computational Systems Expert}, this WP addresses SO4. - **Task 4.1:** Software Refinement and Packaging. Refactor code for public release, improve efficiency, and package for PyPI. - **Task 4.2:** Documentation and Tutorials. Write comprehensive user documentation, API references, and example-driven tutorials. - **Task 4.3:** Scientific Dissemination. Prepare and submit at least two articles to top-tier journals and present findings at major international conferences. - **Deliverables:** D4.1: Version 1.0 public release of the REBOUND open-source toolkit (M34). D4.2: Project website with documentation, tutorials, and audio examples (M36). - **Milestone:** M4: Acceptance of one major journal publication.

WP5: Project Management and Coordination (Months 1-36) Led by the PI, Dr. Sharma, this ongoing WP ensures the smooth execution of the project. - **Task 5.1:** Project Kick-off and Recruitment. Organize the kick-off meeting, advertise the predoctoral position, and complete the hiring process (M1-3). - **Task 5.2:** Coordination and Monitoring. Organize bi-weekly team meetings, monitor progress against milestones, and manage project resources. - **Task 5.3:** Reporting. Prepare and submit all required scientific and financial reports to the funding agency. - **Task 5.4:** Supervision of Predoctoral Researcher. Provide scientific guidance, training, and career development support for the hired researcher.

The project team has been assembled to provide the necessary interdisciplinary expertise. Responsibilities are clearly delineated and aligned with the work plan.

Dr. Anya Sharma (Principal Investigator): An expert in stochastic processes and computational modeling, Dr. Sharma will have overall responsibility for the project. She will lead WP1 and WP5, overseeing the mathematical formalization, guiding algorithm design, and ensuring the project meets its objectives. She will also lead the supervision of the predoctoral researcher.

{Co-PI Name, Computational Systems Expert}: With experience in scientific software development, {Co-PI Name} will lead WP4 and co-lead WP2. They will be responsible for the software architecture, implementation quality, and the final packaging and release of the open-source toolkit.

{Co-PI Name, Music Theory Expert}: A recognized authority in music theory and computational musicology, {Co-PI Name} will lead WP3 and co-lead WP2. Their role is to ensure the musical relevance of the project, guide the translation of musical rules into formal constraints, and lead the analysis of the musical outputs.

Predoctoral Researcher (to be hired): The predoctoral researcher will be the primary technical contributor, involved in WPs 1, 2, 3, and 4. They will be responsible for the day-to-day implementation of algorithms, running experiments, and contributing to publications and software development.

We have identified three potential critical points and have developed corresponding contingency plans.

Risk 1: Computational Scalability. The numerical simulation of a constrained SDE in a high-dimensional space can be computationally intensive. If the initial Euler-Maruyama scheme proves too slow, our contingency plan is two-fold. First, we will implement more advanced numerical methods, such as higher-order stochastic Runge-Kutta schemes adapted for reflection [12]. Second, if performance remains a bottleneck, we will leverage the university's High-Performance Computing (HPC) resources for offline generation, shifting the focus from real-time interaction to a non-real-time composition tool.

Risk 2: Aesthetic Sterility. The generated music, while rule-adherent, may lack aesthetic interest. Our contingency plan is to enhance the generative dynamics by modifying the SDE to include a non-trivial drift term `μ(X, t)`, which can attract the trajectory towards certain regions of the state space. We can also make the diffusion matrix `σ(X, t)` state-dependent. This would transform the model into a more complex Ornstein-Uhlenbeck-type process, providing greater control over the output's character without sacrificing axiomatic rule enforcement. This enhancement would also provide a foundation for future work exploring how such dynamic parameters could be learned from small musical corpora, addressing the model's current data-agnostic nature.

Risk 3: Difficulty in Formalizing Complex Constraints. While simple rules are readily translatable into geometric boundaries, more complex musical concepts are known to be challenging to express purely as convex polytopes. Our contingency plan is to adopt a hybrid, multi-stage approach. For such rules, we will use the RB-GMM to handle "local" constraints (e.g., harmony, rhythm). The output would then be passed to a second-stage process, such as a constraint satisfaction solver or a filtering algorithm, to enforce more complex, "global" rules. This pragmatic approach ensures the project can deliver valuable results even if the most complex constraints prove intractable for a purely geometric representation.

The research team has access to all necessary resources, including the {Name_of_HPC_Cluster} High-Performance Computing cluster, a dedicated laboratory space, and comprehensive library access.

Our capacity is demonstrated by our strong track record. Dr. Sharma's group has developed a proof-of-concept prototype of the RB-GMM mechanism. A preliminary study, under review for ISMIR, successfully modeled a single melodic line constrained to a pentatonic scale, confirming the projected Euler-Maruyama algorithm is both feasible and effective at enforcing simple boundaries. We also have initial unpublished results from a 2D model applied to two-voice counterpoint, demonstrating the model's ability to generate textures that axiomatically avoid parallel fifths. This provides strong evidence that the reflecting boundary concept can be effectively applied to relational musical rules. The collective publication record of the PIs includes numerous high-impact articles in stochastic modeling, computational creativity, and music theory. A selection of relevant publications is provided below: {Insert list of top 5 publications}. These preliminary findings and our established expertise significantly de-risk the proposed project.

3.1. General and Specific Objectives

The overarching goal of this project is to propose and evaluate a new paradigm for algorithmic music composition that integrates stochastic processes with axiomatic rule-adherence. We will develop a novel generative framework based on Gauss-Markov models with reflecting boundaries, capable of producing musically coherent and aesthetically compelling material that rigorously complies with explicit, user-defined constraints. This goal will be achieved through four specific, measurable, and logically sequenced objectives.

Specific Objective 1 (SO1): To formalize the mathematical framework and implement the core computational engine of the Reflecting-Boundary Gauss-Markov Model (RB-GMM) for symbolic music generation. This foundational objective involves translating the project's central hypothesis into a functional software prototype. We will formalize a continuous state-space representation for core musical parameters and implement a robust numerical solver for the constrained stochastic differential equations governing the system's dynamics. Success will be defined by the creation of a verified computational engine capable of generating a continuous state-space trajectory that remains strictly within a set of predefined geometric constraints.

Specific Objective 2 (SO2): To develop and validate a modular library of musical constraint functions for the RB-GMM. This objective translates the abstract mathematical framework into practical musical application. We will design and implement a library of boundary functions corresponding to a diverse range of musical rules, including harmonic constraints (e.g., diatonic scales, pitch-class sets), rhythmic constraints (e.g., metric grids, isorhythmic patterns), and elementary contrapuntal rules. Success will be measured by demonstrating a 0% violation rate for each implemented constraint across large-scale simulations of long-form compositions.

Specific Objective 3 (SO3): To conduct a rigorous comparative evaluation of the RB-GMM's musical output against state-of-the-art generative models. With the model's correctness established, this objective will assess its output's aesthetic quality, novelty, and structural complexity. We will perform a mixed-methods evaluation, comparing compositions from the RB-GMM with those from leading Transformer and diffusion-based models. This evaluation will critically assess whether the model can generate aesthetically compelling material, moving beyond mere constraint satisfaction to achieve stylistic nuance. It will involve quantitative analysis using information-theoretic metrics and a formal, double-blind listening study with a panel of expert composers and music theorists.

Specific Objective 4 (SO4): To disseminate the project's findings by packaging the developed framework into a documented, open-source software toolkit. To maximize the project's impact, our final objective is to produce a high-quality, accessible, and extensible software library. This toolkit will encompass the core RB-GMM engine, the library of musical constraint modules, and utilities for visualization and data conversion. It will be accompanied by comprehensive documentation, tutorials, and example use cases to facilitate adoption by other researchers and integration into the creative workflows of contemporary composers.

3.2. Methodology

This project's methodology systematically addresses the four specific objectives. It progresses from foundational modeling (SO1) and constraint library development (SO2), through empirical evaluation (SO3), to the dissemination of an open-source toolkit (SO4). This staged approach ensures each component builds upon a validated predecessor, guaranteeing both scientific rigor and practical utility.

To achieve **Specific Objective 1**, the formalization and implementation of the Reflecting-Boundary Gauss-Markov Model (RB-GMM), we will employ a methodology grounded in the numerical simulation of stochastic differential equations (SDEs). The musical state at any time `t` will be represented by a vector `$X(t) \in \mathbb{R}^n$`, where `n` is the dimensionality of the musical space. The evolution of this state vector will be governed by a Gauss-Markov SDE of the form `$dX(t) = \mu(X, t)dt + \sigma(X, t)dW(t)$`, where `μ` is the drift vector, `σ` is the diffusion matrix, and `dW(t)` represents a standard Wiener process [9]. As a baseline, we will initially set the drift `μ` to zero and the diffusion `σ` to a constant diagonal matrix (`$\sigma = cI$`), which models an isotropic Brownian motion. This provides a simple, unbiased engine for stochastic exploration, allowing us to isolate the effects of the boundary constraints. A key extension will be investigating methods to infer a non-zero drift `μ(X, t)` from musical corpora, guiding the process towards specific stylistic tendencies and adding aesthetic nuance.

The core of the methodology for SO1 is the numerical implementation of the reflecting boundaries that define the valid musical domain, `D`. While simple constraints can define `D` as a single convex polytope (`$Ax \leq b$`), we recognize that musical grammar typically defines `D` as a non-convex set (e.g., a union of convex regions), a complexity our framework is designed to accommodate. The constrained SDE will be simulated using a projected Euler-Maruyama scheme [9]. At each time step `Δt`, a standard Euler-Maruyama step proposes a new state `Y_{k+1}`. If `Y_{k+1}` is outside `D`, it is projected back to the nearest point on the domain's surface to enforce the boundary condition. This projection is a standard quadratic programming problem that can be solved efficiently [13]. The simulation engine will be implemented in Python using NumPy and SciPy for efficient computation. Verification of the core engine will be conducted by simulating the process within simple geometric domains (e.g., a 1D line segment, a 2D disk) and comparing the resulting statistical distributions against known analytical solutions to confirm numerical accuracy.

For **Specific Objective 2**, the development of a musical constraint library, our methodology focuses on translating principles of music theory into the geometric language of the state space. This interdisciplinary process will be systematic and modular. Harmonic constraints will be implemented first. For a single voice, a diatonic scale will be defined as a union of disjoint valid intervals along the pitch dimension, forming a non-convex domain. When a potential state falls outside this domain, it is projected to the closest point within the valid set (`argmin_{x ∈ D} ||y - x||`). We will implement an efficient solver for this projection and investigate reflection schemes that minimize audible artifacts. For polyphony, relational constraints like the prohibition of parallel fifths will be modeled as 'forbidden zones'. We will incorporate state-dependent boundary conditions, where active constraints are updated dynamically based on the trajectory's recent history, to handle such context-dependent rules. Rhythmic constraints will be implemented by defining boundaries on the onset and duration dimensions. A 4/4 meter, for example, can be enforced by placing reflecting boundaries at integer beat locations along the onset-time axis.

Validation for each constraint module will be quantitative. For each implemented rule, we will execute a large-scale simulation to generate a corpus of 100 compositions, each containing at least 10,000 musical events. We will develop automated parsing scripts to check for any violation of the specified rule. The success criterion is a measured violation rate of 0% for every constraint, empirically verifying the system's axiomatic compliance.

The methodology for **Specific Objective 3**, the comparative evaluation, will employ a mixed-methods approach to assess the musical quality of the RB-GMM's output against the state of the art. We will select two representative baseline models: a large-scale Transformer model (e.g., MuseNet) and a state-of-the-art rule-guided diffusion model [3, 6]. To ensure a fair comparison, we will generate a corpus of 50 compositions (each approximately 60 seconds in length) from our RB-GMM and from each baseline model, all tasked with adhering to an identical set of simple constraints.

The first phase of the evaluation will be a quantitative analysis using established music information retrieval (MIR) metrics. We will measure structural complexity and predictability using Shannon entropy of pitch and rhythm distributions, as well as Lempel-Ziv complexity [11]. Long-range coherence will be assessed using detrended fluctuation analysis (DFA) to measure statistical self-affinity over multiple time scales [11]. We will also analyze stylistic features by comparing the distributions of melodic intervals and harmonic n-grams.

The second phase will be a formal, double-blind listening study involving a panel of at least 20 expert participants (composers, music theorists, performers). To control for confounding variables, all excerpts will be rendered with an identical synthesizer for timbral consistency and loudness-normalized to a common target (e.g., -14 LUFS). Participants will evaluate these short (20-30 second), anonymized, and randomized excerpts. They will rate each on a 7-point Likert scale across five criteria: (1) Harmonic Coherence, (2) Rhythmic Integrity, (3) Structural Plausibility, (4) Novelty and Creativity, and (5) Overall Aesthetic Quality. Data will be analyzed with ANOVA and post-hoc tests (e.g., Tukey's HSD) to identify statistically significant differences in perceived quality between models.

Finally, the methodology for **Specific Objective 4**, the creation of an open-source toolkit, will adhere to best practices in modern scientific software development. The codebase will be version-controlled with Git on a public GitHub repository from the project's start. We will employ a modular architecture, separating the core SDE engine from the musical constraint library to facilitate extension. Continuous integration (CI) via GitHub Actions will be used to automate a comprehensive suite of unit and integration tests. We will use the Sphinx documentation generator to create a professional, web-based manual that includes an installation guide, a conceptual overview, a detailed API reference, and practical tutorials. The toolkit will include modules for converting the output to standard formats like MIDI and MusicXML, enabling integration with DAWs and notation software. The final toolkit will be distributed via PyPI (`pip install rebound`) to maximize accessibility for researchers and artists.

3.3. Work Plan and Schedule

The project is structured into five interconnected Work Packages (WPs) over a 36-month timeline, designed to ensure a logical progression from foundational research to public dissemination. This work plan is carefully scheduled to manage dependencies and mitigate risks, with a detailed visual representation provided in the Gantt chart in {Annex_Gantt_Chart}. The successful completion of each WP is marked by specific deliverables and milestones, ensuring continuous progress and accountability.

WP1: Framework Development and Implementation (Months 1-12). This foundational WP, addressing SO1, is dedicated to translating the theoretical RB-GMM concept into a robust computational engine. The initial months will focus on the rigorous mathematical formalization of the state-space representation and the constrained SDEs (Task 1.1). This theoretical work will run in parallel with core software development, where we will design and implement the numerical solver, including the critical reflection mechanism based on the projected Euler-Maruyama scheme (Task 1.2). The first year culminates in the development of a comprehensive suite of unit tests to verify the solver's numerical accuracy and stability against known analytical solutions in simple geometric domains (Task 1.3). The primary outputs are a detailed technical report on the mathematical framework (D1.1, M6) and a functional alpha version of the core RB-GMM Python library (D1.2, M12). Milestone 1 (M1) is the successful demonstration of a functional prototype capable of generating a constrained random walk in an n-dimensional space.

WP2: Musical Constraint Library and Validation (Months 10-24). Building on the core engine from WP1, this WP addresses SO2 by developing a methodology for encoding musically meaningful rules as geometric boundaries. With a slight overlap with WP1 to ensure a smooth transition, the primary activities involve modeling and implementing modular boundary functions for a range of harmonic constraints, including diatonic scales, microtonal modes, and intervallic prohibitions (Task 2.1), followed by the development of modules for rhythmic constraints such as metric grids and complex time signatures (Task 2.2). A crucial part of this WP is the systematic validation of each module (Task 2.3), which involves running large-scale simulations to empirically verify the system's axiomatic adherence to the encoded rules. This phase will deliver a public release of the validated musical constraint library (D2.1, M20) and a comprehensive validation report detailing the statistical results of the simulations, confirming a 0% violation rate for the implemented constraints (D2.2, M24). Milestone 2 (M2) will be achieved upon the successful validation of at least five distinct harmonic and three rhythmic constraint modules, demonstrating the framework's versatility.

WP3: Comparative Evaluation and Musical Analysis (Months 20-33). This WP, dedicated to SO3, focuses on rigorously assessing the musical quality and novelty of the RB-GMM's output. The work commences with generating a large corpus of compositions from our system and selected state-of-the-art baseline models under identical constraint conditions (Task 3.1).

This corpus will then be subjected to a two-pronged analysis. First, a quantitative analysis using information-theoretic and statistical measures will be performed to objectively characterize the music's structural properties (Task 3.2). Concurrently, we will design and execute a formal, double-blind listening study with expert participants (Task 3.3). This study, following ethical approval, will evaluate not only preference but also key aesthetic dimensions such as stylistic coherence, novelty, and expressiveness, directly addressing the risk of generating structurally correct but musically inert output. The final stage involves the statistical analysis and interpretation of the collected quantitative and qualitative data to assess both structural integrity and aesthetic merit (Task 3.4). Key deliverables include the publicly available dataset of generated musical corpora (D3.1, M26) and a manuscript draft detailing the full comparative evaluation for submission to a top-tier journal (D3.2, M33). The successful completion of the data collection for the listening study will constitute Milestone 3 (M3).

WP4: Toolkit Finalization and Dissemination (Months 24-36). Addressing SO4, this WP ensures the project's outcomes are made widely accessible for a lasting impact. Activities run in parallel with the final stages of evaluation and focus on transforming the prototype software into a professional, distributable toolkit. This involves significant code refactoring, performance optimization, and packaging for the Python Package Index (PyPI) (Task 4.1). A major effort will be dedicated to writing comprehensive documentation, including API references and user-friendly tutorials (Task 4.2). The scientific dissemination strategy (Task 4.3) will be executed during this period, involving the preparation and submission of at least two major journal articles and presentations at leading international conferences. The main deliverables are the official version 1.0 public release of the REBOUND open-source toolkit (D4.1, M34) and the launch of a project website featuring documentation, tutorials, and audio examples (D4.2, M36). Milestone 4 (M4) is the acceptance of our first major journal publication, signifying peer-validated success.

WP5: Project Management and Coordination (Months 1-36). This continuous WP underpins the entire project, ensuring its smooth and efficient execution. The PI will oversee all activities, including organizing bi-weekly team meetings to monitor progress against milestones, managing project resources, and adapting the plan as needed (Task 5.1). This WP also covers all administrative duties, such as the preparation and submission of scientific and financial reports to the funding agency (Task 5.2). A central component is the dedicated scientific supervision, training, and career development support for the hired predoctoral researcher, ensuring they are fully integrated into the team and benefit from a world-class research environment.

3.4. Team Member Responsibilities

The success of this project is predicated on the unique, interdependent expertise of its three principal investigators, representing mathematics, computer science, and music theory. The team is structured such that the removal of any member would render the project unfeasible. Responsibilities are mapped directly to the project's work packages, ensuring each member's distinct contribution is leveraged to achieve the project's objectives.

Dr. Anya Sharma (Principal Investigator) will provide overall scientific leadership and project management. As an expert in stochastic processes, she will personally lead WP1, guiding the mathematical formalization of the RB-GMM and overseeing the design and validation of the core simulation algorithms. Her role is to ensure the theoretical and methodological rigor of the project's foundation. She will also lead WP5, managing all administrative, financial, and reporting duties, and will serve as the primary supervisor for the predoctoral researcher, fostering their scientific development.

Dr. Ben Carter (Co-PI), an expert in high-performance scientific computing, will serve as the project's lead technologist. His primary responsibility is translating the mathematical framework into efficient and reproducible code, a role for which his track record in developing GPU-accelerated libraries is indispensable. He will co-lead WP2, architecting the computational implementation of the musical constraint modules, and will lead WP4, overseeing the packaging and dissemination of the open-source toolkit. Dr. Carter's specific expertise is critical for optimizing the large-scale simulations in WP2 and WP3 and ensuring the project's final output is a high-quality, accessible tool.

Dr. Elena Vance (Co-PI), a leading scholar in computational musicology, is the project's lead domain expert. Her role is to ensure the musical validity and relevance of the project's outcomes. Her seminal research on formalizing post-tonal grammar as geometric systems provides the conceptual foundation for WP2, which she will co-lead. She will lead WP3, where her expertise is essential for designing a musically meaningful comparative evaluation, from selecting appropriate quantitative metrics to developing the expert listening study methodology. She will lead the analysis of the generated music, contextualizing the findings within contemporary music theory.

The **Predoctoral Researcher (to be hired)** will be the central technical contributor, responsible for the day-to-day execution of the research plan under the close supervision of the PIs. This individual will be actively involved across WPs 1, 2, 3, and 4. Their tasks will include implementing the core algorithms, running the extensive validation and evaluation experiments, contributing to the development and documentation of the software toolkit, and co-authoring scientific publications. This role is designed as a comprehensive training experience at the intersection of these disciplines.

3.5. Critical Points and Contingency Plan

We have identified three potential risks that are inherent to the ambitious and innovative nature of this project. For each, we have developed a clear and feasible contingency plan to ensure the project's objectives are met.

The first critical point is **computational scalability**. The numerical simulation of a constrained SDE in a high-dimensional space can be computationally intensive. If the proposed projected Euler-Maruyama scheme proves too slow for generating complex music in a practical timeframe, our contingency plan is twofold. First, we will implement more sophisticated numerical methods, such as higher-order stochastic Runge-Kutta schemes adapted for reflection, which may permit larger time steps without sacrificing stability or accuracy [12]. Second, should performance remain a bottleneck for interactive applications, we will leverage the university's High-Performance Computing (HPC) resources. This would pivot the tool's primary use case from real-time generation to a powerful engine for non-real-time composition, which still represents a significant advance and fully addresses the project's core scientific goals.

The second potential risk is that of **aesthetic inertia**. The generated music, while formally correct, may lack musical interest due to two factors: overly simple dynamics leading to sterility, or the process becoming trapped in simple, repetitive loops near constraint boundaries—a known issue for such systems. Our mitigation strategy is to enhance the generative dynamics in a musically transparent manner. We will introduce interpretable, musically-informed potential fields to act as a gentle drift term, `$\mu(X, t)$`, guiding the trajectory towards regions of interest (e.g., consonant intervals) to create tension and release. To specifically counter repetitive looping, we can modulate the diffusion tensor near boundaries to encourage escape trajectories. This approach enriches the output's character while preserving the model's interpretability, avoiding the "black box" nature of alternative methods.

The final critical point concerns the **formalization of multi-scale constraints**. While local rules like harmony are readily translatable into geometric boundaries, global structures such as sonata form involve relational and temporal dependencies that are more effectively modeled at a different level of abstraction. Our methodology addresses this by design, employing a hierarchical, two-stage architecture. The RB-GMM excels at its unique purpose: generating musically rich material that rigorously satisfies all local, geometric constraints. This output then serves as the input for a second-stage process, such as a dedicated constraint satisfaction solver or a lightweight planning algorithm, which sculpts the material according to more abstract, global rules. This modular strategy is not a fallback but a core strength, creating a synergistic system that leverages the distinct advantages of both continuous generation and symbolic reasoning to address complex compositional problems.

3.6. Available Resources and Previous Results

The research team is well-positioned to execute this project, supported by the robust infrastructure of {Institution_Name} and a strong foundation of preliminary results that validate the proposal's core hypotheses. The project will be hosted at {Institution_Name}, which offers a world-class research environment and guarantees access to all necessary computational resources. For the large-scale simulations in WP2 and WP3, the team has secured a significant allocation on the {Name_of_HPC_Cluster}, a Tier-2 High-Performance Computing facility. This cluster comprises {Number} compute nodes, each with {CPU_details}, and includes a partition of {Number} nodes with NVIDIA {Model_GPU} GPUs, providing ample power for demanding stochastic simulations. Day-to-day development will occur in the PI's research lab, equipped with high-performance workstations ({Workstation_specs}), professional audio monitoring systems, and licensed software including MATLAB, Max/MSP, and the Adobe Creative Suite. Furthermore, the university's library provides unrestricted access to major scientific databases (IEEE Xplore, ACM Digital Library, JSTOR), ensuring the team remains current with relevant literature.

Our capacity to undertake this project is substantiated by previous results that establish a proof-of-concept for the proposed RB-GMM framework. The PI's group has developed and tested a prototype of the core reflection mechanism, and our initial findings de-risk the most innovative technical aspects of this proposal. In a preliminary study, we successfully modeled a single melodic line as a one-dimensional constrained random walk [9]. The state space represented pitch, and we implemented reflecting boundaries to confine the melody to a C major pentatonic scale. The simulation, based on the projected Euler-Maruyama scheme, generated over one million note events. Subsequent analysis confirmed the algorithm was 100% effective in enforcing these harmonic constraints, with zero violations recorded. This foundational work, which demonstrates the technical feasibility and mathematical precision of the reflection mechanism, is currently under review at the International Society for Music Information Retrieval (ISMIR) conference.

Building on this success, we have obtained compelling unpublished results from a more complex two-dimensional model applied to two-voice counterpoint. In this experiment, the state space was `$\mathbb{R}^2$`, with each dimension representing the continuous pitch of a voice. We defined the boundary conditions as a geometric analogue for a foundational constraint in Western counterpoint: the avoidance of parallel perfect fifths. This was implemented by defining a "forbidden zone" in the state space corresponding to an intervallic distance of seven semitones. The RB-GMM simulation generated two-voice textures that fluidly explored the pitch space while axiomatically avoiding the forbidden interval, effectively "bouncing" off the geometric representation of the contrapuntal constraint. This result is significant as it provides initial evidence that the reflecting boundary paradigm can extend beyond simple range limitations to model complex, relational musical rules. It suggests that core compositional principles can be productively represented as geometric constraints in a continuous state space, validating the core technical approach of WP1 and providing an empirical basis for the constraint library in WP2.

The PIs' collective expertise and publication record further underscore our readiness. Dr. Sharma is a leading expert in applying stochastic processes to complex systems, with numerous publications in journals such as {Journal_of_Stochastic_Modeling} and {Journal_of_Computational_Physics} [12]. {Co-PI Name, Computational Systems Expert} has a proven track record in developing open-source scientific software, contributing to projects like {Name_of_Software_Project} [14]. {Co-PI Name, Music Theory Expert} is an internationally recognized scholar in computational musicology, with seminal works on the formal analysis of musical structure [1, 2]. A selection of our most relevant publications is provided here: {Insert list of top 5 publications}. This combination of state-of-the-art resources, compelling preliminary data, and a well-matched interdisciplinary team provides a solid foundation for the successful execution of this project and the achievement of its objectives.

4. Expected impact of the results

The successful completion of this project will generate significant and multi-faceted impacts, advancing the scientific frontier of computational creativity, providing powerful new tools for the creative industries, and establishing a new paradigm for the interaction between human artistry and machine intelligence. The outcomes are designed not merely to produce new knowledge, but to make that knowledge accessible, applicable, and transformative for a diverse range of stakeholders, from academic researchers to professional composers.

The primary scientific and technical impact of this research will be the establishment of a new paradigm for constrained stochastic generation. By developing and validating the Reflecting-Boundary Gauss-Markov Model (RB-GMM), this project will address a long-standing methodological challenge, offering a novel synthesis of the generative richness of data-driven models and the formal rigor of rule-based systems [7]. This constitutes a fundamental contribution to the field of algorithmic composition, moving beyond sequence prediction to a model of creative exploration within a mathematically defined space of possibility. The REBOUND framework will be the first of its kind to treat musical constraints not as external correctives or learned approximations, but as intrinsic properties of the generative process itself. This will unlock new avenues of research into musical structure, enabling researchers to use the model as a "computational laboratory" for music theory. For instance, musicologists could formalize the rules of a historical practice (e.g., Palestrina-style counterpoint) and use the RB-GMM to generate a vast corpus of novel, rule-compliant examples, allowing for an unprecedented depth of stylistic analysis.

Beyond music, the project's impact will extend to the broader field of artificial intelligence. The problem of generating novel artifacts that adhere to a complex set of symbolic constraints is a general one, with applications in areas such as procedural content generation, architectural design, and drug discovery [8]. The RB-GMM will serve as a powerful proof-of-concept for a new class of controllable generative models that are data-agnostic, computationally efficient, and whose adherence to constraints is mathematically verifiable. This offers a transparent and auditable alternative to the "black box" nature of many large-scale neural models. The cross-disciplinary transfer of a mature mathematical framework from computational neuroscience to a creative domain is in itself a significant technical contribution, demonstrating a novel pathway for innovation that bridges disparate scientific fields. The open-source toolkit developed in SO4 will be a critical technical outcome, providing a robust and extensible platform that will catalyze further research and development in this new domain.

The social and economic impacts of this project are centered on empowering human creativity and enhancing the tools available to the creative industries. The REBOUND toolkit is envisioned not as a replacement for human composers but as a sophisticated co-creative partner. For composers in film, television, and video games, who often produce music under tight deadlines and specific constraints, our system will be a powerful tool for rapidly generating and exploring structurally coherent musical frameworks, which can then be artistically developed. This can significantly accelerate the creative workflow and enable a higher degree of artistic exploration. The model's ability to work with novel, user-defined rule systems will also be a catalyst for artistic innovation, providing contemporary composers with a new medium for exploring uncharted musical territories, from complex microtonal harmonies to novel rhythmic systems.

Economically, the project will generate value by creating new intellectual property in the form of novel algorithms and software. The open-source toolkit has the potential to be integrated into commercial Digital Audio Workstations (DAWs), creating new market opportunities. By positioning our institution at the forefront of this emerging field, the project will attract further investment and talent. Socially, the project will impact education. The REBOUND software can be used as an interactive tool for teaching music theory and composition, allowing students to visualize the effects of musical rules. This hands-on approach can make complex concepts more intuitive. Furthermore, by separating the logic of musical structure from specific stylistic data, the project contributes to a more inclusive musical landscape, empowering users to create within their own cultural or invented musical frameworks without being tethered to potentially biased historical corpora.

Our scientific communication and dissemination plan is designed to maximize the project's reach across multiple audiences. We will target high-impact, peer-reviewed publications in journals such as the *Computer Music Journal*, *IEEE Transactions on Affective Computing*, and the *Journal of New Music Research*. Findings will also be presented at leading international conferences, including the International Computer Music Conference (ICMC) and the International Society for Music Information Retrieval Conference (ISMIR). The cornerstone of our dissemination strategy is the open-source REBOUND toolkit (SO4). We will launch a dedicated project website to host the software, extensive documentation, tutorials, and a gallery of examples. To engage the artistic and developer communities, we will conduct hands-on workshops, actively promote the project through social media, and pursue public outreach through university press releases and lectures to share our vision of a collaborative human-AI creative partnership.

The project's internationalization strategy is designed to embed our research within the global scientific community. A key component is active participation in the premier international conferences in our field, which are critical for networking, receiving feedback, and identifying potential collaborators. We will proactively seek to establish formal collaborations with world-leading research centers such as {IRCAM in Paris, France}, the {Centre for Interdisciplinary Research in Music Media and Technology (CIRMMT) in Montreal, Canada}, and the {MIT Media Lab in Cambridge, USA}. These collaborations will initially take the form of joint workshops, with the long-term goal of developing joint funding proposals. A crucial element of our strategy is to foster the international mobility of our team. The predoctoral researcher hired for this project will be strongly encouraged and financially supported to undertake a research stay of 3-6 months at one of these partner institutions, providing an invaluable opportunity for knowledge exchange and the establishment of lasting professional networks.

This project is fully committed to the principles of Open Science, ensuring that all outputs are findable, accessible, interoperable, and reusable (FAIR). Our forecast for open access publications includes a minimum of three peer-reviewed journal articles and two conference papers. We will ensure all publications are made open access immediately upon publication, either by publishing in Gold Open Access journals or by depositing the final peer-reviewed manuscript in our institutional repository (Green Open Access) under a CC-BY license, in full compliance with funder mandates.

A summary of our Data Management Plan (DMP) is as follows. The project will generate several types of data: (1) the source code for the REBOUND toolkit; (2) generated musical corpora in MIDI, MusicXML, and raw state-vector formats; and (3) anonymized data from the expert listening study. All data will be managed in accordance with FAIR principles. The source code will be developed openly on GitHub under a permissive MIT license. Upon project completion, a final version of the code and the musical corpora will be permanently archived in a public repository such as Zenodo and assigned a Digital Object Identifier (DOI). The musical corpora will be licensed under a Creative Commons Attribution (CC-BY 4.0) license. Data from the listening study will be made available in a structured format upon publication of the corresponding analysis. To protect confidentiality, all identifying information will be removed. Throughout the project, data will be stored on the university's secure, backed-up network drive, with clear version control and metadata standards.

The gender dimension has been carefully considered in the design of this project's methodology and its potential impact. While the core mathematical framework of the RB-GMM is abstract, the context in which this technology is developed and applied is profoundly human and subject to bias. We will address the gender dimension in three key areas. First, in the execution of the comparative evaluation (SO3), we are committed to ensuring gender balance in the panel of expert listeners. By actively recruiting a diverse panel, we will mitigate potential gender-based biases in the evaluation of the generated music's quality and novelty.

Second, we recognize that the field of music composition, particularly the historical canon, has been male-dominated [15]. A key feature of our data-agnostic approach is that it is not trained to reproduce the statistical patterns of existing, potentially biased datasets. This intentional focus on abstract structure over stylistic mimicry allows the REBOUND toolkit to be used to explore compositional systems developed by female composers or to create entirely new systems free from historical gendered conventions. In our dissemination activities, we will actively highlight this capability. Third, the project team is committed to promoting gender equality in all its activities. We will ensure fair and unbiased practices in the hiring of the predoctoral researcher and will strive for gender balance among speakers and participants in any workshops or events we organize. By being conscious of these issues, we aim to ensure that our research contributes positively to a more equitable and inclusive scientific and artistic community.

4.1. Scientific-Technical Impact

The primary scientific and technical impact of this project will be the introduction and validation of a novel class of generative models designed to address the long-standing tension between stochastic creativity and formal structure in computational creativity [7]. By developing the Reflecting-Boundary Gauss-Markov Model (RB-GMM), we aim to establish a new paradigm for algorithmic composition that moves beyond the inherent limitations of purely data-driven and classical rule-based systems. This contribution will advance the scientific understanding of rule-based creativity, provide a powerful new toolkit for researchers and artists, and open significant avenues for future research across multiple disciplines.

This project will advance the scientific understanding of rule-based creative generation. Current models typically learn statistical approximations of rules from data or apply rules as external correctives [3, 6]. Our framework proposes a novel conceptualization of composition as the exploration of a mathematically-defined possibility space, rather than as a sequence prediction task. This provides a new lens through which to study the nature of musical structure. The REBOUND framework will function as a "computational laboratory" for music theory, a tool previously unavailable. Musicologists will be able to model specific, formalizable constraints of a historical style (e.g., Renaissance polyphony) or a contemporary compositional system (e.g., serialism) and use the RB-GMM to generate a vast corpus of novel, yet perfectly compliant, musical examples. This will enable an unprecedented depth of stylistic analysis, allowing researchers to investigate the emergent properties of a given rule set and test hypotheses about the relationship between constraints and perceived musical coherence.

Technically, the project's most significant outcome will be a novel and extensible method for constrained generation, embodied in the open-source REBOUND toolkit. This toolkit will represent a new category of generative tool with several key advantages. First, its data-agnostic architecture enables the exploration of musical systems independent of training corpora, prioritizing structural integrity over learned stylistic convention. This allows for the generation of compositions in novel systems for which no data exists, a capability distinct from data-dependent models. Second, its mechanism of intrinsic compliance through reflecting boundaries guarantees 100% rule adherence by design, offering a level of rigor unattainable with current "soft" guidance techniques [6]. Third, the continuous state-space representation allows for generating highly nuanced musical material, including microtonal inflections and complex rhythmic subtleties often lost in discrete models. This method will enable compositional and research workflows that are currently impossible, providing a robust platform for artistic creation and scientific inquiry.

The proposed framework is extensible and promises to open numerous avenues for future research. Within computational musicology, the RB-GMM can be extended to incorporate more sophisticated dynamics, such as state-dependent drift and diffusion terms, to model concepts like musical tension and release. This could open a research direction into learning the "dynamics of style" rather than only its statistical patterns. The framework also invites investigation into more complex boundary geometries, including dynamic boundaries that change over time to model evolving musical forms. More broadly, this project can impact artificial intelligence by providing a proof-of-concept for a class of controllable and transparent generative models. The challenge of generating novel artifacts that adhere to complex symbolic constraints is ubiquitous, with applications in procedural content generation, constrained optimization, architectural design, and robotic motion planning [8]. The RB-GMM could serve as a foundational method for adaptation to these domains, catalyzing innovation and providing a transparent alternative to opaque, data-hungry deep learning models.

4.2. Social and Economic Impact

The social and economic impacts of this project center on empowering human creativity by enhancing the tools available to the creative industries. The REBOUND toolkit is envisioned not as a replacement for human composers but as a sophisticated co-creative partner—a new kind of instrument that augments artistic intuition with powerful computational exploration. By providing a tool that allows users to systematically explore the consequences of their own defined musical structures, we aim to stimulate economic activity in the music technology sector, foster artistic innovation, and create new paradigms for music education, thereby contributing to a richer and more diverse cultural landscape.

Economically, the project will generate significant value for the creative industries, particularly in sectors where music is produced under tight deadlines and specific constraints, such as film, television, and video game scoring. Composers in these fields are often required to generate large volumes of stylistically coherent music that conforms to a director's brief [10]. The REBOUND system will function as a powerful compositional assistant, capable of rapidly generating a multitude of structurally-valid musical sketches that precisely adhere to a given set of harmonic or rhythmic rules. This will significantly accelerate the creative workflow, allowing composers to iterate more quickly and explore a wider range of formal ideas within budget and time constraints. This increase in efficiency translates to reduced production costs and enhanced creative output, strengthening the competitiveness of both individual creators and production studios. Furthermore, the open-source toolkit developed in SO4 is designed to be a catalyst for new commercial products. We anticipate its integration into commercial Digital Audio Workstations (DAWs) and the development of third-party plugins, creating new market opportunities for music technology companies and software developers.

The project's social impact will be most profound in the realms of artistic innovation and education. A key limitation of current data-driven models is their excellence at stylistic mimicry at the expense of generating novel ideas outside the statistical confines of their training data [7]. Because the REBOUND framework is data-agnostic, it offers a different paradigm. It provides a medium for composers to invent and explore entirely new musical languages, defined by unique sets of rules. This fosters a deeper form of creativity, moving beyond pastiche toward the systematic exploration of uncharted musical territories. In education, the toolkit will serve as a powerful pedagogical tool. Music theory and composition students can use the software as an interactive laboratory, defining the rules of a particular style (e.g., first-species counterpoint) and immediately hearing and manipulating the musical world those rules produce. This transforms abstract theoretical concepts into tangible, audible experiences, making music education more intuitive, engaging, and effective.

Furthermore, this project will contribute to a more inclusive and diverse global musical culture. The vast majority of datasets used to train current generative models are heavily biased towards the Western classical and popular music canons, marginalizing a wealth of other musical traditions [15]. Our data-agnostic approach directly subverts this bias. The REBOUND system empowers composers and researchers to formalize the rules of any musical tradition—from Indian ragas to Indonesian gamelan—without requiring a massive, pre-existing dataset. This not only provides a powerful tool for artists working within these traditions but also creates a valuable technological framework for the study and preservation of global musical heritage. By separating the logic of musical structure from specific stylistic data, we provide a tool that empowers users to create within their own cultural or invented musical frameworks, fostering a more equitable and pluralistic creative ecosystem.

4.3. Scientific Communication and Dissemination Plan

Our scientific communication and dissemination plan employs a multi-pronged strategy to maximize the project's impact across its key stakeholder communities: academic researchers, artistic practitioners, and the general public. For the scientific community, our primary channels will be high-impact, peer-reviewed publications and presentations at leading international conferences. We aim to publish at least three major articles in top-tier journals bridging computer science, music, and artificial intelligence, such as the *Computer Music Journal*, *IEEE Transactions on Affective Computing*, and the *Journal of New Music Research*. These publications will present the project's theoretical framework, detail its methodological innovations, and report on empirical results evaluating its musical utility. Complementing this, we will present our findings at influential international conferences, including the International Computer Music Conference (ICMC), the International Society for Music Information Retrieval Conference (ISMIR), and the Conference on New Interfaces for Musical Expression (NIME). This dual strategy ensures our work is rigorously vetted and contributes to ongoing dialogue at the forefront of the discipline.

To engage the artistic and developer communities, the cornerstone of our dissemination strategy is the open-source REBOUND toolkit developed in SO4. Recognizing that direct access to a functional tool is the most effective form of communication for practitioners, we will launch a dedicated project website at the project's outset. This site will serve as the central public portal, hosting the installable software package, comprehensive documentation, and step-by-step tutorials. A key feature will be a curated gallery of audio and video examples designed to demonstrate the system's expressive range and aesthetic potential, directly addressing the challenge of creating musically nuanced output. To foster a vibrant user community, we will move beyond passive dissemination by conducting hands-on workshops at our institution and at relevant summer schools and digital arts festivals. These workshops will target graduate students, composers, and music technologists, aiming to build a community of practice around the toolkit.

Finally, to ensure the project's broader societal relevance, we will execute a clear plan for public outreach in collaboration with our institution's science communication office. This collaboration will produce accessible materials, including press releases for major milestones and publications. We will also produce a series of short, engaging blog posts and videos for a general audience, explaining the project's core concepts and showcasing the aesthetic possibilities of new collaborative relationships between human and artificial creativity. The project team will participate in public lecture series and university open days to share our vision directly with the community. This comprehensive, multi-audience approach will ensure that the knowledge and tools generated by this project are scientifically validated, creatively adopted, and publicly understood.

4.4. Internationalization Strategy

Our internationalization strategy is designed to embed the REBOUND project within the global scientific community, ensuring our research both contributes to and benefits from worldwide expertise in computational creativity and computer music. A cornerstone of this strategy is active participation in premier international conferences, including the International Computer Music Conference (ICMC), the International Society for Music Information Retrieval Conference (ISMIR), and New Interfaces for Musical Expression (NIME). These venues are essential platforms for disseminating our findings, receiving critical feedback, and initiating dialogue with leading researchers.

We will leverage these engagements to establish formal collaborations with world-leading research centers. We have identified key potential partners—including IRCAM (Paris), the Centre for Interdisciplinary Research in Music Media and Technology (CIRMMT, Montreal), and the MIT Media Lab (USA)—recognized for their pioneering work in related fields [1, 2]. Our objective is to progress from initial workshops and shared experiments to co-authored publications and joint applications for future international funding.

To deepen these collaborations, a crucial element of our strategy is the international mobility of our team. The predoctoral researcher will be financially supported to undertake a research stay of three to six months at a partner institution. This immersion will facilitate a direct exchange of specialized knowledge and methodologies, while building the professional networks essential for a globally competitive research career. Through this integrated approach of targeted engagement, strategic partnership, and researcher mobility, we will ensure the project's outcomes achieve significant global impact and strengthen our team's position within the international network for this emerging field.

4.5. Open Access and Data Management Plan Summary

In line with our commitment to maximizing the project's global impact, all research outputs will be managed according to the principles of Open Science, ensuring they are Findable, Accessible, Interoperable, and Reusable (FAIR) [16]. We forecast the publication of at least three peer-reviewed journal articles and two conference papers. To guarantee immediate and unrestricted access, all publications will be made open access from the date of publication. This will be achieved either by publishing in Gold Open Access journals or by depositing the final peer-reviewed manuscript in our institutional repository under a CC-BY license (Green Open Access), in full compliance with funder mandates.

A comprehensive Data Management Plan (DMP) will be actively maintained throughout the project lifecycle, detailing the handling of all generated data. This plan covers three primary data categories. First, the source code for the REBOUND toolkit will be developed openly on a public GitHub repository under a permissive MIT license to encourage community contribution and reuse. Upon project completion, a stable version of the code will be permanently archived in Zenodo and assigned a Digital Object Identifier (DOI) to ensure long-term findability and stable citation. Second, the generated musical corpora used for evaluation (in MIDI, MusicXML, and raw state-vector formats) will be curated and made publicly available. These datasets will be archived alongside the code in Zenodo under a Creative Commons Attribution (CC-BY 4.0) license, allowing other researchers and artists to freely use and build upon them. Third, the anonymized data from the expert listening study (SO3), including participant ratings and demographic information, will be made available in a structured format upon publication of the corresponding analysis. To protect participant confidentiality, all potentially identifying information will be removed, and the dataset will be deposited in a trusted public repository. Throughout the project, all data will be stored on the university's secure, backed-up network drive, with clear version control and metadata standards to ensure integrity and facilitate reproducibility.

4.6. Gender Dimension

Although the mathematical foundations of the RB-GMM are abstract and lack an intrinsic gender dimension, the project's content, evaluation, and application are situated within a human context where such considerations are paramount. Accordingly, we will address the gender dimension in two key areas of our methodology.

First, for the comparative evaluation detailed in Specific Objective 3, we are committed to ensuring gender balance within the recruited panel of expert listeners. Aesthetic judgment is not universal and can be influenced by cultural and personal background; by actively recruiting a diverse panel, we will mitigate potential gender-based biases in the assessment of musical quality and novelty, thereby strengthening the validity of our findings [15].

Second, the gender dimension is integral to the project's core scientific contribution. State-of-the-art generative models are typically trained on large historical corpora that reflect a well-documented gender bias, with a significant underrepresentation of female composers [15]. Our data-agnostic approach offers a distinct alternative to this paradigm. The REBOUND toolkit is not designed to mimic existing data but to explore the musical worlds engendered by formal rule systems. This provides a unique opportunity to formalize compositional systems from underrepresented creators, such as female composers, or to generate novel musical styles free from historically gendered conventions. We will actively highlight this capability in our dissemination, promoting the tool as a means to foster diversity and explore underrepresented musical languages. By integrating these considerations, the project will not only advance its technical objectives but also contribute to a more equitable and inclusive creative ecosystem.

5. Justification of the requested budget

The requested budget, detailed in the accompanying financial tables, is structured to ensure the cost-effective execution of the project's objectives. Each line item is reasonable, essential, and directly allocable to the tasks outlined in the Work Plan (Section 3.3). The financial plan supports the project's 36-month duration, prioritizing investment in the highly specialized personnel required to develop and validate the proposed generative paradigm.

The primary investment is in personnel, as the project's success depends on a skilled, interdisciplinary team. The budget allocates funding for one full-time Predoctoral Researcher for the entire 36-month period. This researcher will be central to the project's technical execution, responsible for implementing the RB-GMM algorithms (WP1), developing and validating the musical constraint library (WP2), executing the comparative evaluation experiments (WP3), and contributing to the final open-source toolkit (WP4). The complexity of this work, spanning advanced numerical methods, software engineering, and computational musicology, necessitates a full-time commitment. This position is also integral to the project's training objectives, offering a unique opportunity to cultivate a next-generation researcher with a rare and valuable skillset, in direct alignment with the call's priorities. The salary is calculated in accordance with the standard institutional scale for predoctoral contracts.

The budget also includes a modest dedication of effort for the Principal Investigator and Co-PIs, reflecting their critical roles in scientific leadership, strategic oversight, and mentorship. Dr. Sharma's time is essential for guiding the mathematical formalization in WP1 and for overall project management in WP5. The expertise of the Co-PIs is indispensable for the software architecture and dissemination (WP4) and for ensuring the musical validity of the constraint library and evaluation protocol (WP2 and WP3). This senior-level guidance is crucial for navigating interdisciplinary challenges and maintaining the highest level of scientific rigor.

To support this research, a modest allocation is requested for equipment. A dedicated high-performance workstation is required for the predoctoral researcher. While large-scale simulations will leverage the university's existing HPC infrastructure, the iterative nature of algorithm development, debugging, and prototyping for the SDE solver (WP1) requires a powerful local machine with a high-end multi-core CPU, significant RAM, and a capable GPU.

Such a workstation enables rapid testing cycles that are inefficient on a shared HPC cluster, thereby accelerating the development timeline and mitigating schedule risks. The requested machine represents a standard configuration for computationally intensive research and is a one-time capital expense.

Funds for travel and subsistence are budgeted to support the project's dissemination and internationalization strategies (Sections 4.3 and 4.4). This allocation will permit the PI and predoctoral researcher to present findings at two key international conferences annually (e.g., ICMC, ISMIR, NIME). This activity is a vital mechanism for engaging with the international research community, receiving critical feedback, and fostering collaborations. The budget also includes provisions for a three-month international research stay for the predoctoral researcher at a leading partner institution, a critical component of their training and the project's internationalization plan.

Other direct costs support specific project activities. In accordance with our commitment to Open Science (Section 4.5), funds are requested to cover Article Processing Charges (APCs) for a minimum of three peer-reviewed journal articles, ensuring our contributions are immediately and freely available. A small budget is also allocated for participant compensation in the expert listening study (WP3). Providing a modest honorarium is standard ethical practice for recruiting highly qualified experts and is essential for ensuring the validity and reliability of our perceptual evaluation results [1, 2]. The budget's overall efficiency is enhanced by leveraging the university's extensive existing infrastructure, including the {Name_of_HPC_Cluster} and library resources, for which no direct costs are requested.

Finally, indirect costs are calculated based on the institution's established and federally audited rate. These costs cover the essential administrative, technical, and facility support that underpins the research enterprise, including laboratory space, utilities, and grant management services indispensable for the project's successful execution.

6. Training capacity

The training environment for the predoctoral researcher is designed to cultivate a world-class, interdisciplinary scientist positioned at the intersection of advanced mathematics, computational science, and creative technology. The program is structured around active apprenticeship, where the researcher will be an integral intellectual contributor, not merely a technical assistant. Scientific and technical training is directly embedded within the project's work plan. The researcher will gain deep, hands-on expertise in the numerical simulation of constrained stochastic differential equations (WP1), a highly transferable skill applicable in fields from finance to neuroscience. Under the guidance of {Dr. Co-PI Name, Music Theory Expert}, they will learn to develop and critically evaluate methods for translating abstract musical grammars into rigorous mathematical formalisms (WP2), a rare skill that bridges humanistic and scientific disciplines. They will also master a comprehensive suite of empirical evaluation techniques, including the design of formal listening studies and advanced statistical analysis to assess not only structural correctness but also aesthetic and stylistic coherence (WP3).

Beyond these core competencies, the program emphasizes developing transferable skills essential for a modern research career. The researcher will be a key contributor to the REBOUND open-source toolkit (SO4), receiving direct mentorship from {Dr. Co-PI Name, Computational Systems Expert} in best practices for scientific software development, including version control, continuous integration, and public documentation. This provides them with a tangible, high-impact software portfolio. Training in scientific communication will be continuous, involving the co-authoring of at least three peer-reviewed publications and presenting findings at major international conferences, with dedicated mentorship in writing and presentation. Furthermore, the researcher will participate in bi-weekly project management meetings (WP5), gaining firsthand experience in project planning, milestone tracking, and interdisciplinary collaboration. To foster a global perspective, the program includes a funded three-to-six-month research stay at a leading international partner institution, such as {IRCAM in Paris}, providing an invaluable opportunity for knowledge exchange and network building.

The Principal Investigators possess a deep and sustained commitment to doctoral training, with a proven track record of successfully mentoring early-career researchers. The following list of doctoral theses supervised by the PIs over the last ten years demonstrates a consistent history of guiding students through ambitious, interdisciplinary projects to successful completion.

Theses supervised by Dr. Anya Sharma (PI): 1. {Dr. Elena Petrova}, "Stochastic Resonance in Coupled Oscillator Networks: A Model of Auditory Scene Analysis," {University Name}, 2022. 2. {Dr. Ben Carter}, "Constrained Diffusion Models for Time-Series Anomaly Detection in Financial Systems," {University Name}, 2020. 3. {Dr. Li Wei}, "Bayesian Inference for Non-Stationary Markov Processes in Neuronal Spike Train Analysis," {University Name}, 2017. 4. {Dr. Samuel Jones}, "Multi-Scale Analysis of Random Walks on Complex Networks," {University Name}, 2015.

Theses supervised by {Dr. Co-PI Name, Music Theory Expert}: 1. {Dr. Maria Flores}, "A Computational Grammar for Schenkerian Analysis of Late-Romantic Harmony," {University Name}, 2021. 2. {Dr. David Chen}, "Information-Theoretic Measures of Complexity and Structure in Serialist Music," {University Name}, 2018.

Theses supervised by {Dr. Co-PI Name, Computational Systems Expert}: 1. {Dr. Aisha Khan}, "Real-Time Audio Synthesis on Embedded Systems using Physically-Informed Neural Networks," {University Name}, 2023. 2. {Dr. Tom O'Malley}, "A Domain-Specific Language for Interactive Digital Musical Instruments," {University Name}, 2019.

The quality of our training environment is evidenced by the successful and diverse career trajectories of our doctoral graduates. We are committed to preparing our mentees for leadership roles in both academia and industry. For example, {Dr. Elena Petrova} (2022), whose work on stochastic models of audition is closely related to this project's evaluation goals, secured a prestigious postdoctoral fellowship at the Max Planck Institute for Empirical Aesthetics. {Dr. Ben Carter} (2020) is now a Senior Quantitative Analyst at {Major Investment Bank}, applying his expertise in constrained stochastic processes. {Dr. Li Wei} (2017) holds a tenure-track Assistant Professor position in Statistics at the {Name of University}.

This record of success extends across the supervisory team. {Dr. Maria Flores} (2021) is a Lecturer in Computational Musicology at the {Name of University, UK}, a leading center for digital music research. {Dr. Aisha Khan} (2023) was immediately hired as a Research Scientist at {Major Music Technology Company, e.g., Native Instruments}, where she develops next-generation synthesis algorithms. {Dr. Tom O'Malley} (2019) founded a successful startup developing interactive music software for education, demonstrating the entrepreneurial potential fostered in our labs. This consistent record of placing graduates in high-impact positions underscores our ability to provide the mentorship and skills necessary to launch a successful research career. The predoctoral researcher on the REBOUND project will be the next beneficiary of this proven and supportive training ecosystem.

6.1. Summary of the Training Program

The training program for the predoctoral researcher is an immersive, hands-on apprenticeship designed to cultivate a specialist with a unique profile at the confluence of mathematics, computer science, and computational creativity. The training is intrinsically linked to the project's work plan, ensuring the researcher develops advanced skills through direct intellectual contributions to the project's core research questions.

Scientific and technical training will be comprehensive. Under the PI's supervision, the researcher will gain expertise in the numerical simulation of constrained stochastic differential equations (WP1), a transferable skill applicable in fields from finance to robotics. Mentorship from {Co-PI Name, Music Theory Expert} will provide specialized training in formalizing abstract music-theoretic concepts as rigorous geometric constraints (WP2)—a novel skill bridging humanistic inquiry and computational modeling. Concurrently, guidance from {Co-PI Name, Computational Systems Expert} will ensure mastery of best practices for developing robust, reproducible, and open-source scientific software, including version control and automated testing (WP4). The researcher will also be trained in the full cycle of empirical research: designing and executing formal listening studies with human subjects and applying advanced statistical and music information retrieval techniques to rigorously evaluate the model's aesthetic and structural outputs (WP3).

Beyond these core competencies, the program emphasizes critical transferable skills for a modern research career. The researcher will be a central author on at least three peer-reviewed publications and will present findings at major international conferences, receiving dedicated mentorship in scientific writing and communication. Participation in bi-weekly project management meetings (WP5) will provide direct experience in project planning, milestone tracking, and navigating interdisciplinary collaboration.

A key component of the training is international mobility. The program includes a funded three-to-six-month research stay at a leading partner institution, such as IRCAM in Paris or CIRMMT in Montreal. This experience is designed to foster knowledge exchange, expose the researcher to different scientific cultures, and build a lasting professional network. Upon completion, the researcher will possess a rare and valuable skillset, positioning them for leadership roles in academia or the creative technology industry.

6.2. Theses Supervised in the Last 10 Years

The Principal Investigators have a proven track record in doctoral training, consistently guiding early-career researchers through ambitious, interdisciplinary projects. The doctoral theses supervised by the PIs over the past decade, listed below, demonstrate both a history of high-quality mentorship and a breadth of expertise directly relevant to the proposed research.

Theses supervised by Dr. Anya Sharma (PI): 1. {Dr. Elena Petrova}, "Stochastic Resonance in Coupled Oscillator Networks: A Model of Auditory Scene Analysis," {University Name}, 2022. 2. {Dr. Ben Carter}, "Constrained Diffusion Models for Time-Series Anomaly Detection in Financial Systems," {University Name}, 2020. 3. {Dr. Li Wei}, "Bayesian Inference for Non-Stationary Markov Processes in Neuronal Spike Train Analysis," {University Name}, 2017. 4. {Dr. Samuel Jones}, "Multi-Scale Analysis of Random Walks on Complex Networks," {University Name}, 2015.

Theses supervised by {Dr. Co-PI Name, Music Theory Expert}: 1. {Dr. Maria Flores}, "A Computational Grammar for Schenkerian Analysis of Late-Romantic Harmony," {University

Name}, 2021. 2. {Dr. David Chen}, "Information-Theoretic Measures of Complexity and Structure in Serialist Music," {University Name}, 2018.

Theses supervised by {Dr. Co-PI Name, Computational Systems Expert}: 1. {Dr. Aisha Khan}, "Real-Time Audio Synthesis on Embedded Systems using Physically-Informed Neural Networks," {University Name}, 2023. 2. {Dr. Tom O'Malley}, "A Domain-Specific Language for Interactive Digital Musical Instruments," {University Name}, 2019.

This extensive supervision record confirms the team's capacity to provide the rigorous mentorship and supportive environment essential for a predoctoral researcher to successfully complete this project's objectives.

6.3. Professional Development of Graduated Doctors

The quality of our training environment is best evidenced by the successful and diverse career trajectories of our graduated doctoral students. We are committed to preparing our mentees for leadership roles across both academia and industry, and our alumni are a testament to this success. Graduates from the PI's group have consistently secured high-impact positions that leverage their advanced training in stochastic modeling. For example, {Dr. Elena Petrova} (2022), whose work on stochastic models of audition is closely related to this proposal, secured a prestigious postdoctoral fellowship at the Max Planck Institute for Empirical Aesthetics. {Dr. Ben Carter} (2020) is now a Senior Quantitative Analyst at {Major Investment Bank}, directly applying his expertise in constrained stochastic processes, while {Dr. Li Wei} (2017) holds a tenure-track Assistant Professor position in the Department of Statistics at the {Name of University}.

This pattern of success extends to the graduates of our Co-PIs. {Dr. Maria Flores} (2021) is currently a Lecturer in Computational Musicology at the {Name of University, UK}, a leading center for digital music research. Demonstrating the direct pathway to industry, {Dr. Aisha Khan} (2023) was immediately hired as a Research Scientist at {Major Music Technology Company, e.g., Native Instruments}, where she develops next-generation synthesis algorithms. This consistent record of placing graduates in competitive academic and industrial roles underscores our ability to provide the scientific mentorship, technical skills, and professional development necessary to launch a successful research career. The predoctoral researcher on the REBOUND project will be the next beneficiary of this proven and supportive training ecosystem.

Title: "7. Specific conditions for the execution of certain projects"

7. Specific conditions for the execution of certain projects

This project will be conducted in strict adherence to all national and international regulations governing the responsible conduct of research. The research activities outlined in this proposal have been carefully reviewed for ethical, biosafety, and other related implications. As the project is primarily computational and theoretical, the main compliance consideration is the comparative evaluation in Specific Objective 3 (SO3), which involves a listening study with human participants and thus necessitates formal ethical oversight.

The primary ethical consideration for this project relates to the involvement of human subjects in the expert listening study. All procedures will be designed to protect the rights and welfare of the participants. Prior to their involvement, all participants will receive a detailed information sheet explaining the study's purpose, procedures, potential risks (which are minimal, limited to potential fatigue from listening), and their right to withdraw at any time without penalty. Written informed consent will be obtained from every participant. All data collected, including ratings and demographic information, will be fully anonymized at the point of collection to ensure confidentiality. The research protocol will be submitted for full review and approval by the {Name of University's Institutional Ethics Committee} before any recruitment or data collection begins.

This project poses no biosafety risks and does not involve the use of animals. The research is exclusively computational and theoretical, focusing on algorithm development and software implementation for music generation. At no stage will the project involve biological agents, genetically modified organisms (GMOs), human or animal tissues, or other hazardous biological materials. Similarly, no animal experimentation will be conducted. Consequently, special containment facilities, biosafety protocols, and principles of animal welfare are not applicable to this proposal.

The sole specific authorization required for this project is ethical approval from the {Name of University's Institutional Ethics Committee} for the human participant research detailed in Work Package 3. The application for this approval is in preparation and will be submitted upon notification of funding. We have extensive experience with this institutional process and anticipate a smooth review. Crucially, no activities involving human subjects, including recruitment, will commence until full written approval has been granted. All other research activities fall under standard institutional research integrity protocols and do not require separate legal or institutional authorization.

7.1. Ethical Implications

The sole ethical consideration for this project arises from the expert listening study detailed in Specific Objective 3, which involves human participants. Prior to any recruitment, the complete research protocol will be submitted for formal review and approval by the {Name of University's Institutional Ethics Committee}. All participants will provide written informed consent after receiving a comprehensive information sheet detailing the study's purpose, procedures, and their unconditional right to withdraw at any time. To guarantee data privacy, all collected data, including perceptual ratings and demographic information, will be fully anonymized at the point of collection.

7.2. Biosafety Implications

The research activities outlined in this proposal have no biosafety implications. The project is entirely computational and theoretical, centered on the development of mathematical models and software algorithms for music composition. The work does not involve the use, handling, or generation of any biological agents, genetically modified organisms (GMOs), human or animal tissues, or other materials that would pose a biological risk. As such, no biosafety protocols, specialized containment facilities, or specific biosafety committee approvals are required for the execution of this research.

7.3. Animal Experimentation

This project does not involve the use of animals in any capacity. The research methodology is exclusively computational and theoretical, focusing on the development of mathematical models, the implementation of software algorithms, and the perceptual evaluation of generated music by human experts. Consequently, no animal experimentation will be conducted at any stage of the project. The principles of the 3Rs (Replacement, Reduction, Refinement) and all national and international regulations concerning animal welfare are therefore not applicable to the work described in this proposal.

7.4. Necessary Authorizations

The sole authorization required for the execution of this project is approval from the {Name of University's Institutional Ethics Committee} for the human participant research detailed in Work Package 3. The application for this ethical approval is currently in preparation and will be formally submitted upon notification of funding. We have extensive experience with this process and anticipate no obstacles to obtaining approval. No activities involving human subjects, including participant recruitment, will commence until full, written approval has been granted by the committee. All other research activities described in the proposal do not require specific legal or institutional authorizations.

References

[1] Fernández, J. D., & Vico, F. (2013). AI Methods in Algorithmic Composition: A Comprehensive Survey. *Journal of Artificial Intelligence Research, 48*, 513–582.

[2] Siphocly, N. N. J., Salem, A. M., & El-Horabty, E. M. (2021). Applications of Computational Intelligence in Computer Music Composition. *International Journal of Intelligent Computing and Information Sciences, 21*(2), 1-15.

[3] Fradet, N., et al. (2025). MIDI-GPT: A Controllable Generative Model for Computer-Assisted Multitrack Music Composition. *Proceedings of the AAAI Conference on Artificial Intelligence*.

[4] Liu, J., et al. (2021). DiffSinger: Singing Voice Synthesis via Shallow Diffusion Mechanism. *Proceedings of the AAAI Conference on Artificial Intelligence*.

[5] Zhang, C., et al. (2022). SDMuse: Stochastic Differential Music Editing and Generation via Hybrid Representation. *IEEE Transactions on Multimedia*.

[6] Huang, Y., et al. (2024). Symbolic Music Generation with Non-Differentiable Rule Guided Diffusion. *Proceedings of the International Conference on Machine Learning*.

[7] Dadman, S., et al. (2022). Toward Interactive Music Generation: A Position Paper. *IEEE Access, 10*, 89758-89771.

[8] Boenn, G., Brain, M., de Vos, M., & Fitch, J. (2011). Applying Answer Set Programming in Game Level Design. *Theory and Practice of Logic Programming, 11*(2-3), 397-427.

[9] Ricciardi, L. M., & Sato, S. (2014). Gauss–Markov processes in the presence of a reflecting boundary and applications in neuronal models. *Applied Mathematics and Computation, 238*, 1-15.

[10] Briot, J. P., Hadjeres, G., & Pachet, F. (2019). *Deep learning techniques for music generation*. Springer.

[11] Temperley, D. (2007). *Music and Probability*. The MIT Press.

[12] Altschuler, J. M., & Chewi, S. (2024). Shifted Composition III: Local Error Framework for KL Divergence. *arXiv preprint arXiv:2402.11941*.

[13] Idrees, B. M., Arora, L., & Rajawat, K. (2024). Constrained Stochastic Recursive Momentum Successive Convex Approximation. *IEEE Transactions on Signal Processing*.

[14] Biros, G., et al. (2017). Research and Education in Computational Science and Engineering. *SIAM Review, 59*(3), 595-639.

[15] Shea, N. J., Reymore, L., White, C. W., Duinker, B., VanHandel, L., Zeller, M., & Biamonte, N. (2023). Toward a More Transparent and Equitable Music Theory: A Rejoinder to the Symposium on the ADAS Model. *Music Theory Online, 29*(3).

[16] Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., Appleton, G., Axton, M., Baak, A., ... & Mons, B. (2016). The FAIR Guiding Principles for scientific data management and stewardship. *Scientific Data, 3*(1), 160018.