



HAL
open science

PAC-Bayesian Learning and Domain Adaptation

Pascal Germain, Amaury Habrard, François Laviolette, Emilie Morvant

► **To cite this version:**

Pascal Germain, Amaury Habrard, François Laviolette, Emilie Morvant. PAC-Bayesian Learning and Domain Adaptation. Multi-Trade-offs in Machine Learning, NIPS 2012 Workshop, Dec 2012, Lake Tahoe, United States. hal-00749366

HAL Id: hal-00749366

<https://hal.science/hal-00749366>

Submitted on 10 Dec 2012

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

PAC-Bayesian Learning and Domain Adaptation

Pascal Germain

Département d'informatique et de génie logiciel
Université Laval, Québec, Canada
pascal.germain@ift.ulaval.ca

Amaury Habrard

Laboratoire Hubert Curien UMR CNRS 5516,
Univ. Jean Monnet, 42000 St-Etienne, France
amaury.habrard@univ-st-etienne.fr

François Laviolette

Département d'informatique et de génie logiciel
Université Laval, Québec, Canada
francois.laviolette@ift.ulaval.ca

Emilie Morvant

Aix-Marseille Univ., LIF-QARMA, CNRS,
UMR 7279, 13013, Marseille, France
emilie.morvant@lif.univ-mrs.fr

Abstract

In machine learning, *Domain Adaptation* (DA) arises when the distribution generating the test (target) data differs from the one generating the learning (source) data. It is well known that DA is an hard task even under strong assumptions [1], among which the *covariate-shift* where the source and target distributions diverge only in their marginals, *i.e.* they have the same labeling function. Another popular approach is to consider an hypothesis class that moves closer the two distributions while implying a low-error for both tasks [2]. This is a *VC-dim* approach that restricts the complexity of an hypothesis class in order to get good generalization. Instead, we propose a *PAC-Bayesian* approach that seeks for suitable weights to be given to each hypothesis in order to build a majority vote. We prove a new DA bound in the PAC-Bayesian context. This leads us to design the first DA-PAC-Bayesian algorithm based on the minimization of the proposed bound. Doing so, we seek for a ρ -weighted majority vote that takes into account a trade-off between three quantities. The first two quantities being, as usual in the PAC-Bayesian approach, (a) the complexity of the majority vote (measured by a Kullback-Leibler divergence) and (b) its empirical risk (measured by the ρ -average errors on the source sample). The third quantity is (c) the capacity of the majority vote to distinguish some structural difference between the source and target samples.

Preliminaries

Domain Adaptation. We consider DA for binary classification tasks where $X \subseteq \mathbb{R}^d$ is the input space of dimension d and $Y = \{-1, 1\}$ is the label set. We have two different distributions over $X \times Y$ called the *source domain* P_S and the *target domain* P_T . D_S and D_T are the respective marginal distributions over X . We tackle the challenging task where we have no information about the label on P_T . A learning algorithm is then provided with a *labeled source sample* $S = \{(\mathbf{x}_i^s, y_i^s)\}_{i=1}^m$ drawn *i.i.d.* from P_S , and an *unlabeled target sample* $T = \{\mathbf{x}_j^t\}_{j=1}^{m'}$ drawn *i.i.d.* from D_T . Let $h: X \rightarrow Y$ be an hypothesis function. The *expected source error* of h over P_S is the probability that h commits an error, $R_{P_S}(h) = \mathbf{E}_{(\mathbf{x}^s, y^s) \sim P_S} I(h(\mathbf{x}^s) \neq y^s)$, where $I(a) = 1$ if predicate a is true and 0 otherwise. The *expected target error* R_{P_T} over P_T is defined in a similar way. R_S is the *empirical source error*.

The DA objective is then to find a low error target hypothesis, even if no label information is available about the target domain. Clearly this task can be infeasible in general. However, under the assumption that there exists hypothesis in the hypothesis class \mathcal{H} that do perform well on both the source and the target domain, Ben David et al. [2] provide the following guarantee,

$$\forall h \in \mathcal{H}, R_{P_T}(h) \leq R_{P_S}(h) + \frac{1}{2} d_{\mathcal{H}\Delta\mathcal{H}}(D_S, D_T) + \nu, \quad (1)$$

where $\nu \stackrel{\text{def}}{=} \operatorname{argmin}_{h \in \mathcal{H}} (R_{P_S}(h) + R_{P_T}(h))$ is the error of the best joint hypothesis, and $d_{\mathcal{H}\Delta\mathcal{H}}(D_S, D_T)$, called the $\mathcal{H}\Delta\mathcal{H}$ -distance between the domain marginal distributions, quantifies how hypothesis from \mathcal{H} can “detect” differences between those two distributions. According to Equation (1), the lower this detection capability is for some given \mathcal{H} , the better are the generalization guarantees. Hence, as pointed out in [2], Equation (1) together with the usual VC-bound theory, express a multiple trade-off between the accuracy of some particular hypothesis h , the complexity of the hypothesis class \mathcal{H} , and the “incapacity” of hypothesis of \mathcal{H} to detect difference between the source and the target domain.

PAC-Bayesian Learning of Linear Classifier. The PAC-Bayesian theory, first introduced by McAllester [3], traditionally considers majority votes over a set \mathcal{H} of binary hypothesis. Given a prior distribution π over \mathcal{H} and a training set S , the learning process consists in finding the posterior distribution ρ over \mathcal{H} leading to a good generalization. Indeed, the essence of this theory is to bound the risk of the stochastic Gibbs classifier G_ρ associated with ρ . In order to predict the label of an example \mathbf{x} , the Gibbs classifier first draws a hypothesis h from \mathcal{H} according to ρ , then returns $h(\mathbf{x})$ as the predicted label. Note that the error of the Gibbs classifier corresponds to the expectation of the errors over ρ : $R_{P_S}(G_\rho) = \mathbf{E}_{h \sim \rho} R_{P_S}(h)$. The classical PAC-Bayesian theorem bounds the expectation of error $R_{P_S}(G_\rho)$ in term of two major quantities: The empirical error $R_S(G_\rho) = \mathbf{E}_{h \sim \rho} R_S(h)$ on a sample S and the Kullback-Leibler divergence $\text{KL}(\rho \parallel \pi) \stackrel{\text{def}}{=} \mathbf{E}_{h \sim \rho} \frac{\rho(h)}{\pi(h)}$.

Theorem 1 (as presented in [4]). *For any domain $P_S \subseteq X \times Y$, for any set \mathcal{H} of hypothesis, for any prior distribution π over \mathcal{H} , and any $\delta \in (0, 1]$, we have,*

$$\Pr_{S \sim (P_S)^m} \left(\forall \rho \text{ on } \mathcal{H} : \text{kl} \left(R_S(G_\rho) \parallel R_{P_S}(G_\rho) \right) \leq \frac{1}{m} \left[\text{KL}(\rho \parallel \pi) + \ln \frac{\xi(m)}{\delta} \right] \right) \geq 1 - \delta,$$

where $\text{kl}(q \parallel p) \stackrel{\text{def}}{=} q \ln \frac{q}{p} + (1 - q) \ln \frac{1 - q}{1 - p}$, and $\xi(m) \stackrel{\text{def}}{=} \sum_{k=0}^m \binom{m}{k} \left(\frac{k}{m}\right)^k \left(1 - \frac{k}{m}\right)^{m-k}$.

Now, let \mathcal{H} be a set of linear classifiers $h_{\mathbf{v}}(\mathbf{x}) \stackrel{\text{def}}{=} \text{sgn}(\mathbf{v} \cdot \mathbf{x})$ such that $\mathbf{v} \in \mathbb{R}^d$ is a weight vector. By restricting the prior and the posterior to be Gaussian distributions, Langford and Shawe-Taylor [5] have specialized the PAC-Bayesian theory in order to bound the expected risk of any linear classifier $h_{\mathbf{w}} \in \mathcal{H}$ identified by a weight vector \mathbf{w} . More precisely, for a prior $\pi_{\mathbf{0}}$ and a posterior $\rho_{\mathbf{w}}$ defined as spherical Gaussians with identity covariance matrix respectively centered on vectors $\mathbf{0}$ and \mathbf{w} , *i.e.*

$$\text{for any } h_{\mathbf{v}} \in \mathcal{H}, \quad \pi_{\mathbf{0}}(h_{\mathbf{v}}) \stackrel{\text{def}}{=} \left(\frac{1}{\sqrt{2\pi}}\right)^d e^{-\frac{1}{2}\|\mathbf{v}\|^2} \quad \text{and} \quad \rho_{\mathbf{w}}(h_{\mathbf{v}}) \stackrel{\text{def}}{=} \left(\frac{1}{\sqrt{2\pi}}\right)^d e^{-\frac{1}{2}\|\mathbf{v} - \mathbf{w}\|^2}, \quad (2)$$

we obtain that the expected risk of the Gibbs classifier $G_{\rho_{\mathbf{w}}}$ on a domain P_S is given by,

$$R_{P_S}(G_{\rho_{\mathbf{w}}}) = \mathbf{E}_{(\mathbf{x}, y) \sim P_S} \mathbf{E}_{h_{\mathbf{v}} \sim \rho_{\mathbf{w}}} I(h_{\mathbf{v}} \neq y) = \mathbf{E}_{(\mathbf{x}, y) \sim P_S} \Phi \left(y \frac{\mathbf{w} \cdot \mathbf{x}}{\|\mathbf{x}\|} \right),$$

where $\Phi(a) \stackrel{\text{def}}{=} \frac{1}{2} [1 - \text{Erf}(\frac{a}{\sqrt{2}})]$. Moreover, the KL-divergence between the posterior and the prior distributions becomes simply $\text{KL}(\rho_{\mathbf{w}} \parallel \pi_{\mathbf{0}}) = \frac{1}{2} \|\mathbf{w}\|^2$. In this context, Theorem 1 becomes,

Corollary 1. *For any domain $P_S \subseteq \mathbb{R}^d \times Y$ and any $\delta \in (0, 1]$, we have,*

$$\Pr_{S \sim (P_S)^m} \left(\forall \mathbf{w} \in \mathbb{R}^d : \text{kl} \left(R_S(G_{\rho_{\mathbf{w}}}) \parallel R_{P_S}(G_{\rho_{\mathbf{w}}}) \right) \leq \frac{1}{m} \left[\frac{1}{2} \|\mathbf{w}\|^2 + \ln \frac{\xi(m)}{\delta} \right] \right) \geq 1 - \delta.$$

Based on this specialization of the PAC-Bayesian theory to linear classifiers, Germain et al. [4] suggested to minimize the bound on $R_{P_S}(G_{\rho_{\mathbf{w}}})$ given by Corollary 1. The resulting learning algorithm, called PBGD, performs a gradient descent in order to find an optimal weight vector \mathbf{w} . Doing so, PBGD realizes a trade-off between the empirical accuracy (expressed by $R_S(G_{\rho_{\mathbf{w}}})$) and the complexity (expressed by $\|\mathbf{w}\|^2$) of the learned linear classifier.

PAC-Bayesian Learning of Adapted Linear Classifier

DA Bound for the Gibbs Classifier. The originality of our contribution is to combine PAC-Bayesian and DA frameworks. We define the notion of *domain disagreement* $\text{dis}_\rho(D_S, D_T)$ to measure the structural difference between domain marginals in terms of posterior distribution $\rho \sim \mathcal{H}$,

$$\text{dis}_\rho(D_S, D_T) \stackrel{\text{def}}{=} \mathbf{E}_{h_1, h_2 \sim \rho^2} [R_{D_T}(h_1, h_2) - R_{D_S}(h_1, h_2)],$$

where $R_{D'}(h_1, h_2) \stackrel{\text{def}}{=} \mathbf{E}_{\mathbf{x} \sim D'} I(h_1(\mathbf{x}) \neq h_2(\mathbf{x}))$. Unlike the distance $d_{\mathcal{H}\Delta\mathcal{H}}$ suggested by [2], our “distance” measure dis_ρ takes into account a ρ -average over all pairs of hypothesis in \mathcal{H} instead of focusing on a single particular pair of hypothesis. However, it nevertheless allows us to derive the following bound which proposes a similar trade-off as in Equation (1), but relates the source and target errors of the Gibbs classifier. For all probability distribution ρ on \mathcal{H} , we have,

$$R_{P_T}(G_\rho) \leq R_{P_S}(G_\rho) + \text{dis}_\rho(D_S, D_T) + \lambda_\rho, \quad (3)$$

where $\lambda_\rho = R_{P_T}(h^*) + R_{P_S}(h^*)$, with $h^* = \text{argmin}_{h \in \mathcal{H}} \{\mathbf{E}_{h' \sim \rho} (R_{D_T}(h, h') - R_{D_S}(h, h'))\}$, measures the joint error of the hypothesis which minimizes the domain disagreement. Hence, similarly to Equation (1), we provide evidences that a good DA is possible if $\text{dis}_\rho(D_S, D_T)$ and λ_ρ are low. Under this assumption, we propose to design the first DA-PAC-Bayesian algorithm inspired from the PAC-Bayesian learning of linear classifiers [4]. We focus on the two first terms of Inequality (3), and we refer to this quantity as the *expected adaptation loss*,

$$B_{P_{\langle S, T \rangle}}(G_\rho) \stackrel{\text{def}}{=} R_{P_S}(G_\rho) + \text{dis}_\rho(D_S, D_T),$$

where $P_{\langle S, T \rangle}$ denotes the joint distribution over $P_S \times P_T$. The independence of each draw from P_S and P_T allows us to rewrite $B_{P_{\langle S, T \rangle}}$ as the expectation of the *domain adaptation loss* \mathcal{L}_{DA} ,

$$B_{P_{\langle S, T \rangle}}(G_\rho) = \mathbf{E}_{h_1, h_2 \sim \rho^2} \mathbf{E}_{(\mathbf{x}^s, y^s, \mathbf{x}^t) \sim P_{\langle S, T \rangle}} \mathcal{L}_{DA}(h_1, h_2, \mathbf{x}^s, y^s, \mathbf{x}^t), \quad (4)$$

$$\mathcal{L}_{DA}(h_1, h_2, \mathbf{x}^s, y^s, \mathbf{x}^t) \stackrel{\text{def}}{=} I(h_1(\mathbf{x}^s) \neq y^s) + I(h_1(\mathbf{x}^t) \neq h_2(\mathbf{x}^t)) - I(h_1(\mathbf{x}^s) \neq h_2(\mathbf{x}^s)).$$

Given $\langle S, T \rangle = \{(\mathbf{x}_i^s, y_i^s, \mathbf{x}_i^t)\}_{i=1}^m$, a sample of m source-target pairs drawn *i.i.d.* from $P_{\langle S, T \rangle}$, the *empirical adaptation loss* of G_ρ is $B_{\langle S, T \rangle}(G_\rho) = \mathbf{E}_{h_1, h_2 \sim \rho^2} \sum_{i=1}^m \mathcal{L}_{DA}(h_1, h_2, \mathbf{x}_i^s, y_i^s, \mathbf{x}_i^t)$.

PAC-Bayesian Bounds for Domain Adaptation. We restrict ourselves to the case exhibited by Equation (2) where \mathcal{H} is a set of linear classifiers, and posterior and prior distributions are Gaussians. First, we compute the expected adaptation loss $B_{P_{\langle S, T \rangle}}(G_{\rho_{\mathbf{w}}})$ of the Gibbs classifier $G_{\rho_{\mathbf{w}}}$ (remember that the posterior distribution is centered on the linear $h_{\mathbf{w}}$). With $\Phi_{\text{dis}}(a) \stackrel{\text{def}}{=} 2\Phi(a)\Phi(-a)$, we obtain,

$$B_{P_{\langle S, T \rangle}}(G_{\rho_{\mathbf{w}}}) = \mathbf{E}_{(\mathbf{x}^s, y^s, \mathbf{x}^t) \sim P_{\langle S, T \rangle}} \left[\Phi\left(y^s \frac{\mathbf{w} \cdot \mathbf{x}^s}{\|\mathbf{x}^s\|}\right) + \Phi_{\text{dis}}\left(\frac{\mathbf{w} \cdot \mathbf{x}^s}{\|\mathbf{x}^s\|}\right) - \Phi_{\text{dis}}\left(\frac{\mathbf{w} \cdot \mathbf{x}^t}{\|\mathbf{x}^t\|}\right) \right].$$

Now, we derive a new PAC-Bayesian theorem to bound the expected adaptation loss of linear classifiers. Theorem 2 is obtained by two key results. First, we use the specialization of the PAC-Bayesian theory to linear classifiers introduced by Corollary 1. Second, we need the methodology developed by [6, Theorem 5] to bound a loss relying on a pair of hypothesis $h_1, h_2 \sim \rho^2$ (like our domain adaptation loss of Equation (4)). We then obtain $\text{KL}(\rho_{\mathbf{w}}^2 \parallel \pi_0^2) = 2 \text{KL}(\rho_{\mathbf{w}} \parallel \pi_0) = \|\mathbf{w}\|^2$.

Theorem 2. *For any domain $P_{\langle S, T \rangle} \subseteq \mathbb{R}^d \times Y \times \mathbb{R}^d$ and any $\delta \in (0, 1]$, we have,*

$$\Pr_{\langle S, T \rangle \sim (P_{\langle S, T \rangle})^m} \left(\forall \mathbf{w} \in \mathbb{R}^d : \text{kl}\left(B_{\langle S, T \rangle}^* \parallel B_{P_{\langle S, T \rangle}}^*\right) \leq \frac{1}{m} \left[\|\mathbf{w}\|^2 + \ln \frac{\xi(m)}{\delta} \right] \right) \geq 1 - \delta,$$

where $B_{\langle S, T \rangle}^* \stackrel{\text{def}}{=} \frac{1}{2} B_{\langle S, T \rangle}(G_{\rho_{\mathbf{w}}}) + \frac{1}{4}$ and $B_{P_{\langle S, T \rangle}}^* \stackrel{\text{def}}{=} \frac{1}{2} B_{P_{\langle S, T \rangle}}(G_{\rho_{\mathbf{w}}}) + \frac{1}{4}$ ensure that the values provided to the $\text{kl}(\cdot \parallel \cdot)$ function are in interval $[0, 1]$.

Designing the Algorithm. The algorithm DA-PBGD, described here, minimizes the upper bound given by Theorem 2 by gradient descent. The corresponding objective function is,

$$\mathcal{B}(\langle S, T \rangle, \mathbf{w}, \delta) \stackrel{\text{def}}{=} \sup \left\{ \epsilon : \text{kl}(B_{\langle S, T \rangle}^* \parallel \epsilon) \leq \frac{1}{m} \left[\|\mathbf{w}\|^2 + \ln \frac{\xi(m)}{\delta} \right] \right\},$$

for a fixed value of δ . Consequently, our problem is to find weight vector \mathbf{w}^* that minimizes \mathcal{B} subject to the constraints $\mathcal{B} > B_{\langle S, T \rangle}^*$ and $\text{kl}(B_{\langle S, T \rangle}^* \parallel \mathcal{B}) = \frac{1}{m} [\|\mathbf{w}\|^2 + \ln \frac{\xi(m)}{\delta}]$. The gradient is obtained by computing the partial derivative of both sides of the latter equation with respect to w_j (the j^{th} component of \mathbf{w}). After solving for $\partial \mathcal{B} / \partial w_j$, we find that the gradient is,

$$\frac{\mathcal{B}(1-\mathcal{B})}{2m(\mathcal{B}-B_{\langle S, T \rangle}^*)} \left[4\mathbf{w} + \ln \left(\frac{\mathcal{B}(1-B_{\langle S, T \rangle}^*)}{B_{\langle S, T \rangle}^*(1-\mathcal{B})} \right) \sum_{i=1}^m \left[\Phi' \left(\frac{y_i^s \mathbf{w} \cdot \mathbf{x}_i^s}{\|\mathbf{x}_i^s\|} \right) \frac{y_i^s \mathbf{x}_i^s}{\|\mathbf{x}_i^s\|} + \Phi'_{\text{dis}} \left(\frac{\mathbf{w} \cdot \mathbf{x}_i^s}{\|\mathbf{x}_i^s\|} \right) \frac{\mathbf{x}_i^s}{\|\mathbf{x}_i^s\|} - \Phi'_{\text{dis}} \left(\frac{\mathbf{w} \cdot \mathbf{x}_i^t}{\|\mathbf{x}_i^t\|} \right) \frac{\mathbf{x}_i^t}{\|\mathbf{x}_i^t\|} \right] \right],$$

where $\Phi'(a)$ and $\Phi'_{\text{dis}}(a)$ denote respectively the derivatives of Φ and Φ_{dis} evaluated at a . The kernel trick applied to DA-PBGD allows us to work with dual weight vector $\boldsymbol{\alpha} \in \mathbb{R}^d$ that is a linear classifier in an augmented space. Given a kernel $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$, we have $h_{\mathbf{w}}(\mathbf{x}) = \sum_{i=1}^m \alpha_i k(\mathbf{x}_i, \mathbf{x})$.




Figure 1: Illustration of the decision of DA-PBGD on 4 rotations angles: From left to right 20°, 30°, 40°, 50°. In green and pink is the source sample, in grey is the target sample.

Experimental Results. Our DA-PBGD has been evaluated on a toy problem called inter-twinning moon and compared with: PBGD and SVM with no adaptation, the semi-supervised Transductive-SVM (TSVM) [7], the iterative DA algorithms DASVM [8] and the non-iterative version of DASF [9] based on the bound (1). We used a Gaussian kernel for all the methods. These preliminary results – illustrated on Tab. 1 and on Fig. 1 – are very promising. Moreover on Fig. 2, we clearly see the trade-off between the difficulty of the task and the minimization of the source risk in action: When the DA task is feasible DA-PBGD prefers to minimize the domain disagreement even if it implies an increase of the empirical source error, but when this minimization becomes hard, *i.e.* the complexity of the task is high, it prefers to focus only on the empirical source error.

Among all the possible exciting perspectives, we notably aim to theoretically define elegant and relevant assumptions allowing one to control the λ_ρ term of Eq. (3) to make our DA bound very tight.

Table 1: Average accuracy results for 4 rotation angles. DA-PBGD is more stable than the others and outperforms all the methods for 2 angles.

Rotation angle	20°	30°	40°	50°
PBGD	99.5	89.8	78.6	60
SVM	89.6	76	68.8	60
TSVM	100	78.9	74.6	70.9
DASVM	100	78.4	71.6	66.6
DASF	98	92	83	70
DA-PBGD	97.7	97.6	97.4	53.2




Figure 2: The trade-off between target and source errors according to the difficulty of the task (*i.e.* the rotation angle).

Acknowledgments This work was supported in part by the french project VideoSense ANR-09-CORD-026, in part by the IST Programme of the European Community, under the PASCAL2 Network of Excellence IST-2007-216886 and in part by NSERC discovery grant 262067. This publication only reflects authors’ views.

References

- [1] S. Ben-David and R. Uner. On the hardness of domain adaptation and the utility of unlabeled target samples. In *Proceedings of Algorithmic Learning Theory*, pages 139–153, 2012.
- [2] S. Ben-David, J. Blitzer, K. Crammer, A. Kulesza, F. Pereira, and J.W. Vaughan. A theory of learning from different domains. *Machine Learning Journal*, 79(1-2):151–175, 2010.
- [3] David A. McAllester. Some PAC-bayesian theorems. *Machine Learning*, 37:355–363, 1999.
- [4] P. Germain, A. Lacasse, F. Laviolette, and M. Marchand. PAC-Bayesian Learning of Linear Classifiers. In *Proceedings of ICML*, 2009.
- [5] J. Langford and J. Shawe-Taylor. PAC-bayes & margins. In *Advances in Neural Information Processing Systems 15*, pages 439–446. MIT Press, 2002.
- [6] Alexandre Lacasse, François Laviolette, Mario Marchand, Pascal Germain, and Nicolas Usunier. PAC-bayes bounds for the risk of the majority vote and the variance of the Gibbs classifier. In *NIPS*, 2007.
- [7] T. Joachims. Transductive inference for text classification using support vector machines. In *ICML*, 1999.
- [8] L. Bruzzone and M. Marconcini. Domain adaptation problems: A DASVM classification technique and a circular validation strategy. *IEEE Trans. Pattern Anal. Mach. Intell.*, 32(5), 2010.
- [9] E. Morvant, A. Habrard, and S. Ayache. Parsimonious Unsupervised and Semi-Supervised Domain Adaptation with Good Similarity Functions. *Knowledge and Information Systems*, 33(2):309–349, 2012.