The Influence Limiter: Provably Manipulation-Resistant Recommender Systems (Appendix)

Paul Resnick
University of Michigan
School of Information
presnick@umich.edu

Rahul Sami
University of Michigan
School of Information
rsami@umich.edu

August 3, 2007

A Appendix

This appendix should be read in conjunction with the article by Resnick and Sami [1]. Here, we include the proofs that were omitted from the main article due to shortage of space.

A.1 Lemma 5

Lemma 5: *For the quadratic scoring rule (MSE) loss, for all* $q,u \in [0,1]$, $GF(q||u) \geq \frac{D(q||u)}{2}$.

Proof of Lemma 5: Because both $D(q||u) = D(1-q||1-u)$ and $GF(q||u) = GF(1-q||1-u)$, we can assume $u \geq q$ without loss of generality. Keeping q fixed, we want to show that the result holds for all u. Note that $D(q||q) = GF(q||q) = 0$. Thus, differentiating with respect to u, it is sufficient to prove that $GF'(q||u) \geq D'(q||u)/2$ for all $u \geq q, u \leq 1$.

We change variables by setting $y = u - q$. We use the notation $D'(y)$ to denote $D'(q||u)|_{u=q+y}$, treating q as fixed and implicit. Likewise, we use the notation $GF'(y)$. For brevity, we use $\overline{q}$ to denote $(1-q)$.

$$\begin{aligned} D(q||u) &= q[(\overline{q}-y)^2 - \overline{q}^2] + \overline{q}[(q+y)^2 - q^2] \\ &= q[y^2 - 2y\overline{q}] + \overline{q}[y^2 + 2qy] \\ &= y^2 \\ \Rightarrow D'(y) &= 2y \end{aligned}$$

$$
\begin{aligned}
GF(q||u) &= q\log(1+y^2-2\overline{q}y)+\overline{q}\log(1+y^2+2qy) \\
\Rightarrow GF'(y) &= q\frac{2y-2\overline{q}}{1+y^2-2\overline{q}y}+\overline{q}\frac{2y+2q}{1+y^2+2qy} \\
&= 2y\left[\frac{q}{1+y^2-2\overline{q}y}+\frac{\overline{q}}{1+y^2+2qy}\right]-2q\overline{q}\left[\frac{1}{1+y^2-2\overline{q}y}-\frac{1}{1+y^2+2qy}\right] \\
&= 2y\frac{1+y^2}{(1+y^2-2\overline{q}y)(1+y^2+2qy)}-2q\overline{q}\frac{2y}{(1+y^2-2\overline{q}y)(1+y^2+2qy)} \\
&= 2y\left[\frac{1+y^2-2q\overline{q}}{(1+y^2-2\overline{q}y)(1+y^2+2qy)}\right]
\end{aligned}
$$

To complete the proof, we show that the term within the brackets is at least $1/2$, by showing that $2*\textit{numerator}-\textit{denominator}\geq 0$:

$$
\begin{aligned}
2(1+y^2-2q\overline{q}) &- (1+y^2-2\overline{q}y)(1+y^2+2\overline{q}y) \\
&= 2+2y^2-4q\overline{q}-(1+y^2)^2+4q\overline{q}y^2 \\
&= 1-y^4-4q\overline{q}(1-y^2) \\
&= (1+y^2-4q\overline{q})(1-y^2)\geq 0
\end{aligned}
$$

The last inequality follows because $y^2\leq 1$ and $4q\overline{q}\leq 1$. This shows that $GF'(y)\geq 0.5D'(y)$, and thus, $GF(y)\geq D(y)/2$ for all y. □

Theorem 7 also requires a slightly different version of Lemma 5, in order to show that $\log R_j$ does not decrease in expectation even when $R_j=t>1$. For this, we define

$$
\begin{aligned}
GF_t(q||u) \stackrel{def}{=} &[q\log(t+L(HI,u)-L(HI,q))+ \\
&(1-q)\log(t+L(LO,u)-L(LO,q))]-\log t
\end{aligned}
$$

Note that $GF(q||u)=GF_1(q||u)$.

Lemma 5b: *For the quadratic loss function, for all $q,u\in[0,1]$, and for any $t\geq 1$, $GF_t(q||u)\geq 0$.*

Proof of Lemma 5b: The proof follows the proof of Lemma 5, with $t>1$ instead of 1; it is easy to see that $GF_t'(y)\geq 0$, which is sufficient to show that $GF_t(q||u)$ is nonnegative, because $GF_t(q,q)=0$. □

A.2 Lemma 6

Lemma 6: *Suppose $\hat{\pi}_j$ and $\hat{\pi}_{j-1}$ are two partitions such that $\hat{\pi}_j$ is a refinement of $\hat{\pi}_{j-1}$. For each state ω, let $\mathbf{q}_j(\omega)=E(l(\omega)|\hat{\pi}_j)$ be the optimal prediction function given partition $\hat{\pi}_j$.*

Proof of Lemma 6: Let $\mathbf{u}(\omega)$ be *any* function that is constant on each component of $\hat{\pi}_{j-1}$. Then, $EGF(\mathbf{q}_j||\mathbf{u})\geq I(\hat{\pi}_j||\hat{\pi}_{j-1})/2$ in the quadratic loss model.

Proof of Lemma 6: For simplicity, we prove the result assuming for a single component s_{j-1} of $\hat{\pi}_{j-1}$; the stated result follows easily by averaging over all such components. Let $u = \mathbf{u}(s_{j-1})$ be the value of the function u on s_{j-1}. Let $q_{j-1} = E(l|\hat{\pi}_{j-1}$ be the optimal value of u on s_{j-1}.

Corresponding to s_{j-1}, there is some set S of components of $\hat{\pi}_j$, such that $\cup_{s\in S}\, s = s_{j-1}$. Let q_s be the value of function $\mathbf{q}_j$ on component $s \in S$. Further, let p_s be the probability of component s, so that $\sum_{s\in S} p_s = 1$. Now,

$$\begin{aligned} EGF(\hat{\pi}||u) &= \sum_{s\in S} p_s GF(q_s||u) \\ &\geq (1/2)\sum_s p_s D(q_s||u) \ \text{(from Lemma 5)} \\ &\geq (1/2)\sum_s p_s D(q_s||q_{j-1}) \end{aligned}$$

The last inequality follows from the fact that q_{j-1} is the optimal prediction given information $\hat{\pi}_{j-1}$, *i.e.*, it minimizes expected loss among all functions $\mathbf{u}$ that are constant on components of $\hat{\pi}_{j-1}$. In other words, the expected improvement is smallest when the best prior prediction is made. If u is not optimal, $D(q_s||u)$ can only go up in expectation (over components s). □

A.3 Theorem 7

Theorem 7: *Suppose rater j has rated m items, and suppose the informativeness of rater j is $I(q_j||q_{j-1}) = h$. Then, for all $m \geq (2\lambda+1)/h$, rater i's expected reputation (with the quadratic scoring rule) is bounded below by*

$$E(R_j) \geq mh - 2\lambda - 2\log(mh - 2\lambda)$$

Proof of Theorem 7: The reputation of an informative rater j grows in two phases: While the reputation is low, it tends to grow exponentially as j's influence β_j also grows. Once $\beta_j = 1$ is reached, j's reputation grows linearly. The intuition provided in the main paper is that it takes $2\lambda/h$ rounds to reach full credibility, after which h bits are gained in each round. Unfortunately, the proof gets messy because, after reaching a reputation of 1, it is possible to temporarily lose sufficient reputation that the influence limits apply again.

An alternative approach is to work with a transformed reputation in which the two phases are balanced. To this end, we define a function $G : \Re^+ \to \Re$:

$$G(x) \quad = x + 2\log x$$

We observe that $G(x)$ is increasing, invertible, and concave. Also, note that $G(e^{-\lambda}) = -2\lambda + e^{-\lambda}$, and $G(1) = 1$.

Now, let $\overline{G}^{(i)}$ denote the expected value of $G(R_j)$ after items $1, 2, \cdots, i$. $\overline{G}^{(0)} = G(e^{-\lambda}) > -2\lambda$.

Claim: $\overline{G}^{(i+1)} \geq \overline{G}^{(i)} + h$

Proof of claim: Consider the distribution of $G(R_j)$ after i rounds. Further, consider any possible value g in this distribution. We show that, conditioning on $G(R_j) = g$ after round i, the expected value of G after round $i+1$ is at least $g+h$. Note that $E(G(R_j)) = E(R_j) + 2E(\log R_j)$. When $R_j > 1$, we will show that

$E(R_j)$ increases by h in expectation. When $R_j \leq 1$, we will show that $2\log R_j$ increases by h in expectation. In either case, the other term does not decrease in expectation. *Case (i):* $g < 1$. In this case, $R_j < 1$ after the ith round, so j is influence limited. Consider the logarithmic term $2\log x$ in $G(x)$. The raters before j have combined information represented by the partition $\hat{\pi}_{j-1}$; thus, none of them could have distinguished between two states in the same component of $\hat{\pi}_{j-1}$. Thus, the rating on item i just before j's is a function that is constant on components of $\hat{\pi}_{j-1}$. The influence-limited rating $\tilde{q}_{j-1}$ is therefore also a function that is constant on $\hat{\pi}_{j-1}$. Thus, by Lemma 6, the expected value of the $\log R_j$ term after the $i+1$st round is at least $h/2$ higher than it was at the start of the $i+1$st round. The linear term also increases in expectation (possibly by a very small amount), and so the expectation of G after $i+1$ rounds must be at least $g+h$.

Case (ii): $g \geq 1$ In this case, j has full credibility $\beta_j = 1$, and we show that the linear term in G increases by at least h in expectation. By the same argument as in case (i), the previous rating $\tilde{q}_{j-1}$ is a function that is constant on $\hat{\pi}_{j-1}$. Thus, by definition of the informativeness $I(\hat{\pi}_j||\hat{\pi}_{j-1})$, the expected value of the linear term increases by at least h. By lemma 5b, with $t = R_j \geq 1$, the expected value of the logarithmic term does not reduce. Thus, in this case too, the expected value of G after $i+1$ rounds must be at least $g+h$.

As this is true conditioned on any value of g, it must be true in expectation. Thus, $\overline{G}^{(i+1)} \geq \overline{G}^{(i)} + h$, and hence, $\overline{G}^{(m)} \geq \overline{G}^{(0)} + mh > mh - 2\lambda$. Because of the concavity of the function G, Jensen's inequality implies that

$$G(E(R_j)) \geq E(G(R_j)) = \overline{G}^{m} > mh - 2\lambda$$

For $m \geq (2\lambda+1)/h$, $mh - 2\lambda \geq 1$. Thus, $\log(mh-2\lambda) \geq 0$. Setting $r = mh - 2\lambda - 2\log(mh-2\lambda)$, we see that $G(r) = r + 2\log r \leq r + 2\log(mh-2\lambda)$. Thus, $G(r) \leq mh - 2\lambda$, which means that $G(E(R_j)) > G(r)$. As G is an increasing function, it follows that $E(R_j) \geq r = mh - 2\lambda - 2\log(mh-2\lambda)$. □

Remark: The last term in the bound, $-2\log(mh-2\lambda)$, appears to be loose; it should be possible to tighten it to a small constant.

References

[1] P. Resnick and R. Sami. The influence limiter: Provably manipulation-resistant recommender systems. In *To appear in Proceedings of the ACM Recommender Systems Conference (RecSys07)*, 2007.