

Comparative Analysis of Key Inference Models for Musical Metacreation

Paul M. Bodily and Dan Ventura

Department of Computer Science
Brigham Young University
Provo, UT 84602-6576

Abstract

Creative musical systems must be equipped with certain intelligent abilities to understand fundamental aspects of music, particularly in order to autonomously interact with other creative agents. Key inference, a relatively simple task for trained human experts, is one such intelligent ability required to normalize and analyze melodic compositions. We assess the accuracy of several traditional and machine learning approaches to the key and key signature inference problems on a dataset of 480 melodies in MIDI format. We compare these accuracies with those of trained human musicians on the same tasks. We evaluate the impact of including note duration and note repetition as learning features. We find machine learning approaches outperform traditional key inference methods. The highest accuracies (0.729 for key inference and 0.896 for key signature inference) were achieved using a 4-gram language model. Including note duration improved the results of traditional approaches when inferring key, but had the opposite effect for key signature inference. Our findings suggest that the key of a melodic passage depends more heavily on the sequence of the notes rather than their frequency or distribution.

Introduction

A fundamental element of musical metacreation is the quest to “[endow] machines with creative behavior” (Pasquier, Eigenfeldt, and Bown 2012). Such behavior in humans draws on a full spectrum of intelligent abilities (Colton, Wiggins, and others 2012). Often there are overlaps in the intelligent abilities that are required for different tasks. This has been noted to be especially true of musical computational creative systems (Bodily and Ventura 2017).

An example of such an ability is *musical key inference*: the ability to recognize the group of pitches or scale that form the basis of a particular composition. This ability is a relatively simple task for trained human musicians. The key or tonic of a passage is used to transpose passages, motifs, and progressions to a common tonal center in order to be able to compare relationships and functions of pitched elements across musical selections in different keys.

This work is licensed under the Creative Commons “Attribution 4.0 International” licence.

In computational systems data often comes pre-labeled with the key. However, humans and music systems alike are often placed in unsupervised training scenarios in which the key is not labeled, but must rather be inferred. This inference is critical, particularly for interactive agents which must interpret and respond to other music agents. The key in which systems improvise or compose is often not explicitly communicated by other performing agents but is rather inferred from what *is* explicitly communicated (i.e., notes, chords, etc.). In the quest for creating autonomous, interactive musical agents, to what extent can this skill of inferring musical key be learned?

The *key* or tonality is a root pitch or *tonic* and modality (e.g., major, minor, etc.) which forms the structural basis for Western art music from Baroque to Romantic, along with much modern popular music. This tonality provides a context within which “the melodic and harmonic unfolding of a composition takes place” (Vos and Van Geenen 1996). Even the untrained ear appreciates the structure, tension, and resolution that tonality provides. Each key is associated with a *key signature* representing the specific pitches that belong to the scale represented by the key. There exists a many-to-one relationship between keys and key signatures. For example, each major key and its *relative natural minor key* (whose root is three half steps lower) are associated with the same key signature. Though music certainly exists with keys whose modes are not strictly major or minor, we have as a simplifying assumption chosen to focus solely on these two common modes. Note that we are not interested in inferring harmonic accompaniment (cf. (Groves 2013)), but rather in determining the main key and key signature of a piece of music.

Our interest in this problem arises as both a matter of necessity and as a philosophical inquiry. In efforts to apply machine learning to the generation of pop lead sheets, much of the freely available training data in the pop genre exists in MIDI format. Unlike more amenable formats (e.g., Music XML), MIDI files do not typically contain information about key or key signature. In order to train machine learning models on this data, inferring key and/or key signature is useful for being able to “normalize” pitches into a common key. Also for viewing MIDIs with music notation software (e.g., MuseScore), key signature is an important notational element.

Besides these practical applications, the question is of broader philosophical interest: how is it that a human musician can listen to a musical selection and infer its key? There is a profound irony in the contrast between knowing how and being able to explain how to infer the key of a musical selection. Expert musicians routinely and accurately perform this task. However the description for their methodology is often inexact and derives from finding the key that “feels” right, provides a sense of “finish”, or the key where the music “lands”. One common supposition is that the solution involves finding the key which minimizes the total number of accidentals in a composition (Gorow 2011). Part of our purpose is to test this hypothesis by comparing this approach with several machine learning algorithms. Examining algorithms that perform well on this task may help provide some insight into how human musicians infer key.

We hypothesize that a monophonic melodic sequence alone is sufficient in many cases to determine key. We have chosen to focus on monophonic melodies, again, for both practical and philosophical reasons. First, though compositions have varying amounts of harmonic context, the monophonic melody acts as a common denominator in all compositions in our dataset. Second, humans—in singing, humming, whistling, or playing a monophonic instrument—are able to communicate and interpret music based on monophonic melody. Picture, for example, singing a melody to a musical friend that had never heard it. Without specifying either the key or tonic, the friend is able to infer this information with some probability in order to make sense of what (s)he is hearing (e.g., melody starts on the major third, ends with the tonic, etc.). Thus the ability to infer tonality for monophonic melodies represents an AI task if computers are to effectively understand and communicate with humans.

We test our hypothesis on a dataset of 480 pop melodies in MIDI format¹. Data were selected to exclude melodies with modulation or key changes. As a fundamental building block of music knowledge and creativity, effective key inference models stand to add increase autonomy in musical metacreation systems.

Related Work

Several previous studies have examined key inference in various contexts, though to our knowledge ours is the first attempt to do so using n -gram models with monophonic pop melodies.

Krumhansl matches the relative frequencies and durations with which tones are sounded (which she terms a *tonal hierarchy*) of a composition against the known tonal hierarchies of each key (2001). This algorithm was applied to infer keys for compositions from three classical composers: Bach, Chopin, and Shostakovich.

The key-finding algorithm of Longuet-Higgins and Steedman successively eliminates keys based on the presence or absence of the notes of the composition in each of the major and minor scales (Longuet-Higgins and Steedman 1971). Holtzman (1977) infers key from the prevalence

¹Metadata are available at http://popstar.cs.byu.edu/metadata_for_midis.csv

of common key-defining intervals (e.g., triads, tonic-fifths, tonic-thirds). Both algorithms were applied to Bach’s *Well-Tempered Clavier*.

Hu and Saul use Latent Dirichlet Allocation (LDA)—a statistical approach for discovering hidden topics in large corpora of text—for key-finding, looking for common co-occurrences of notes in music. Their model essentially treats keys like topics. They then model compositions as a random mixture of key-profiles, allowing them to track modulations (2009). Temperley interprets the traditional key-profile model (proposed by Krumhansl (2001) and subsequently modified by Temperley (1999)) as a Bayesian probabilistic model and discusses the implications of the connection between these two models (2002). All applications of the model are on the Kostka-Payne corpus, a collection of textbook excerpts of tonal music. Vos and Van Geenen present a parallel search key-finding algorithm for single-voiced music. Notes are evaluated against both the scalar and the chordal structures of each key. They demonstrate the model’s effectiveness on Bach’s *Well-Tempered Clavier* (1996). Zhu et al. present a method for key estimation in audio pop and classical music (2005). Their method performs marginally higher with pop music than with classical music.

Much recent work has been devoted to inferring key from audio music. Shenoy et al. outline a rule-based algorithm for finding key from audio musical signals using chroma based frequency analysis and chord progression patterns (2004). Mauch and Dixon infer chords and key simultaneously from audio (2010). Chafe et al. extract symbolic music from audio from which meter and key are inferred (in that order) (1982). The key-recognizer assumes that rhythmic and melodic accent points are significant features for inferring key.

Temperley and Clercq (2013) address key-finding in rock songs from harmonic, melodic, and combined harmonic-melodic data. They compare a key-profile method (a normative distribution of scale-degrees) and a key-profile method weighted by note durations for key-finding in monophonic rock melodies. Their findings suggested that the weighted method worked somewhat better. Their work provides an excellent survey of the work done on key-finding algorithms in both pop and classical music, both from audio and sheet music.

Methods

Data

We collected 480 melodies from 278 pop artists representing music spanning several decades (see Figure 1). Songs were compiled from several online MIDI databases and keys (inferred from melody together with harmonic context) were manually labeled by trained musicians. We were pleased to find that every key was represented by at least two songs (see Figure 1). Only songs with a single key were selected for our experiments (no modulations). Melody notes were isolated from the MIDI files. The melody for the entire song was used for inference purposes.




Figure 1: *Key distribution of pop melodies*. Every key is represented at least twice in the subset of melodies. In general keys with fewer sharps and flats are more highly represented.

Table 1: Highly Represented Pop Artists

Artist	# of songs
Beatles	28
Elvis Presley	13
Kiss	11
Madonna	8
Eagles	6
Aerosmith	6
Elton John	6
U2	6
Beach Boys	5
Michael Jackson	5
Pink Floyd	5
Bobby Vee	5
Adele	5
Queen	4
Kinks	4

Implementation

We compared the accuracy of 4 traditional and 5 machine learning methods against the accuracy of trained musicians (different than those employed for labeling ground truth) on the tasks of inferring key and key signature for pop melodies². For each key signature there is a major and minor key and thus key inference—which precludes key signature inference—represents the more difficult of the two tasks. For the machine learning methods, we used 10-fold cross validation for training and testing.

Minimize Accidentals by Count. This algorithm represents the common theory that the best key is that which minimizes the number of resulting accidentals.

Minimize Accidentals by Duration. Similar to the previous method but finds the key which minimizes the total duration of accidentals.

Root Mean Squared Error (RMSE) of Pitch Count Profiles. Similar to the method followed by Krumhansl (2001) and others, we generated pitch profiles. Rather than generate profiles for each major and minor key, we chose to generate a single profile for all major keys and a second for all minor keys. This decision is based on the assumption that the variation in pitch distribution varies very little (if at all) as compared to the distributional variation between the major and minor modes. Major and minor mode pitch profiles were generated from the pitch counts in training instances transposed to either the C major or A minor keys depending on whether the instance was major or minor. A pitch profile is created for each test instance, transposed into each of the 12 possible keys. Each transposed profile is compared to both the major and minor generic pitch profiles using RMSE. The transposed pitch profile and the major/minor pitch profile with the minimum RMSE value are used to infer key and modality.

RMSE of Pitch Duration Profiles. Similar to the previous method but pitch profiles are generated from pitch durations rather than mere counts (see Figure 2).

n-gram Models. An n-gram model calculates the probability of the next token given some context window of length n (Brown et al. 1992). These probabilities are learned from the sequence of notes in the training instances and then used to calculate the probability of note sequences in the test instances. We trained n -gram models for values of n from 1 to 5, using Laplace smoothing and a pseudocount alpha value of 1. For each value of n a single n -gram model was trained for melodies in major keys and another for melodies in minor keys. Probabilities were normalized across both models. Test instances were then transposed and scored by each trained model. The transposition and model which maximized the probability of the test instance determined the key and modality.

In addition to the methods described above, we also report accuracy from four other sources. First, the *baseline* accuracy represents the approach of always guessing the most common class. *MIDI Annotations* refers to the key signature that was originally given in the MIDI file (if any was pro-

²source available at http://popstar.cs.byu.edu/find_key.py




Figure 2: *Weighted Major/Minor Pitch Profiles*. Profiles are based on the duration of pitches in each of the major and minor modes. Profiles for major keys (blue) have been normalized to C major and those for minor keys (red) to A minor. Note that the A pitch is relatively more frequent for minor keys (where it functions as the tonic) than for major keys (where it functions as the 6th from the root). Likewise the G pitch is relatively more frequent for major keys (where it functions as the dominant) than for minor keys (where it functions as the 7th from the root).

vided). We report the accuracy of a third-party MIDI-reader called *MuseScore*. Insofar as MIDI is not a symbolic music format, many files fail to include (accurate) information about key or time signature, thus motivating the need to infer this information from the notes themselves. This functionality is built in with varying success to many programs which render MIDI files as sheet music. *MuseScore* (version 2.0.3) is one such program. Despite not having access to the *MuseScore* algorithm, we manually opened and examined each song from the database in *MuseScore* to record the result of their key inference algorithm. We include the accuracy of *MuseScore*’s inferred key-signature in our results for comparison.

Lastly, we enlisted 15 trained musicians to complete the same melodic key inference task. Each musician was afforded access to a keyboard and asked to listen to 20 melodies (the first of which was a control melody, *God Rest Ye Merry Gentlemen*—all respondents answered the control key correctly) without access to printed scores. In addition to labeling the key and key signature for each melody, respondents were asked to indicate how familiar they were with the melody³. This was to control for bias that might arise as a result of prior knowledge of the harmonic context of the melody. We report the accuracies on the melodic key inference task as a function of the degree of familiarity as well as an overall weighted average of these three scores.

³Survey results and self-reported musician training information are available at http://popstar.cs.byu.edu/raw_survey_results_name_key.csv

Table 2: Key Finding Accuracy for Pop Melodies

Method	Key	Key Signature
Baseline (C)	0.144	0.202
MIDI Annotations	n/a	0.483
MuseScore	n/a	0.746
Minimize Accidentals by Count	0.494	0.660
Minimize Accidentals by Duration	0.490	0.665
RMSE of Pitch Count Profiles	0.606	0.796
RMSE of Pitch Duration Profiles	0.613	0.765
Unigram model	0.629	0.815
Bigram model	0.654	0.883
Trigram model	0.679	0.883
4-gram model	0.729	0.896
5-gram model	0.700	0.885
Human (“I know this song”)	0.846	0.865
Human (“Sounds familiar”)	0.75	0.786
Human (“I don’t know this song”)	0.676	0.815
Human Average	0.719	0.822

Results

Results are shown in Table 2. We report accuracy both for inferring the key and for inferring the key signature. Inasmuch as key signature is a more generic classification of key (e.g., C major and A minor both have a key signature with no flats or sharps), accuracy for key signature will always be better than accuracy for key.

Discussion

Normalizing to a common key really only requires that we identify the key signature for a composition without regard for whether the key is the major key associated with the key signature or its relative minor. We chose to model the major and minor separately based on the hypothesis that the melodies that each produces would be sufficiently different to warrant creating individual profiles. The key inference methods which minimize the frequency of accidentals show the most dramatic improvement because these methods inherently fail to provide a way of distinguishing between a major key and its relative minor.

We encountered several challenges, primarily centered around assumptions made about the modalities present in our dataset. Songs based on blues scales often include a flat seven (which would suggest a key a fifth below the actual tonic) or both the major and minor third. Hard rock songs often exclude the third all-together, making it difficult to infer whether a major or minor key signature is more accurate. This finding supports the argument made by Temperley and Clercq (2013) and others against the applicability of the major/minor dichotomy to rock music. These confounding influences are reflected in the confusion matrices for classifications of songs in these genres (see Figure 3).

As regards the n -gram models, we note that the unigram model is essentially equivalent to a pitch profile rendered as an applied probability distribution, and thus it seems reason-

		Predicted																							
		Major						Minor																	
Actual		C	C#	D	D#	E	F	F#	G	G#	A	A#	B	C	C#	D	D#	E	F	F#	G	G#	A	A#	B
		Major	C	57	0	0	0	0	2	0	3	0	0	0	0	6	0	0	0	0	1	0	0	0	0
C#	0		9	0	0	0	0	0	0	0	0	0	0	0	3	0	0	0	0	0	0	0	0	0	0
D	0		0	26	0	0	0	0	2	0	1	0	0	0	0	5	0	0	0	0	0	0	0	0	0
D#	0		0	0	15	0	0	0	0	0	0	1	0	0	0	0	4	0	0	1	0	0	0	0	0
E	0		0	0	0	20	0	0	1	0	0	0	1	0	0	1	0	2	0	0	5	0	0	0	0
F	0		0	0	1	0	40	0	0	0	0	1	0	0	0	0	0	0	1	0	0	3	0	0	0
F#	0		0	0	0	0	0	8	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1
G	0		0	1	0	0	0	0	49	0	0	0	0	0	0	0	1	0	0	0	5	0	0	1	0
G#	0		0	0	0	0	0	0	0	10	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0
A	1		0	3	0	1	1	0	0	0	27	0	0	1	0	1	0	0	0	0	1	0	3	0	0
A#	0	0	0	0	0	0	0	0	0	0	21	0	0	1	0	0	0	0	0	0	0	0	1	0	
B	0	0	0	0	1	0	0	0	0	0	0	5	0	0	0	0	0	0	0	0	0	0	0	0	
Minor	C	11	0	0	0	0	0	0	2	0	0	0	0	15	0	0	0	0	0	0	0	0	0	0	0
	C#	0	2	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0
	D	0	0	5	0	0	0	0	0	0	0	0	0	0	0	3	0	0	0	0	0	0	0	0	0
	D#	0	0	0	9	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0
	E	0	0	0	0	5	0	0	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0
	F	0	0	0	0	0	3	0	0	0	0	0	0	1	0	0	0	0	9	0	0	0	0	0	0
	F#	0	1	0	0	0	0	3	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0
	G	0	0	1	0	1	0	0	5	0	1	0	0	0	0	1	0	0	0	0	19	0	0	0	0
	G#	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	2	0	0	0
	A	0	0	0	0	0	0	0	0	0	2	0	0	0	0	0	0	0	0	0	0	0	5	0	0
A#	0	0	0	0	0	0	0	0	0	0	2	0	0	0	0	0	0	0	0	0	0	0	7	0	
B	0	0	0	0	0	0	0	0	0	0	0	3	0	0	0	0	0	0	0	0	0	0	0	1	

Figure 3: *Key Inference Confusion Matrix for 4-gram model*. Much of the mis-classification confuses major and minor modes for the same key. In cases such as hard rock which use open fifths and exclude the third completely the mode (i.e., major or minor) is difficult to establish.

able that it should perform about on par with the RMSE of Pitch Count Profile method.

We increased the value of n until we observed a decrease in accuracy. As is typical of n -gram models, as n increases, the model begins to essentially memorize the training data at which point the differences between test and training sets become as significant as the differences between the inferred key classes.

It is important to note that although our best accuracy on key-finding (.729) is significantly below the values reported by studies mentioned in the related works, the task of inferring key from melody is significantly more challenging than inferring key from songs which include harmony. This is evidenced by the relative accuracies of the algorithmic and human performances on the melodic key inference problem.

It should also be considered that human listeners that are familiar with a melody may more accurately infer key from having familiarity with the harmony also. This is evidenced in our results in the fact that the more familiar the human lis-

teners were with the song being played, the higher the accuracy on the key inference problem. Notably the same was not necessarily true with the key signature inference, suggesting that without knowledge of the context, humans may confuse a major and relative minor for an unknown melody but still infer the correct key signature. This is also evidenced by the confusion matrix for the 4-gram model, shown in Figure 3. We therefore express confidence in the reported accuracies of these models.

It is of interest to note that the n -gram models generally outperformed the human models on the key signature inference task. The 4- and 5-gram models outperformed the human respondents on the key inference task when the human was otherwise unfamiliar with the song being played.

Our results suggest that considering pitch counts or durations is less effective than a model which considers the sequence of pitches. This agrees with our intuition insofar as many pop songs spend the bulk of their duration progressing through chords which are not the root and may not even be

closely related to the root. Notions of resolution and finding where the song “lands” inherently suggest that the contour and progression of the notes matters more than their frequency. The key is often most clearly defined at the beginning and ends of musical phrases or the beginning and end of the song itself. Whereas the method of counting note frequencies fails to give higher weight to these defining regions of the melodic passage, this information is embedded within the probabilistic framework of n -gram models which employ a special n -gram for the beginning and end of sequences.

We find the superior accuracy of n -gram models to the MuseScore key-signature inference model to be particularly promising inasmuch as it suggests that an n -gram model might be used to improve the state of the art in industrial MIDI-reading software.

As regards key inference from monophonic pop melodies, we find that machine learning methods (n -gram language models in particular) perform better than traditional key-finding algorithms, though both improve upon baseline accuracy. We also envision developing a framework for detecting key *modulations* in pop melodies and using normalized unlabeled melodic data for compositional analysis.

Conclusion

We have examined several solutions to the task of inferring key and key signature in monophonic music sequences. We found that for the particular case of monophonic sequences, the 4-gram model performed best at both identifying key and key signature. Endowing musical creative systems with intelligent abilities like key inference increases their autonomy. As key is a fundamental concept to many genres of music, this computational subconcept model can be modularly reused to improve musical metacreation across several domains.

References

Bodily, P. M., and Ventura, D. 2017. HBPL: a framework for debating, developing, and reusing foundational models of musical metacreativity. In *Proceedings of the Fifth International Workshop on Musical Metacreation*.

Brown, P. F.; Desouza, P. V.; Mercer, R. L.; Pietra, V. J. D.; and Lai, J. C. 1992. Class-based n -gram models of natural language. *Computational linguistics* 18(4):467–479.

Chafe, C.; Mont-Reynaud, B.; and Rush, L. 1982. Toward an intelligent editor of digital audio: Recognition of musical constructs. *Computer Music Journal* 6(1):30–41.

Colton, S.; Wiggins, G. A.; et al. 2012. Computational creativity: The final frontier? In *ECAI*, volume 12, 21–26.

Gorow, R. 2011. *Hearing and writing music: professional training for today's musician*. SCB Distributors.

Groves, R. 2013. Automatic harmonization using a hidden semi-markov model. In *Proceedings of the Artificial Intelligence and Interactive Digital Entertainment Conference (Boston, MA)*, 48–54.

Holtzman, S. R. 1977. A program for key determination. *Journal of New Music Research* 6(1):29–56.

Hu, D., and Saul, L. K. 2009. A probabilistic topic model for unsupervised learning of musical key-profiles. In *ISMIR*, 441–446. Citeseer.

Krumhansl, C. L. 2001. *Cognitive foundations of musical pitch*. Oxford University Press.

Longuet-Higgins, H. C., and Steedman, M. J. 1971. On interpreting bach. *Machine intelligence* 6:221–241.

Mauch, M., and Dixon, S. 2010. Simultaneous estimation of chords and musical context from audio. *IEEE Transactions on Audio, Speech, and Language Processing* 18(6):1280–1289.

Pasquier, P.; Eigenfeldt, A.; and Bown, O. 2012. Preface. In *Proceedings of the First International Workshop on Musical Metacreation*.

Shenoy, A.; Mohapatra, R.; and Wang, Y. 2004. Key determination of acoustic musical signals. In *Multimedia and Expo, 2004. ICME'04. 2004 IEEE International Conference on*, volume 3, 1771–1774. IEEE.

Temperley, D., and Clercq, T. d. 2013. Statistical analysis of harmony and melody in rock music. *Journal of New Music Research* 42(3):187–204.

Temperley, D. 1999. What's key for key? the krumhansl-schmuckler key-finding algorithm reconsidered. *Music Perception: An Interdisciplinary Journal* 17(1):65–100.

Temperley, D. 2002. A bayesian approach to key-finding. In *Music and artificial intelligence*. Springer. 195–206.

Vos, P. G., and Van Geenen, E. W. 1996. A parallel-processing key-finding model. *Music Perception: An Interdisciplinary Journal* 14(2):185–223.

Zhu, Y.; Kankanhalli, M. S.; and Gao, S. 2005. Music key detection for musical audio. In *11th International Multimedia Modelling Conference*, 30–37. IEEE.