# Audio Retrieval by Relevance Feedback

Michael Mandel & Graham Poliner

January 26, 2005

**Abstract**

In order to manage growing music collections, a personal music recommender could find new music, appropriate to the user's mood, that he or she would like to listen to. This paper approaches these goals using the flexible search technique of active SVM learning that adapts to users' perceptions instead of vice versa. In the best case, active SVM learning requires fewer than half the number of training examples a normal SVM classification would require to achieve the same precision and recall. In addition to the idea of applying active SVM learning to the audio domain, the paper has contributed a collection of ground truth classification of popular songs and a preliminary software implementation of this recommender.

## 1 Introduction

With the sizes of personal digital music collections growing, choosing music appropriate to a mood gets more difficult. Furthermore, finding music one would like to listen to both within a personal collection and in general is also a difficult task. Professional disk jockeys arrange similar music in particular orders to create and manipulate emotions in their audience. Our vision is to make an automatic personalized DJ for anyone's music collection that will find interesting, appropriate songs from a user's collection that are appropriate to that user's current state of mind and arrange them in an order that will both entertain and adjust the user's mood to a desired end point.

This project is a first step toward that vision, the step of finding acoustically similar music. We leave for later the problem of ordering that music and tracing trajectories through a mood space. Since finding songs that are similar to each other is time consuming and varies across individuals, we would like to create a flexible, open-ended approach to this task.

Our solution is to use relevance feedback, specifically Active Support Vector Machine learning, to learn a classifier for each individual query. With acoustic features calculated offline, a search is both a mapping from those low level features to higher level concepts and a way for the user to clarify the search to his or herself. Each search begins with a seed song or songs followed by iteration between a new classifier being trained on labeled data and the user's labeling informative examples. The search thus proceeds quickly and at every stage the system supplies its best estimate of songs similar to the seed song.

Using ground truth information on musical similarity culled from Allmusic.com, we have developed an automatic tester for our system. In tests with this harness the system usually outperforms a normal support vector machine (SVM) trained on the same number of data points and only rarely performs worse.

# 2 Ground Truth

As with any multimedia system, the second question to ask (after "what does it do?") is "how do you know that it works?" because in general ground truth data is hard to come by, biased, hard to access, or hard to align with an experiment. Part of what distinguishes a good experiment is solid proof of its success on test cases that are diverse enough to simulate most or all of the situations the system would encounter in real use. This proof requires accurate, relevant ground truth information. Our ground truth data comes from the All Music Guide website.

The All Music Guide (AMG) is a website (www.allmusic.com) and book that reviews, rates, and categorizes music and musicians. As ground truth for our automatic testing we used the "moods" AMG lists for artists, albums, and some songs. They defined moods as "adjectives that describe the sound and feel of a song, album, or overall body of work."[1] In particular, there are currently 178 moods including, for example, acerbic, campy, cerebral, silly, sleazy, rollicking, rustic, and hypnotic.

## 2.1 Collecting Moods

While AMG in general only ascribes moods to albums and artists, it does list the most exemplary 50 or so songs for each mood. Unfortunately for this experiment, they do not publish their database in a convenient form (without payment) and they do not like web robots crawling their site. A large part of collecting the ground truth, then, was trying to get the information useful to us off of AMG's web site. For this purpose we wrote a number of perl scripts, which grabbed the names of these top songs for every mood. Other perl scripts compared listings of our music collections with AMG's lists and compiled them into a large, uniform collection.

After comparing the contents of our music collections to the top songs from every mood, we settled on the 32 moods for which we had over 15 songs each. This led to a collection of 535 songs total, some of which, generally the most musically important, occupied the top spots in multiple moods.

# 3 Audio Features

Our representation of all of these songs is based on mel frequency cepstral coefficients (MFCCs) clustered into a Gaussian mixture model (GMM) and processed through the Fisher kernel approach first described by Jaakkola and Haussler in [3] and described for use in multimedia applications in [5].

MFCCs are a type of time-frequency decomposition of an audio signal meant to approximate the human perception of hearing. While originally intended for use in automatic speech recognition, it has found applications in other auditory domains like music retrieval [4]. Calculating MFCCs involves five steps. After breaking the signal into approximately 20ms long frames, a time scale at which the signal is assumed to be stationary, each frame is processed independently. A frame goes through a discrete cosine transform to work in the frequency domain, then the logarithm of the spectrum is taken to imitate the amplitude sensitivity of human hearing, then the frequency is resampled on the mel scale which emphasizes the frequencies most important to human hearing, and finally the frame goes through another discrete cosine transform to decorrelate any sinusoidal components from one another. All of the MFCCs contain 13 frequency bands and, depending on the length of the song, approximately 20,000 temporal frames.

Using 1% of the MFCCs selected at random from every song, we then created a Gaussian mixture model of all of the music. This GMM is a generative model for all of the music and assumes that every time frame is independent of every of other frame and that they are all sampled from the same distribution, an assumption that almost certainly is not true, but is still the state of the art.

Jaakkola and Haussler describe the Fisher kernel as a method for summarizing the influence of the parameters of a generative model on a collection of samples from that model. In this case, the parameters we consider are the means of the Gaussians in the GMM, reducing the dimensionality of the data to the partial derivatives of the log posterior probabilities of each Gaussian with respect to the data. From [5],

$$\nabla_{\mu_k} \log(P(X|\mu_k)) = \sum_{t=1}^{m} P(k|\mathbf{x}_t)\Sigma_k^{-1}(\mathbf{x}_t - \mu_k). \tag{1}$$

Where $P(k|\mathbf{x}_t)$ is the posterior probability of the $k$th Gaussian in the mixture, given data point $\mathbf{x}_t$, and $\mu_k$ and $\Sigma_k$ are the mean and variance of the $k$th Gaussian. This process then reduces arbitrarily sized songs to 260 dimensional vectors (20 means with 13 dimensions each).

Each song's feature vector contains one more dimension, for its estimated tempo.

## 3.1 Tempo Tracker

The tempo tracker used in this experiment is based on Eric Scheirer's tempo tracker, as described in [7], but has been simplified in a number of respects.

The tempo tracker first starts with MFCCs as a band-separated version of the signal. The envelope in each band is then calculated and all of the signals are downsampled to speed computation. The tracker detects onsets in each channel and then combines the channels by adding all of

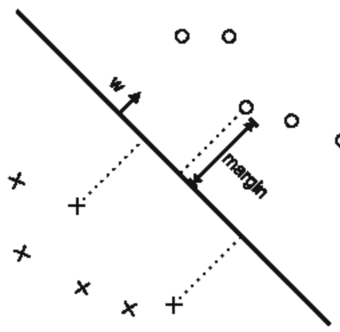


Figure 1: Linear Support Vector Machine [8].

their onsets together. The resulting signal is a series of pulses in time, hopefully corresponding to beats in the song. To translate this pulse train into a number, its DFT is taken and the largest peak within a reasonable range of tempos (50-400 beats per minute) is chosen as the dominant tempo.

Experiments with a similar system gave rather noisy results. The system's estimate was correct 30% of the time, it was incorrect 30% of the time, and it was off of the correct tempo by a factor of 2 or 1/2 40% of the time. This last case is a rather unfortunate failure mode when the tempo estimate is used to determine if a song is generally fast or slow. We plugged it in to the SVM anyway, assuming that it would be ignored if uninformative.

# 4 Support Vector Machines

A number of classification instances exist in which the task is so complicated a programmer cannot specify the output of a system based on the input data. In these situations, specific algorithms may be employed to learn a mapping from the input features to output classes similarly to how humans learn to perform categorization. The support vector machine (SVM) is a supervised learning system that uses a hypothesis space of linear functions in a high dimensional feature space in order to learn separating hyperplanes. As such, SVM classification attempts to generalize an optimal decision boundary between classes based on low-level representations of data.

Binary classification, such as whether or not a song is in a particular mood, of linearly separable data is the simplest implementation example of an SVM. As displayed in Figure 1, labeled training data in a given space may be separated by an maximal margin hyperplane through SVM classification. One side of the decision boundary consists of the positive instances, and the other side of the decision boundary consists of the negative instances. The training data is projected into a higher dimensional feature space via a Mercer kernel operator K [8]. For nonlinearly separable data, the original feature space may be mapped to a higher-dimensional, separable space by choosing a different kernel function.

In our implementation, we selected a Gaussian kernel. This means the hypotheses are radial basis function networks as

$$K(\mathbf{u}, \mathbf{v}) = e^{-\gamma(\mathbf{u}-\mathbf{v})\cdot(\mathbf{u}-\mathbf{v})}, \tag{2}$$

which creates boundaries by placing weighted Gaussians on key training instances[2].

# 5 Active Learning

In an active learning system, the user becomes an integral part of the learning and classification process. As opposed to passive SVM classification where a classifier is trained on a set of labeled examples and tested on another set of labeled examples, in an active learning system the user is presented with a set of unlabeled instances to label from the pool of all songs. Classification is performed based on the feedback of the user, and relevant responses are determined based on the individual users preferences and interpretations.

The active learning system we developed consisted of the following components: the binary classifier function, a query function, and the set of labeled and unlabeled instances. The classifier was trained on the set of labeled examples using a RBF kernel as described in the preceding section. The query function component selects the next set of instances to be labeled by the user. In order to reduce the set of hypotheses as quickly as possible, the user is presented with the instances that approximately halve the hypothesis space at each round of feedback [8]. This approximation was accomplished by choosing the previously unseen instances closest to the hyperplane decision boundary. Better approximations have been reported; however, the query function we used was selected for its speed and ease of implementation. The most relevant songs as determined by the SVM are classified by returning the positively labeled or predicted instances farthest from the decision boundary.

This active learning implementation has a number of advantages over traditional SVM classification. Primarily, by presenting the user with and training on the most relevant data of the set, the algorithm could potentially converge to the desired mapping function twice as fast as a passive SVM system. Secondly, by allowing the user to dynamically label the set of instances, any number of different classification retrieval problems may be performed using the same set of features and classifier. For example, the system may be used to query a specific AMG mood, jazz music, or female artists. As a result, the active SVM system efficiently maps the range of high-level descriptors people use to describe their perception of music to the low-level features with which we can represent music digitally.

## 5.1 Algorithm Summary

1. The user seeds the system with a representative song(s).

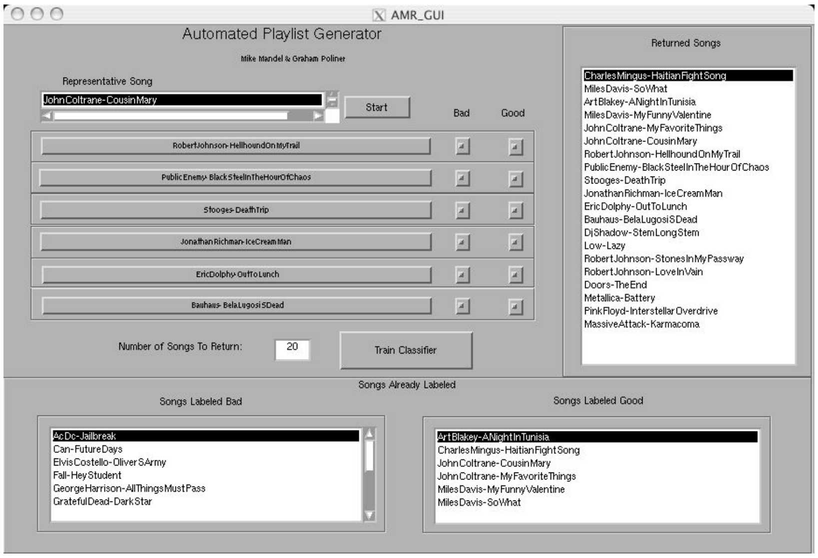


Figure 2: Active Learning Graphical User Interface.

2. If this is the first feedback round, the user is presented with randomly selected songs to label (in order to maintain stability). Else, the user is presented with the songs returned by the query function (instances closest to the decision boundary).
3. Learn an SVM on the labeled training instances.
4. Return to the user the most relevant songs (positive instances farthest from the decision boundary).
5. Repeat steps 2 to 4 iteratively until the user is satisfied with the results or a predetermined number of iterations has been reached.

# 6 User Interface

The graphical user interface is displayed in Figure 2. As described in the algorithm summary, the user selects a representative seed song and begins the active retrieval system. The user is then presented with six songs to label as similar to the representative song or dissimilar according to the user's desired heuristic (no decision is allowed). Next, the user selects the number of songs to return and begins the classification process. The songs the user has labeled are displayed at the bottom of the interface, and the songs returned by the classifier are displayed in the list on the right hand side of the interface. At any time, the user may click on any song displayed in the interface to hear a representative segment. After each classification round, the user is presented with six new songs to label and may perform the process iteratively as many times as desired.

# 7 Experiments

The ability of the features to discriminate between the AMG moods was evaluated using SVM classification with a RBF kernel. We performed empirical classification evaluations by randomly selecting two AMG moods and training a binary classifier. 100 random pairs were selected, and for each pair 100 random folds of cross validation were performed on the data with training size of 66% of the data. The evaluation was repeated with 100 permutations of 5 randomly selected mood sets and with all of the potential AMG moods.

A simulated user was developed in order to compare the active retrieval SVM algorithm with passive SVM classification. The simulated user program seeds the active retrieval system with a representative song(s) from a given AMG mood as a query concept. Next, the simulated user performs relevance feedback using the corresponding AMG mood set for relevance decisions. The returned results of the active learner were recorded after each iteration in order to generate statistical data for comparison with the passive SVM classifier. The ground truth for the evaluation was defined as the representative set of the AMG mood.

A few audio processing researchers were encouraged to evaluate the performance of the active retrieval system based on the ability of the system to generate relevant response to their queries. The evaluators searched for descriptors such as jazz, rap, rock, punk, female artists, beat, etc.

# 8 Results

For over half of the moods the active SVM learner outperformed a regular SVM trained on the same number of examples. See Table 1 for a list of the moods and the number of training examples saved using an active SVM instead of a regular SVM. Only in a few cases did the regular SVM outperform the active SVM completely, although sometimes one classifier performed better on high precision, low recall tasks while the other outperformed it on low precision, high recall tasks.

## 8.1 Subsystem Testing

After obtaining these results we wanted to determine the causes of error in the active SVM system by testing its various parts separately. To determine if the active SVM system worked, we tested it with a 2-Dimensional, linearly separable toy data set, for which it was easily able to increase its performance after each round of feedback.

To determine if the features provided enough information to separate different moods from one another, we ran a number of regular SVM multiclass cross-validation tests. For many combinations of two-class tests, the percent of songs correctly classified was on average 53.1%. For five-class tests, 21.9% of songs were correctly classified. These numbers are very close to classifying songs randomly, thus leading to the conclusion that

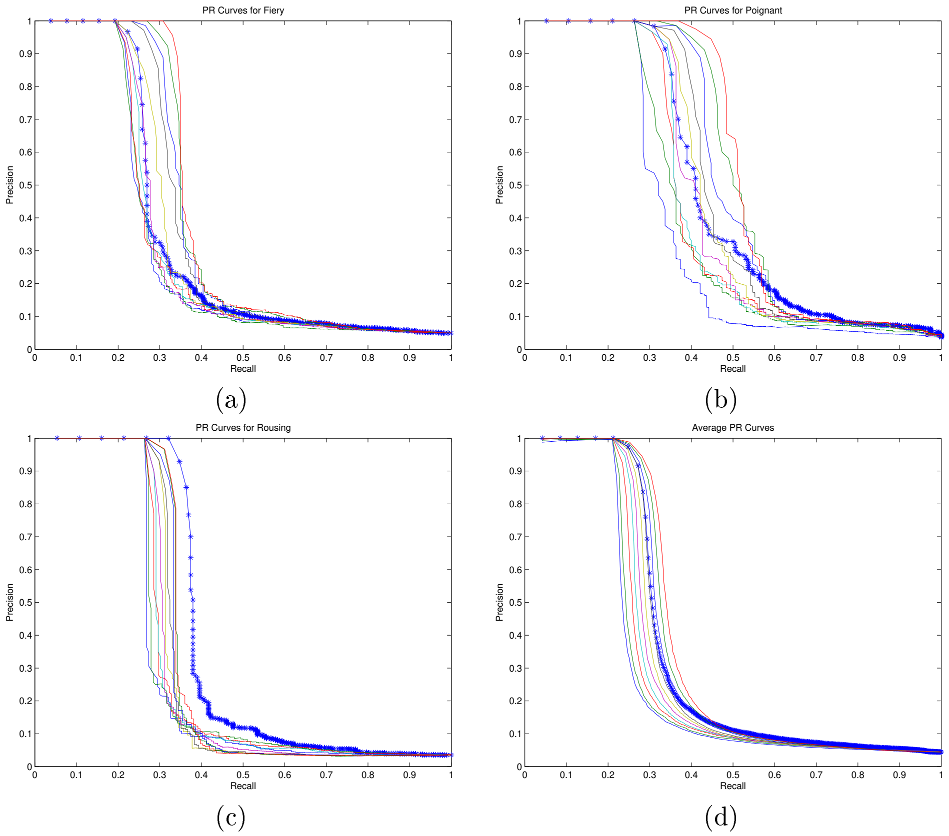


Figure 3: Precision-Recall curves for particularly good moods: (a) Fiery, (b) Poignant; a particularly bad mood: (c) Rousing; and (d) across all moods. The starred line is the performance of the normal SVM, the other lines (from left to right) are for each training iteration of the active SVM.

our features probably do not provide enough information to distinguish moods from one another.

# 9 Conclusions

We have demonstrated that active learning with support vector machines could be useful for content-based retrieval of audio in databases. In the best case, active learning requires fewer than half as many training examples for equivalent results to passive SVM classification. In limited user trials, subjects reported positive results in searching for semantic descriptors and overall satisfaction with the user interface. However, the poor performance results of the passive SVM classifiers suggests a limited correlation between the features and the AMG moods. Therefore, the short-time features may not be effective for classification of the AMG moods. In addition, the ground truth may not accurately represent the output classes. Since the labels only depict the top 50 songs in a given mood, the absence of a song from the ground truth does not necessarily preclude it from being in the class. However, the results of the active retrieval system on representative features and the success of the system in user trials show promise for future research in content-based audio retrieval.

## 9.1 Future Work

We hope to pursue a number of enhancements to the active-retrieval system. New features such as the pitch histogram, melody contour, and average posterior probability of GMMs could potentially yield better classification results and enable more diverse queries. As noted in [8], we could also explore different methods for choosing query points and relevant examples. Better ground truth data could be bootstrapped from a knowledge of the moods of albums and the audio similarity of the songs on those albums.

# References

[1] http://www.allmusic.com/cg/amg.dll?p=amg&sql=32:amg/info_pages/a_siteglossary.html

[2] N. Cristianini, J. Shawe-Taylor. *An Introduction to Support Vector Machines*, Cambridge University Press 2000.

[3] T. Jaakkola, D. Haussler, "Exploiting Generative Models in Discriminative Classifiers." In Advances in Neural Information Processing Systems 11, 1998.

[4] B. Logan, "Mel Frequency Cepstral Coefficients for Music Modelling," in Proc. Int. Symp. Music Information Retrieval (ISMIR), 2000.

[5] P. Moreno, P. Ho, N. Vasconcelos, "A Kullback-Leibler Divergence Based Kernel for SVM Classification in Multimedia Applications." Proc. of Neural Information Processing Systems, Vancouver Canada, 2003.

[6] Y. Rui, T. Huang, M. Ortega, and S. Mehrotra, "Relevance Feedback: A Power Tool for Interactive Content-Based Image Retrieval." IEEE Trans. on Circuits and Systems for Video Technology, 8(5):644-655, Sep. 1998.

[7] E. Scheirer. Tempo and beat analysis of acoustic musical signals. Journal of the Acoustical Society of America, 103(1):588-601, 1998.

[8] S. Tong and E. Chang, "Support Vector Machine Active Learning for Image Retrieval," ACM International Conference on Multimedia, pp. 107-118, Ottawa, October 2001.

Table 1: Results for all of the moods included in our experiment. Num Songs is the total number of songs labelled with that mood, Training Savings is the number of examples fewer than 50 giving the same results as a regular SVM trained with 50 examples. Dashes indicate that the active SVM performed worse than the regular SVM when trained on 50 examples. The final row is not calculated directly from the other numbers in the table, but from Figure 3.

| Mood Name | Num Songs | Training Savings |
|---|---|---|
| Acerbic | 25 | 35 |
| Aggressive | 30 | 20 |
| Angry | 26 | 0 |
| Angst-Ridden | 30 | 35 |
| Bleak | 20 | 0 |
| Boisterous | 19 | 25 |
| Brooding | 21 | 25 |
| Cathartic | 31 | – |
| Confident | 22 | 25 |
| Detached | 26 | 20 |
| Druggy | 22 | – |
| Eerie | 24 | – |
| Energetic | 27 | 20 |
| Fiery | 26 | 25 |
| Gloomy | 22 | – |
| Intimate | 12 | 35 |
| Ironic | 28 | 0 |
| Irreverent | 23 | 30 |
| Malevolent | 24 | 10 |
| Ominous | 34 | 15 |
| Party/Celebratory | 19 | 25 |
| Poignant | 19 | 30 |
| Quirky | 23 | 10 |
| Rebellious | 34 | 15 |
| Reflective | 20 | 30 |
| Rousing | 19 | – |
| Sexy | 23 | 5 |
| Stylish | 21 | 0 |
| Theatrical | 23 | – |
| Visceral | 23 | 0 |
| Volatile | 21 | 0 |
| Witty | 21 | 0 |
| **Average** | **23.7** | **20** |

Table 2: List of all of the songs used in this project, sorted alphabetically by artist. The artist and song names have been "canonicalized" to remove punction and spaces and changed to Title Case.

| Artists | Songs |
|---|---|
| Abba | DancingQueen, Waterloo |
| AcDc | BackInBlack, DirtyDeedsDoneDirtCheap, HighwayToHell, Jailbreak |
| Aerosmith | BackInTheSaddle, DreamOn, WalkThisWay |
| AfghanWhigs | Debonair, Gentlemen |
| Air | SexyBoy |
| AlGreen | CallMeComeBackHome, HereIAmComeAndTakeMe, IMStillInLoveWithYou, LetSStayTogether |
| AlanisMorissette | AllIReallyWant, YouOughtaKnow |
| AliceInChains | GotMeWrong, IStayAway, ManInTheBox, Rooster, Would, Rise |
| Anthrax | GotTheTime |
| ArtBlakey | ANightInTunisia, Moanin |
| AtDriveIn | Quarantined |
| Autechre | Eutow |
| Auteurs | JunkShopClothes, ShowGirl |
| BadBrains | PayToCum |
| BadCompany | FeelLikeMakinLove |
| BarryWhite | CanTGetEnoughOfYourLoveBabe |
| Bauhaus | BelaLugosiSDead |
| Bauhaus | InTheFlatField, SheSInParties |
| BeastieBoys | FightForYourRight, HeyLadies |
| Beatles | AndILoveHer, CanTBuyMeLove, DriveMyCar, GotToGetYouIntoMyLife, HelterSkelter, HereThereAndEverywhere, IAmTheWalrus, IMDown, ISawHerStandingThere, IWantToHoldYourHand, InMyLife, ItWonTBeLong, LetItBe, MarthaMyDear, PaperbackWriter, SheLovesYou, SheSaidSheSaid, StrawberryFieldsForever |
| Beck | WhereItSAt |
| BeeGees | HowCanYouMendABrokenHeart, IVeGottaGetAMessageToYou, NewYorkMiningDisaster1941 |
| BelleAndSebastian | ItCouldHaveBeenABrilliantCareer, TheStateIAmIn |
| BillyIdol | EyesWithoutAFace, RebelYell |
| BirthdayParty | DeadJoe, NickTheStripper |
| BlackSabbath | BlackSabbath, IronMan, Paranoid, SweetLeaf, WarPigs |
| BlindFaith | CanTFindMyWayHome |
| Blondie | HangingOnTheTelephone, HeartOfGlass, InTheFlesh, OneWayOrAnother, RipHerToShreds |
| BluesBrothers | OpeningICanTTurnYouLoose |
| Blur | Beetlebum, CountryHouse, GirlsAndBoys, PopScene, Song2, ThisIsALow |
| BobDylan | BalladOfAThinMan, Highway61Revisited, IThrewItAllAway, IfYouSeeHerSayHello, LikeARollingStone, MillionDollarBash, MrTambourineMan, SimpleTwistOfFate, SubterraneanHomesickBlues |
| BobMarley | GetUpStandUp |
| BoogieDownProductions | YouMustLearn |
| Breeders | Cannonball |
| BrianEno | BabySOnFire, ILlComeRunning |
| BruceSpringsteen | BornInTheUSA, BornToRun |
| BurtBacharach | TheyLongToBeCloseToYou, WalkOnBy |
| Bush | EverythingZen |
| Buzzcocks | EverFallenInLove, WhatDoIGet |
| CamperVanBeethoven | TakeTheSkinheadsBowling, WhereTheHellIsBill |
| Can | FutureDays, Mushroom |
| CannonballAdderley | JiveSamba |
| Cars | GoodTimesRoll, JustWhatINeeded |
| CatStevens | MorningHasBroken |
| CharlesMingus | HaitianFightSong |
| CheapTrick | IWantYouToWantMe, Surrender |
| ChemicalBrothers | BlockRockinBeats |
| ChristinaAguilera | GenieInABottle |
| ChuckBerry | NoParticularPlaceToGo, ReelinAndRockin |
| Clash | CompleteControl, Garageland, IMSoBoredWithTheUSA, JanieJones, LondonCalling, LondonSBurning, WhiteManInHammersmithPalais, WhiteRiot |
| CocteauTwins | CarolynSFingers |
| CreedenceClearwaterRevival | FortunateSon, UpAroundTheBend |
| CrosbyStillsAndNash | LongTimeGone |
| CrowdedHouse | DonTDreamItSOver |
| Cure | FascinationStreet, JustLikeHeaven, Pornography |
| CyndiLauper | GirlsJustWantToHaveFun |
| CypressHill | InsaneInTheBrain |
| DaftPunk | DaFunk, OneMoreTime |
| Damned | NewRose |
| DavidBowie | AlwaysCrashingInTheSameCar, DiamondDogs, DriveInSaturday, GoldenYears, JohnIMOnlyDancing, LifeOnMars, MoonageDaydream, RebelRebel, SpaceOddity, TheManWhoSoldTheWorld, YoungAmericans, ZiggyStardust |
| DeadKennedys | HolidayInCambodia, KillThePoor |
| DeadMilkmen | PunkRockGirl |
| DeeeLite | GrooveIsInTheHeart |
| DefLeppard | PourSomeSugarOnMe |

Table 2: *Continued*

| Artists | Songs |
|---|---|
| DepecheMode | EnjoyTheSilence, EverythingCounts, NeverLetMeDownAgain |
| Devo | GirlUWant, JockoHomo, Mongoloid, WhipIt |
| DigitalUnderground | TheHumptyDance |
| DivineComedy | BecomingMoreLikeAlfie |
| DjShadow | MidnightInAPerfectWorld, StemLongStem |
| Dmx | GetAtMeDog |
| DoobieBrothers | RockinDownTheHighway |
| Doors | BreakOnThroughToTheOtherSide, FiveToOne, LAWoman, PeopleAreStrange, RidersOnTheStorm, RoadhouseBlues, TheEnd |
| DrDre | NuthinButAGThang |
| DrOctagon | BlueFlowers |
| DukeEllington | Caravan |
| DuranDuran | GirlsOnFilm, Rio |
| Elastica | 21, Connection |
| ElmoreJames | TheSkyIsCrying |
| EltonJohn | RocketManIThinkItSGoingToBeALongLongTime, YourSong |
| ElvisCostello | Alison, LipService, NoAction, OliverSArmy, PumpItUp, RadioRadio |
| Eminem | MyNameIs |
| Enya | OrinocoFlow |
| EricDolphy | OutToLunch |
| Eurythmics | HereComesTheRainAgain, SweetDreamsAreMadeOfThis |
| EverlyBrothers | AllIHaveToDoIsDream |
| FaithNoMore | Epic |
| Fall | HeyStudent |
| FatboySlim | TheRockafellerSkank |
| FlamingLips | PilotCanAtTheQueerOfGod, SheDonTUseJelly |
| FleetwoodMac | Landslide |
| FrankSinatra | FlyMeToTheMoon, InTheWeeSmallHoursOfTheMorning, NightAndDay |
| FrankZappa | DonTEatTheYellowSnow, MyGuitarWantsToKillYourMama, ValleyGirl |
| Fugazi | MarginWalker |
| Garbage | StupidGirl |
| Genesis | InTooDeep |
| GeorgeHarrison | AllThingsMustPass |
| GeorgeMichael | FatherFigure |
| GetUpKids | RedLetterDay |
| GratefulDead | DarkStar |
| GreenDay | BasketCase |
| GunClub | ForTheLoveOfIvy, SexBeat |
| GunsNRoses | Nightrain, ParadiseCity, WelcomeToTheJungle |
| Hole | Violet |
| HowlinWolf | SmokestackLightning, Spoonful |
| HumanLeague | DonTYouWantMe, KeepFeelingFascination |
| IceCube | CheckYoSelf |
| IceT | NewJackHustler |
| IggyPop | LustForLife |
| Inxs | WhatYouNeed |
| IronButterfly | InAGaddaDaVida |
| JamesBrown | GetUpIFeelLikeBeingASexMachine, IGotYouIFeelGood, PleasePleasePlease |
| JamesTaylor | FireAndRain |
| JanetJackson | MissYouMuch |
| JayZ | BigPimpin |
| JeffBuckley | LastGoodbye |
| JeffersonAirplane | WhiteRabbit |
| JerryLeeLewis | Breathless, HighSchoolConfidential, WholeLottaShakinGoinOn |
| JesusAndMaryChain | BluesFromAGun |
| JethroTull | Aqualung |
| JimiHendrix | CastlesMadeOfSand, Fire, PurpleHaze |
| JimmyEatWorld | BleedAmerican |
| JohnColtrane | Alabama, CousinMary, GiantSteps, MyFavoriteThings |
| JohnnyCash | FolsomPrisonBlues, ISeeADarkness, IStillMissSomeone |
| JonathanRichman | IceCreamMan |
| JoyDivision | Atmosphere, LoveWillTearUsApart, SheSLostControl |
| Kansas | CarryOnWaywardSon |
| KidRock | Bawitdaba, Cowboy |
| Kinks | AWellRespectedMan, DedicatedFollowerOfFashion, SunnyAfternoon, WhereHaveAllTheGoodTimesGone, YouReallyGotMe |
| Korn | FreakOnALeash |
| LedZeppelin | BabeIMGonnaLeaveYou, CommunicationBreakdown, DazedAndConfused, Heartbreaker, Kashmir, NoQuarter, WholeLottaLove |
| LennyKravitz | AreYouGonnaGoMyWay |
| LeonardCohen | EverybodyKnows |
| Live | LightningCrashes |
| LizPhair | FuckAndRun |
| LlCoolJ | MamaSaidKnockYouOut |
| LouReed | Vicious |
| Low | Lazy |
| MC5 | KickOutTheJams |
| Madonna | ExpressYourself, Holiday, IntoTheGroove, LikeAVirgin, Vogue |

Table 2: *Continued*

| Artists | Songs |
|---|---|
| MariahCarey | Fantasy |
| MarvinGaye | InnerCityBluesMakeMeWannaHoller, LetSGetItOn, SexualHealing, WhatSGoingOn |
| MassiveAttack | InertiaCreeps, Karmacoma, Risingson, UnfinishedSympathy |
| MenAtWork | WhoCanItBeNow |
| Metallica | Battery, CreepingDeath, EnterSandman, FadeToBlack, MasterOfPuppets, One |
| Meters | CissyStrut |
| MethodMan | BringThePain |
| MichaelJackson | Bad, DonTStopTilYouGetEnough, WannaBeStartinSomethin |
| MilesDavis | ItAinTNecessarilySo, MilesRunsTheVoodooDown, MyFunnyValentine, SoWhat |
| Ministry | JesusBuiltMyHotrod, Stigmata |
| MinorThreat | StraightEdge |
| Misfits | Attitude, LastCaress |
| MissionOfBurma | ThatSWhenIReachForMyRevolver |
| Mogwai | MogwaiFearSatan |
| Monkees | IMNotYourSteppinStone |
| MonsterMagnet | NegasonicTeenageWarhead |
| MoodyBlues | NightsInWhiteSatin |
| MoogCookbook | BlackHoleSun |
| Morrissey | CertainPeopleIKnow |
| MuddyWaters | MannishBoy |
| NeilYoung | AfterTheGoldRush, DownByTheRiver, SugarMountain, TheNeedleAndTheDamage-Done |
| NewOrder | BizarreLoveTriangle, BlueMonday |
| NineInchNails | Closer, HeadLikeAHole, Wish |
| Nirvana | AllApologies, Aneurysm, ComeAsYouAre, Dive, HeartShapedBox, InBloom, Smells-LikeTeenSpirit, SomethingInTheWay, TerritorialPissings |
| Nsync | ByeByeBye |
| Oasis | ChampagneSupernova, LiveForever |
| OrnetteColeman | Congeniality |
| PJHarvey | 50ftQueenie |
| Parliament | GiveUpTheFunkTearTheRoofOffTheSucker |
| PearlJam | EvenFlow, Jeremy |
| PetShopBoys | BeingBoring, Rent, Suburbia, WestEndGirls, WhereTheStreetsHaveNoNameICanT-TakeMyEyesOffYou |
| PeterGabriel | GamesWithoutFrontiers, HereComesTheFlood, ShockTheMonkey |
| PinkFloyd | AstronomyDomine, InterstellarOverdrive, Money, OneOfTheseDays, Time |
| Pixies | BoneMachine, Debaser, HereComesYourMan, MonkeyGoneToHeaven, SomethingA-gainstYou, Tame |
| PizzicatoFive | HappySad |
| Poison | TalkDirtyToMe |
| Pretenders | BrassInPocket |
| PrimalScream | MbvArkestraIfTheyMoveKillEm |
| Primus | JerryWasARaceCarDriver |
| Prince | Delirious, Kiss, LittleRedCorvette |
| Prodigy | Firestarter, SmackMyBitchUp |
| PsychedelicFurs | LoveMyWay, PrettyInPink |
| PublicEnemy | 911IsAJoke, BlackSteelInTheHourOfChaos, BringTheNoise, FightThePower, Party-ForYourRightToFight, RebelWithoutAPause, WelcomeToTheTerrordome |
| Pulp | Babies, CommonPeople, ISpy, Underwear |
| PussyGalore | DickJohnson |
| Queen | BohemianRhapsody, KillerQueen, WeWillRockYou, BetterLivingThroughChemistry |
| REM | EverybodyHurts, ItSTheEndOfTheWorldAsWeKnowItAndIFeelFine, TalkAbout-ThePassion |
| Radiohead | Creep, FakePlasticTrees, HighAndDry, KarmaPolice, NoSurprises, ParanoidAndroid |
| Rammstein | DuHast |
| Ramones | BlitzkriegBop, IWannaBeSedated, RockawayBeach, SheenaIsAPunkRocker |
| RedHotChiliPeppers | GiveItAway |
| Replacements | IHateMusic, IWillDare, LeftOfTheDial, WaitressInTheSky |
| RobertJohnson | HellhoundOnMyTrail, LastFairDealGoneDown, LoveInVain, StonesInMyPassway |
| RodStewart | MandolinWind |
| RollingStones | 19thNervousBreakdown, GetOffOfMyCloud, GimmeShelter, HeartOfStone, ICanT-GetNoSatisfaction, LetSSpendTheNightTogether, PaintItBlack, StreetFightingMan, UndercoverOfTheNight |
| RoyOrbison | Crying, InDreams, RunningScared |
| RufusWainwright | CigarettesAndChocolateMilk |
| RunDMC | KingOfRock |
| Sade | SmoothOperator |
| SaintEtienne | OnlyLoveCanBreakYourHeart |
| SexPistols | AnarchyInTheUK, GodSaveTheQueen |
| SimonAndGarfunkel | HomewardBound, TheSoundOfSilence |
| Slayer | SouthOfHeaven |
| SmashingPumpkins | Disarm |
| Smiths | HeavenKnowsIMMiserableNow, HowSoonIsNow, Panic, PleasePleasePleaseLet-MeGetWhatIWant, TheBoyWithTheThornInHisSide, ThereIsALightThatNeverGoe-sOut, ThisCharmingMan |
| SonicYouth | DeathValley69, ExpresswayToYrSkull, KoolThing, ShadowOfADoubt, TeenAgeRiot |
| SoulCoughing | DownToThis |
| Soundgarden | BlackHoleSun, JesusChristPose, Outshined, RustyCage |

Table 2: *Continued*

| Artists | Songs |
|---|---|
| Specials | NiteKlub |
| Spiritualized | LadiesAndGentlemenWeAreFloatingInSpace |
| Stereolab | CybeleSReverie, LoBoobOscillator |
| StevieWonder | MyCherieAmour |
| StoneRoses | IAmTheResurrection, IWannaBeAdored |
| Stooges | DeathTrip, FunHouse, IWannaBeYourDog, LABlues, Loose, NoFun, RawPower, SearchAndDestroy, TVEye |
| Suede | BeautifulOnes, SleepingPills, SoYoung, TheDrowners |
| Sugarcubes | Birthday |
| Supergrass | CaughtByTheFuzz |
| Supertramp | BloodyWellRight |
| SuzanneVega | MarleneOnTheWall |
| TalkingHeads | PsychoKiller |
| TeenageFanclub | Alcoholiday |
| Television | MarqueeMoon |
| TheyMightBeGiants | BirdhouseInYourSoul |
| Tindersticks | TinyTears |
| TomPetty | AmericanGirl |
| TomWaits | ClapHands, HangOnStChristopher, Ol55 |
| ToriAmos | Crucify, SilentAllTheseYears |
| Tricky | Aftermath, BlackSteel, Christiansands, HellIsAroundTheCorner, Overcome, Ponderosa |
| U2 | PrideInTheNameOfLove |
| Unkle | RabbitInYourHeadlights |
| UrgeOverkill | SisterHavana |
| VanHalen | AinTTalkinBoutLove, Panama |
| VelvetUnderground | Heroin, VenusInFurs |
| VerucaSalt | Seether, OnYourOwn |
| ViolentFemmes | AddItUp, BlisterInTheSun, GoneDaddyGone, KissOff |
| Ween | ICanTPutMyFingerOnIt, PushThLittleDaisies |
| WhiteStripes | FellInLoveWithAGirl, HotelYorba |
| Who | AQuickOneWhileHeSAway, AnywayAnyhowAnywhere, BabaORiley, BorisTheSpider, MyGeneration, Substitute, SummertimeBlues, WonTGetFooledAgain |
| Wire | IAmTheFly, ThreeGirlRhumba |